The AI Entrepreneur

Innovating with Artificial Intelligence

Luca Grilli

The AI Entrepreneur

Artificial intelligence (AI) is often portrayed as a force set to radically transform entrepreneurship. This book takes a more analytical approach, asking what AI truly changes – and what it does not – about the role of innovative entrepreneurs in advanced capitalist economies. Integrating AI into established economic theories of entrepreneurship, Luca Grilli develops new conceptual frameworks informed by emerging empirical evidence on the nature of AI entrepreneurship. His analysis shows how AI frequently reinforces incumbent advantage while also generating forms of systemic lock-in around large, AI-driven ecosystems. These dynamics risk narrowing the space for genuinely innovative ideas, thereby reshaping the conditions under which entrepreneurship can thrive. Against this backdrop, *The AI Entrepreneur* reflects on how institutions and economic policy can safeguard space for entrepreneurial agency, preventing the AI entrepreneur from becoming a postmodern simulacrum confined within increasingly 'fenced' forms of capitalism.

LUCA GRILLI is Full Professor of Business and Industrial Economics at Politecnico di Milano. He is among the most cited social scientists in innovative entrepreneurship and has advised government in this field. This book builds on more than two decades of scholarly research with policy-relevant implications.

The AI Entrepreneur

Innovating with Artificial Intelligence

LUCA GRILLI
Politecnico di Milano

CAMBRIDGE
UNIVERSITY PRESS



Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre,
New Delhi – 110025, India

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781009621267
DOI: 10.1017/9781009621274

© Luca Grilli 2026

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press & Assessment.

When citing this work, please include a reference to the DOI 10.1017/9781009621274

First published 2026

Cover image: © Josh Hawley / Moment / Getty Images.

A catalogue record for this publication is available from the British Library

A Cataloging-in-Publication data record for this book is available from the Library of Congress

ISBN 978-1-009-62126-7 Hardback
ISBN 978-1-009-62125-0 Paperback

Cambridge University Press & Assessment has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

For EU product safety concerns, contact us at Calle de José Abascal, 56, 1°, 28003 Madrid, Spain, or email eugpsr@cambridge.org.

To
Evelina,
the least artificial and
the most entrepreneurial
human being
I have ever known

Contents

List of Figures *page* viii
List of Tables ix
List of Boxes x
Preface xi
Acknowledgements xiii

Introduction 1

Part I Basics

1 Capitalism and Innovative Entrepreneurship 11
2 Artificial Intelligence: Setting the Scene 36

Part II The Encounter

3 Artificial Intelligence and the Entrepreneurial Idea 73
4 Can Artificial Intelligence Be Entrepreneurial? 98
5 Artificial Intelligence and Success in Innovative Entrepreneurship 124
6 Phenomenology of the Start-Up in the AI Era 158

Part III Policy Challenges

7 The Perils for the AI Entrepreneur 189
8 AI Entrepreneurship Policy and Institutional Reforms 213

Conclusion 243

References 245
Index 282

Figures

2.1 Deep learning in context: a socio-technical representation *page* 45
2.2 AI supply chains 61
3.1 The 'SoleScribe Snack Shoe' 92
3.2 Entrepreneurial ideas and AI 97
4.1 Entrepreneurial action and AI 102
4.2 The 'AI robotic' version of Michelangelo's *The Creation of Adam* 114
4.3 The mind and the body of the AI entrepreneur: forces shaping the action space 123
5.1 AI and start-ups' human partners to manage 146
6.1 AI and start-ups: a taxonomy 160
6.2 '(Golden) talent trap' and 'adverse selection' in advanced AI-based entrepreneurship 175
6.3 Evolution of media and science narratives on start-up success (1997–2025): (a) news media and (b) science 184
7.1 The 'Systemic Lock-in Model' of AI-based entrepreneurship 191
7.2 The AI entrepreneur and the risk of entrapment 197
8.1 The vicious-cycle risk of postmodern entrepreneurship 218
8.2 Comparison EU vs. US on the number of AI-based start-ups created (2005–2023): (a) sample (year precision) and (b) subsample (month-year precision) 225

Tables

5.1 Role of founders' human capital in the AI-related innovative entrepreneurship space *page* 142
8.1 Institutional reforms and AI entrepreneurship: a policy framework 228

Boxes

1.1 How VC works *page* 33
2.1 AI breakthroughs: the 'duo' of academy and industry in AI science 39
2.2 On the elusive taxonomies of AI 42
3.1 Entry deterrence of predictive AI 89
3.2 In search of an entrepreneurial idea with ChatGPT 91
4.1 Implementing the entrepreneurial idea with AI 121
5.1 Deep learning and founders' human capital 131
5.2 From tools to orchestration: making the 'SoleScribe Snack Shoe' stack work 136
5.3 Deep learning and VC 156
6.1 From AI scientific papers to AI-based start-ups 171
6.2 MAMAA as VCs and acquirers of AI-based innovative start-ups 181
6.3 What has 'success' in entrepreneurship become? Evidence from press and science 183
7.1 AI 'killer' acquisitions 196
7.2 AI freelancers in a crowdwork platform: the case of Upwork 200
8.1 Where is entrepreneurship in AI regulation(s)? 222
8.2 The 'unintended consequences' of the EU AI Act on European AI-based start-ups 224

Preface

Artificial intelligence is literally on everyone's lips. There is no topic, even seemingly the most unrelated, that doesn't have a good chance these days of being magically drawn to the magnet of artificial intelligence. Its acronym, AI, is now paired with an almost endless sequence of other terms, from healthcare and education to entertainment, sports, cooking, housing, military spending, security, and beyond. AI is increasingly regarded as a potential tool for addressing some of the major challenges the world is facing – from climate change to poverty – although, as this book will argue, its actual contribution critically depends on the institutional context in which it is deployed.

Its pervasiveness manifests in the most diverse contexts: from simple conversations overheard on the street to newspaper articles and from social media posts to videos that quickly go viral. Clearly, the wave has also swept through all scientific fields. I have often questioned myself (especially my snobbish spirit) whether it was worth joining the caravan with my own wagon (or rather, the wagon I know best): (innovative) entrepreneurship. In my (partial) defence, it should be noted that my personal interest in AI dates back several years before the incredible hype triggered by the release of ChatGPT in November 2022. This interest was fuelled both by cinematic imagination (e.g. *Her* from 2013) and documentary reality (e.g. *AlphaGo* from 2017, which I will discuss in Chapter 3), at a time when broader public attention towards AI was already rising but still far from the current frenzy.

In the end, three main reasons led me to take the plunge. The first is *technological* and delves into the conviction that AI (which, for the sake of clarity, is synonymous with 'deep learning' in this work) is an incredibly powerful technology, capable of transforming many aspects of our lives and daily routines, including the economic sphere. But, after an initial fascination and a consequent deeper technical exploration of how it works, I realize that this immensely powerful AI is

far from omnipotent. Thus, there is a need to demystify AI, cutting through the hype to better understand its immense potential while avoiding disappointment from unrealistic expectations. The second reason is more *economic* in nature. I am firmly convinced (in excellent company) that entrepreneurship is the primary driver of the capitalist system. Consequently, I quickly sensed that, perhaps more than other technological revolutions, AI raises profound questions about the nature and defining traits of the entrepreneurial process. For this reason, I felt it was worth studying the complex relationship between AI and entrepreneurship, also to better understand the capitalist system that lies ahead. This book was born from this purely exploratory drive. Clearly, the task seems (and probably is) daunting, largely because we find ourselves in the challenging position of 'analyzing a storm while still in the eye of it'. Yet this is precisely where the third, more *institutional*, reason comes in: the conviction that appropriate policies can allow us – even in the age of AI, or perhaps precisely because of AI – to preserve a form of entrepreneurial capitalism that remains dynamic, creative, and, why not, also oriented towards addressing societal grand challenges.

In this book, I offer my own vision – as scientifically grounded as possible – on these themes. Although this is a book about AI, I chose to write it in a somewhat traditional way. AI tools were used, unless otherwise specified, mainly for language refinement, the formulation of certain examples (which I always cross-checked), and occasionally to test ideas; yet every argument and interpretation (whether right or wrong) ultimately originates from me. In short, if the future of writing is moving towards a form of 'distant writing', defined more by the ability to pose questions to AI systems than to provide answers, what I present here is best described as an exercise in 'close writing', guided by my own sustained reflection on the themes at hand.

Acknowledgements

I wish to express my sincere gratitude to Thomas Åstebro, Stefano H. Baruffaldi, Vincenzo Butticè, Sergio Mariotti, and Mattia Pedota for their valuable and detailed comments on draft versions of this book, which have, at least from my perspective, greatly contributed to improving it. I am also grateful to Simone Scardapane for his careful reading and thorough review of earlier versions of the more technical chapter, as well as for his helpful suggestions for improvement. I equally wish to thank Walter Quattrociocchi for his attentive reading of this chapter. Alessandra Colombelli, Tommaso Minola, and John P. Walsh are also acknowledged for closely accompanying the development of this work and for the valuable exchanges throughout the process. I am also grateful to my departmental colleagues, within the context of the HumanTech Project under the 'Departments of Excellence' programme, for their continued goodwill while writing this book, and to those colleagues – fortunately few – whose collaboration with me was delayed during its completion. Great research assistance made it possible to conduct many empirical analyses, most of which are included in the 'Boxes'. I wish to thank the following for their contribution in this respect: Donatello Larocca, Miriam I. Marchesani, Matteo Napodano, Lorenzo Rapini, Giacomo Seratoni Gualdoni, and Gianluca Tartaglia. I am deeply grateful to Marco Mari for his final reading of the book. Any remaining errors or inaccuracies are, of course, entirely my own responsibility. Finally, this (intensive) research endeavour would not have been possible without the unwavering support of my family. My deepest thanks go to my wife, Silvia, and to my daughter, Sofia, for their patience, understanding, and constant encouragement. I also wish to thank my parents, Paolo Ilio and Anna Maria, and my sister, Elisa, for their unfailing commitment over the years.

Introduction

Can artificial intelligence (AI) truly generate new entrepreneurial ideas – and, perhaps more importantly, help turn them into reality? How will entrepreneurial finance evolve in this new era, and what will become of the venture capital (VC) industry as AI tools increasingly permeate investment decisions? At the same time, will AI redefine the importance of founders' human and social capital for the success of innovative ventures? What kinds of start-ups are now emerging as a result of AI, and will any of them genuinely challenge established industrial paradigms or threaten dominant positions, including those of Big Tech firms? More broadly, how will capitalism itself, and the competitive dynamics that sustain it, be reshaped by AI-driven transformations in the entrepreneurial process? And what role should institutions and policy-making play in guiding this transition, while also addressing broader societal grand challenges?

These are the questions that frame the discussion developed in the following chapters. In a nutshell, this book seeks to examine the entrepreneurial phenomenon (and innovative entrepreneurship in particular) through the lens of AI, with the aim of deriving broader reflections on the evolution of modern advanced capitalist economies. In doing so, our analytical perspective is deliberately multidimensional: it looks simultaneously at the micro-foundations of entrepreneurship (the individual, the team, and the firm), the meso-level of industrial and technological ecosystems, and the macro and institutional settings that shape incentives and constraints for entrepreneurial activity. This multidimensional approach (which, admittedly, characterizes much of the research in entrepreneurship) stems from the conviction that perhaps more than ever, understanding AI's impact on entrepreneurship – and, in turn, on modern advanced capitalism – requires crossing disciplinary boundaries. It demands the combination of insights from economics, psychology, management,

sociology, and public policy, in order to capture the complex interplay between technological change, entrepreneurial dynamics, market outcomes, and institutional evolution.

Recognizing the complexity of this task, we are aware that such an analysis must, on the one hand, already strive to offer some normative guidance and, on the other, will inevitably be superseded over time. Yet it remains essential that it be undertaken now, while the contours of this transformation are still being drawn. This awareness also connects to a broader and increasingly shared view – both in the media and in academia – that public policy (e.g. antitrust, regulation) must now, and before it is too late, play a fundamental role in 'steering the ship through the storm to the harbour' of social welfare (e.g. Hauser et al., 2025). And if this is true, we are deeply convinced that this debate cannot overlook the impact that AI may have on *innovative entrepreneurship*. Thus, fully aware of the challenges ahead, we now embark on this exploration in the following chapters – not blind to its limitations, yet reassured by the words of Max Weber, who reminds us of the inevitability that every scientific endeavour of today will be surpassed tomorrow. And yet, he also underscores the necessity of undertaking it nonetheless – not only for the practical implications it may yield today but, more simply, because *it must be done*:

> In science, each of us knows that what he has accomplished will be antiquated […]; it is the very meaning of scientific work, to which it is devoted in a quite specific sense, as compared with other spheres of culture for which in general the same holds. […] Whoever wishes to serve science has to resign himself to this fact. Scientific works certainly can last as 'gratifications' because of their artistic quality, or they may remain important as a means of training. Yet they will be surpassed scientifically – let that be repeated – for it is our common fate and, more, our common goal. We cannot work without hoping that others will advance further than we have. In principle, this progress goes on ad infinitum. And with this, we come to enquire into the meaning of science. For, after all, it is not self-evident that something subordinate to such a law is sensible and meaningful in itself. Why does one engage in doing something that in reality never comes, and never can come, to an end?
>
> One does it, first, for purely practical, in the broader sense of the word, for technical, purposes: in order to be able to orient our practical activities to the expectations that scientific experience places at our disposal. Good. Yet this has meaning only to practitioners. What is the attitude of the

academic man towards his vocation – that is, if he is at all in quest of such a personal attitude? He maintains that he engages in 'science for science's sake' […]. (Weber, 1922/1946, p. 138)

Analytical Orientation

This book seeks, as far as possible, to maintain analytical balance and critical distance in examining the various topics addressed. When evaluative judgements are advanced and may appear to take a particular stance, they are always supported by scientific evidence and grounded in the relevant literature (as is to be expected in any scientific endeavour). Particular care is devoted to avoiding the emphatic or dramatic tones often found in popular narratives, even though such accounts may at times offer valuable insights (e.g. Harari, 2024). While the analysis deliberately refrains from teleological or sensationalist framings, AI is nonetheless regarded as a genuine 'revolution' – arguably the only concession to the tradition of 'techno-optimism', as critically discussed by Vaclav Smil (2023). The fact that AI may also be subject to certain shadows – from the potential burst of stock market bubbles to gains that are not immediately significant in its use – only serves to make it even more akin to previous technological revolutions, such as the Internet and, before that, computerization. This, of course, does not mean assuming from the outset that AI will reshape every socio-economic dynamic.

As the subsequent chapters illustrate, the technological revolution brought about by AI is not thought to transform all dimensions of innovative entrepreneurship in a deterministic or uniform manner, nor to produce an inevitable rupture with past patterns of entrepreneurial activity. Rather, this work can be viewed – in the spirit of what Obschonka et al. (2024) describe as 'prospecting' – as an attempt to discern what is changing and what remains stable in the entrepreneurial process as a result of AI, and consequently, to consider which forms of capitalism may lie ahead. In doing so, it also reflects on the forms of public and entrepreneurship-oriented policy that may be required to preserve, or, where necessary, to restore, those institutional configurations most conducive to the 'productive' forms of capitalism envisioned by Baumol et al. (2007), in which large firms and new innovative start-ups coexist, and the role of the latter remains essential for the collective good.

Synopsis

The book is organized into three parts.

The first, entitled 'Basics', establishes the two foundational axes of the analysis: on the one hand, innovative entrepreneurship and its role within the capitalist system (Chapter 1); on the other, artificial intelligence as a transformative technological domain (Chapter 2).

The second part, 'The Encounter', explores the relationship between these two axes. It examines whether, to what extent, and in which dimensions the advent of AI alters – or leaves unchanged – the processes underpinning innovative entrepreneurship and the broader dynamics of start-up formation and success (Chapters 3–6).

Finally, the third part, 'Policy Challenges', addresses the potential risks that AI may introduce for innovative entrepreneurship, particularly from the standpoint of social welfare (Chapter 7). Building on this analysis, it develops a set of reflections and policy recommendations aimed at preserving and revitalizing the role of productive, innovation-driven entrepreneurship in modern economies, while steering it towards the pursuit of genuinely meaningful societal objectives (Chapter 8).

The book concludes with some final reflections. What follows is a brief overview of the content – and, in some cases, the necessary premises – of each of the eight chapters.

In Chapter 1, we will start with a portrait of the salient features of innovative entrepreneurship, and its role in advanced economies. Apart from some necessary introductory remarks, this chapter is very much 'AI-free', serving as a (brief) overview of the state of the art regarding the main micro-founded drivers underlying the phenomenon of innovative entrepreneurship. It focuses on what we consider (so far) to be the three pillars for innovative start-ups to succeed in the market: the human capital of their founders, their social capital, and financial capital, with a specific emphasis on venture capital. In this analysis, we implicitly distance ourselves from a perspective that, although present in the literature (see, e.g., the 'gambler's ruin' theory applied to entrepreneurship, e.g. Coad et al., 2013) views entrepreneurial success as largely, if not entirely, governed by random processes. With this, we certainly do not intend to claim that the process by which an innovative start-up emerges in the economic sphere is deterministic, nor that possessing certain characteristics necessarily

guarantees entrepreneurial success, regardless of how success is defined. Predicting with certainty ex-ante whether a specific start-up will survive, grow, and ultimately succeed remains – and, let us state this upfront, will continue to remain – simply impossible, even in the age of AI. Needless to say, we are also convinced that, as with everything in life, luck plays an important role.

However, we are equally convinced, to paraphrase the famous entrepreneur Peter Thiel in his book *Zero to One* (co-written with Blake Masters, 2014), that entrepreneurs are not simply 'lottery tickets'. Certainly, luck matters, but only to a certain extent, as evidenced by the existence of numerous successful serial founders. Or, if we prefer to stay within the metaphor, we believe (on the basis of a large body of scientific evidence produced throughout several decades) that certain characteristics – particularly the three pillars mentioned earlier – certainly allow entrepreneurs (on average) to acquire a significantly larger number of lottery tickets to succeed in the market. In this sense, it resonates with a broader belief embedded in capitalist ideology, namely that the capitalist system is perceived by its supporters as fair only insofar as (emphasis mine):

> [W]hoever wins does so because of a combination of *good skill*, *hard work* and *good luck*. These are the attributes which are [or better to say 'should', author's note] rewarded under capitalism – not class privilege, for example. (Bowles, 2013, p. 24)

Having established the economic foundations of our discourse, Chapter 2 lays the groundwork for AI. If Chapter 1 is, to a great extent, intentionally 'AI-free', this chapter, conversely, is deliberately rather 'entrepreneurship-free'. The aim is to analyze the technological aspects in detail, paying particular attention to their socio-economic features and implications, but without making any specific reference to the entrepreneurial dynamics that may arise from them. After a brief introductory overview of AI, it seeks to produce a representation of 'deep learning' – showing how these technologies work, what their potential (and their limitations) are, so as to also encompass the foreseeable further developments of the technology. Following this technological but also socio-economic analysis, the chapter focuses on the 'new capabilities' of this digital 'statistical machine' that is AI (i.e. predictive, generative, agentic), and the ways it is deployed within the economic system.

It is from Chapter 3 onwards that we begin to blend the two fundamental constructs of this book – entrepreneurship and AI – with the aim of investigating whether, and how, the technological revolution brought about by AI may affect innovative entrepreneurship and, by extension, modern advanced economies. Chapter 3 explores how AI technologies influence the identification and creation of innovative entrepreneurial ideas. It highlights that, in principle, creative potential is not limited to generative AI; predictive AI systems can also display creativity by breaking established patterns and revealing non-obvious business ideas. The chapter distinguishes between 'Kirznerian' (entrepreneurial) ideas, which originate from market inefficiencies and unmet consumer needs and are likely to shrink as incumbents gain a predictive advantage, and 'Schumpeterian' ideas, which are more radical and disruptive and may see their space expanded through the generative capabilities of AI. While AI may raise the average level of creativity and potentially reduce its diversity, human divergent thinking remains relevant, pointing to the importance of hybrid or symbiotic intelligence.

Chapter 4 revolves around a fundamental question: Can AI be entrepreneurial, or rather, can it assist in the implementation stage of an entrepreneurial idea? In addressing this question, the chapter distinguishes between epistemic and ontological uncertainty, arguing that while AI can reduce the former by leveraging data and predictive models, it cannot eliminate the latter, which preserves – especially when combined with the generative capabilities highlighted in Chapter 3 – a vital space for human entrepreneurship. The discussion also highlights psychological challenges, such as the 'Not Thought Here' syndrome, in which entrepreneurs may feel less ownership over AI-generated ideas, potentially undermining motivation and leadership. It also underscores the enduring importance of human qualities and the embodied, social, and persuasive dimensions of entrepreneurship, drawing analogies to theatre rather than cinema. While agentic AI can automate and support many entrepreneurial tasks, human agency remains essential – particularly in innovative contexts – because of the irreducible elements of Knightian uncertainty, trust-building, and authentic human interaction.

Building on the findings of the previous two chapters, Chapter 5 examines how AI may reshape the three pillars of success in innovative entrepreneurship: human, social, and financial capital. It argues that while AI lowers entry barriers and substitutes for some skills, specific knowledge and the ability to orchestrate human and technological resources

remain vital, captured by the new competence of *dynamic AI-enabled orchestration capability*. Social capital dynamics are also changing: AI tends to reduce team and network size but increases the importance of legitimacy and reputation. Financial capital, particularly VC, continues to play a central role, with AI expanding access to investors and information but not replacing the core functions of selection and coaching.

Chapter 6 explores the evolving landscape of start-ups in the AI era, proposing a taxonomy based on whether AI serves as a core offering or an enabling tool, and distinguishing between predictive and generative capabilities. It identifies four archetypes – predictive analytics ventures, creativity-centric generative AI start-ups, business intelligence-enhanced and creativity-augmented AI-served start-ups – while simultaneously highlighting how the boundaries between these archetypes are becoming increasingly blurred, especially between the first two, as firms combine multiple capabilities, including agentic ones. The chapter emphasizes how incumbents benefit from proprietary data in predictive domains while generative AI tends to lower entry barriers. At the end, three stylized facts emerge: the tight link between AI science and entrepreneurship, the 'talent trap' drawing top engineers to large high-tech firms, and the pervasive influence of MAMAA (i.e. Microsoft – Apple – Meta – Alphabet – Amazon) and other 'AI giants' over innovative AI-based start-ups.[1] The chapter concludes by suggesting a possible cultural shift in which entrepreneurial success is increasingly measured by acquisition rather than independent growth, with significant implications for innovation and market dynamics in the AI economy.

The analysis developed so far provides the foundation for Chapter 7, which starts from the premise that, from a purely technological standpoint, AI transforms but does not undermine the entrepreneur's

[1] The five United States (US) firms are often referred to using the alternative label GAFAM, derived from the original names of Google and Facebook. We prefer the acronym MAMAA, which we consider better suited to capturing the relevant boundaries of the group. In line with other contributions (e.g. Jacobides et al., 2021; Lundvall and Rikap, 2022; Srnicek, 2025), we employ the expression 'AI giants' in an intentionally broad and evocative manner, acknowledging that its boundaries are neither precisely defined nor stable over time. The term is used to refer to those Big Tech firms that, as will be examined in greater detail, play a leading role in AI. At present, this group (narrowly defined) certainly includes the MAMAA companies as its core protagonists, but also firms such as IBM, NVIDIA, Oracle and the three major Chinese players commonly referred to by the acronym BAT (Baidu, Alibaba, and Tencent), as well as other Chinese heavyweights such as ByteDance and Huawei.

fundamental role within capitalism. The real risks lie elsewhere and have a more human component. They stem from the dominance of large AI giants whose systemic lock-in – spanning talent, science, venture capital, acquisitions, and infrastructure – constrains competition and may limit start-ups' independent success. Additional pressures arise from macroeconomic and political factors, and sociological dynamics such as algorithmic polarization and '(socially) thin innovation', where novelty lacks genuine social value. Together, these trends risk fostering multiple 'postmodern' forms of entrepreneurship confined within the boundaries of dominant platforms, giving rise to a kind of 'fenced capitalism' in which the space for truly disruptive and innovative entrepreneurship is, in practice, increasingly entrapped.

Chapter 8 revisits the classical 'market failure' rationale for public intervention for innovative entrepreneurship, arguing that AI does not overturn it. The analysis advocates integrating AI tools into entrepreneurship policy-making – mirroring practices already widespread in the private sector – while warning against simplistic entry subsidies and calling for more comprehensive support mechanisms. The discussion also addresses the unintended consequences of broad regulatory interventions, which may disproportionately burden start-ups, and highlights the need to avoid the weaponization of industrial policy in the field of AI, as this can distort the playing field for innovative entrepreneurship. Starting from the premise that the risks identified in Chapter 7 – *incumbent advantage*, *systemic lock-in*, *labour market turmoil*, and *postmodern entrepreneurship* – are both real and significant, the chapter outlines four policy propositions aimed at ensuring that AI acts as an enabler rather than a constraint for innovative entrepreneurship.

A brief concluding chapter synthesizes the main arguments of the book and reflects on the persistence of innovative entrepreneurship in the AI era, as well as on the institutional conditions that shape its broader economic and social role.

We now embark on this journey. What follows is an endeavour to understand how a general-purpose technology of (we believe) unprecedented scope is beginning to reshape – sometimes subtly, sometimes more profoundly – the drivers, the constraints, and the possibilities of innovative entrepreneurship, and how these shifts extend to the broader trajectories of modern advanced capitalist economies.

PART I

Basics

1 *Capitalism and Innovative Entrepreneurship*

Capitalism and the 'AI Revolution'

Capitalism is an ambiguous term, where definitions also change and evolve throughout time (e.g. see Kocka, 2016 and Levy, 2017, among others). Broadly speaking, trying to focus on its usually agreed fundamentals, we can say that this term typically identifies an economic system based on the private ownership of the means of production for profit purposes and the use of markets (and prices) to allocate scarce goods and resources (Bowles, 2013). But apart from these two traits, which could be considered the common denominator of most modern advanced economies (with the possible sole exception of China), it is challenging to provide a single profile of capitalism and, in fact, many refer to different types (Baumol et al., 2007; Bowles, 2013) or even varieties (Hall and Soskice, 2001) of the capitalist system. Even the place and date of birth are of course uncertain: some take the High Middle Ages in Italy as a reference, others the long-duration expeditions to India in the early 1600s, and still others make the period coincide with the beginning of the first industrial revolution in England in the late eighteenth century (see Fulcher, 2004 and Kocka, 2016 for a brief review). Thus, if we take the latter view, which is probably the dominant one, then modern capitalism emerged from a revolution and has certainly gone through several phases and sometimes dramatic and abrupt changes and even crises. Although the term 'revolution' is often overused, it is widely agreed that *artificial intelligence*, along with its integration with other technologies (e.g. the Internet of Things (IoT), blockchain, 3D printing, and robotics) has a disruptive power in the way we (will) live and do business. Indeed, AI itself is an 'umbrella' term, which involves a variety of techniques and technologies, even looking at its smallest (but now prevalent and dominant) domain, that is, deep learning (Narayanan and Kapoor, 2024). Accordingly,

also definitions vary, ranging from specific (Kaplan and Haenlein, 2019) to more general ones (Goodfellow et al., 2018), where one commonly agreed upon (see Collins et al., 2021; Grilli and Pedota, 2024) is the one offered by Rai et al. (2019, p. iii): '[AI is, author's note] the ability of a machine to perform cognitive functions that we associate with human minds, such as perceiving, reasoning, learning, interacting with the environment, problem-solving, decision-making, and even demonstrating creativity'.

But what is so special about AI? For our purposes, and as we will argue more deeply later in the book, it is its autonomous *predictive* and *generative* capabilities, and the power associated with these features (which may also extend to the ability to take action, i.e. *agentic* AI) that distinguish it from all other technologies that preceded it and characterized previous revolutions.[1] In fact, other technologies in the past have shown the same characteristic of having a broad potential spectrum of application – namely *general-purpose technologies* (Jovanovic and Rousseau, 2005), such as electricity in the 1920s. Others have even risen to the role of *inventing a method of inventing* (Griliches, 1957), such as nanotechnology at the end of the last century (Darby and Zucker, 2005), which in turn opened up an unprecedented range of commercial opportunities. What seems new about AI, compared to all previous technological waves, however, is the apparent shift in the human–machine relationship of the technological element from a tool, however complex and sophisticated, into something that can, in principle, be elevated to at least a co-leading actor in the innovation process. This means that if we were accustomed to always seeing a human being behind a technology and an innovation, recognizing this with AI could increasingly be perceived as more difficult. In 1986, in a seminal article, Melvin Kranzberg expounded six laws on technological development that would later become famous, with the sixth and final one reading:

> Man is a constituent element of the technical process. Machines are made and used by human beings. Behind every machine, I see a face – indeed, many faces: the engineer, the worker, the businessman or businesswoman, and, sometimes, the general and admiral. Furthermore, the function of the

[1] After the first marked by mechanization, the second by electrification, and the third by computerization, AI is seen as a central element for the fourth and fifth industrial transformations.

technology is its use by human beings – and sometimes, alas, its abuse and misuse. To those who identify technology simply with the machines themselves, I use the computer as a metaphor to show the importance of the interaction of human and social factors with the technical elements – for computers require both the mechanical element, the 'hardware', and the human element, the 'software'; without the software, the machine is simply an inert device, but without the hardware, the software is meaningless. We need both, the human and the purely technical components, in order to make the computer a usable and useful piece of technology. (Kranzberg, 1986, p. 558)

It is clear that without necessarily going so far as to contradict this sixth law, or moving towards rather extreme theses of singularity (which many experts might consider sci-fi or pseudo-scientific), with the AI revolution, the balance in the relationship between man and machine may shift more and more in favour of the latter (e.g. Makridakis, 2017; Bouschery et al., 2023; Grilli and Pedota, 2024). To use Kranzberg's own metaphor, now the 'hardware' can make itself at least some of the 'software' part. In other words, the ability of artificially intelligent systems, through deep learning and neural networks models, to learn and imagine 'new things' or discover 'new patterns', without (or with limited) human intervention justifies the eradication of the learning and generative dimensions from the sphere of human exclusivity.[2]

Clearly, this shift requires scholars to rethink the complex nexus of relationships between the material and conceptual antecedents of innovative activity (Amabile, 2020; Grilli and Pedota, 2024). In this

[2] That said, our approach here is certainly not at odds with a critical view of the prevailing trend to 'anthropomorphize the unknown' and to attribute autonomous minds to non-human entities (see our Chapter 4), today represented by computing power, data centres, machine learning, and algorithmic modelling (Pasquinelli, 2016, p. 1). What we argue, however, is that while it may be true that AI is not dissimilar to other technological revolutions in allowing for greater labour automation than in the past, we also argue that this labour automation is at an all-time high, going for the first time in human history to affect high-profile functions and tasks (Pasquinelli, 2023). It is sufficient here to recall the words of the 2024 Nobel Prize in Physics laureate, Geoffrey Hinton, who, in the aftermath of receiving the award for his research on machine learning and neural networks, stated: 'I think it [i.e. AI, author's note] will have a huge influence. It will be comparable with the industrial revolution. But instead of exceeding people in physical strength, it's going to exceed people in intellectual ability' (The Nobel Prize, 2024).

particular area, we focus on the process of innovative entrepreneurship, which we need to describe in detail in the following pages of this chapter, before asking how it might change with the advent of AI, and how this in turn might change capitalism, perhaps in a new type or variety.

'Hamlet' at Centre Stage

Entrepreneurship, and innovative entrepreneurship in particular, lies at the very core of any capitalist society. In the words of Bowles (2013), already recalled in the Introduction, it is the entrepreneur – whose success, though partly due to luck, is above all the result of merit, stemming from skills and effort – who makes the capitalist system socially acceptable. Yet this centrality of the entrepreneurial figure, which may sound simple, even commonplace, to many ears today, has not always been so firmly established in economic thought. It was not until 1968, in a later influential article, that Baumol argued how, with few exceptions, the enterprise in formal economic models had hitherto been considered 'entrepreneurless' and posed his famous parallelism saying that it was as 'the Prince of Denmark ha[d] been expunged from the discussion of Hamlet' (p. 66). Fortunately, things have changed a lot since then, and today, both in public opinion and in the scientific literature, the 'Prince of Denmark', if not the centre of the stage, has gained a leading role. Capitalism thrives on competition, and entrepreneurs are considered to be key players in this competitive landscape.[3] The pursuit of profits should drive entrepreneurs to improve productive efficiency, reduce costs, and enhance the quality of goods and services. Start-ups have the potential to influence the conduct of existing companies, guiding prices towards competitive levels and ensuring allocative efficiency. Particularly in high-tech sectors, new ventures may challenge existing technologies controlled by established industry leaders, create new market segments from sketch, and foster dynamic efficiency (Audretsch, 1995). In a nutshell, entrepreneurship is reputed

[3] It is important to emphasize, however, that for Baumol, entrepreneurship – more precisely, the entrepreneurial spirit – is by no means an exclusive prerogative of capitalist society, but rather a fundamentally human condition that transcends epochs and social systems (Baumol, 1990).

to benefit consumers by offering them a variety of choices and promoting continuous improvement in products and services.[4]

In this respect, what really matters for scholars studying entrepreneurship is often not just the rates of entrepreneurial activity themselves in an economic system but, more specifically, the extent to which an economic system is capable to generate truly ground-breaking entrepreneurial actions stemming from the identification, creation, and exploitation of genuinely innovative opportunities (Baumol and Strom, 2007). In other words, assuming that having an entrepreneurial spirit is probably beneficial (or at least not harmful) in many human activities, it is nevertheless true that a capitalist society should not be primarily concerned with the 'quantity' of entrepreneurs – that is, the number of individuals who decide to pursue an entrepreneurial career – but rather with their 'quality' (Shane, 2009), namely that those who have promising innovative ideas and the capability to realize them are placed in the best conditions to try (and possibly to try again, given the uncertainty that naturally surrounds any truly innovative activity, and the not insignificant possibility of failure). Thus, an 'entrepreneurial society' (see Economidou et al., 2018) should create the conditions so that those who have the skills to experiment and (potentially) succeed as innovative entrepreneurs are in the best position to do so (and probably discourage others).[5] In turn, this has to do with the formal and informal institutions (North, 1990, 1991) in place in a given society, which, to paraphrase another famous work of Baumol (1990), may direct the entrepreneurial talents of individuals towards activities that can be productive, unproductive, or destructive.

[4] Following a well-established literature in industrial organization, it is worth recalling that social welfare can be characterized either by the sum of consumer and producer surplus (in a static setting) or by reference to the three fundamental notions of efficiency – namely, productive and allocative efficiency (static) and dynamic efficiency (i.e. innovation). See, among many others, Padilla (2024). It is clear that innovative entrepreneurship can benefit all three dimensions of efficiency, with a greater emphasis, almost by definition, on the latter.

[5] The literature has highlighted and subsequently addressed a potential paradox concerning entrepreneurship. While entrepreneurship significantly contributes to economic growth and performance, the least developed countries exhibit higher rates of entrepreneurship. The 'entrepreneurship paradox' is resolved by differentiating entrepreneurship based on quality, specifically high and low, and acknowledging that only the former (also referred to as ambitious) has an impact (Acs et al., 2023).

But we will return later in both this chapter and the book on the role of institutions, now it is good to focus on the figure of the entrepreneur in the words of the economist who first recognized this actor as the leading role in the capitalist plot.

The Schumpeterian Entrepreneur: Early and Late Versions

Schumpeter Mark I

Joseph Alois Schumpeter, in his seminal work, *The Theory of Economic Development*, first published in German in 1911 (and later translated into English in 1934), emphasized the entrepreneurial role in introducing innovations that disrupt existing economic structures. While the term 'creative destruction' was actually popularized only later (in his book, *Capitalism, Socialism and Democracy*, 1942), the process he describes is the one in which entrepreneurs, through the introduction of new technologies, products, or business models, destroy old ways of doing things and create new opportunities. Thus, Schumpeter saw the entrepreneur as the driving force behind economic development. In his own words, it is a sort of modern hero who is capable to cope with uncertainty, make crucial decisions that shape the direction of the start-ups, and bring together all various factors of production and necessary components for entering new markets with (not yet proven) ideas. Hence, *entrepreneur* is for him synonymous with *innovator*, and entrepreneurship is more 'a state of mind' rather than an innate and immutable characteristic that an individual possess:

> [E]veryone is an entrepreneur only when he actually 'carries out new combinations', and loses that character as soon as he has built up his business, when he settles down to running it as other people run their businesses. This is the rule, of course, and hence it is just as rare for anyone always to remain an entrepreneur throughout the decades of his active life as it is for a businessman never to have a moment in which he is an entrepreneur, to however modest a degree. (Schumpeter, 1911/1934, p. 78)

What motivates entrepreneurs to be entrepreneurial in Schumpeter's sense? What are the driving incentives? Basically, we could say that there is more than economics, but surely, economics, or to better say, economic incentives, play a major role. Firstly, the entrepreneur in Schumpeter's vision (Schumpeter, 1911/1934, pp. 93–94) is primarily

moved by 'the dream and will to found a private kingdom'; therefore, expected supranormal profits arising from the successful introduction of innovation serve as a reward for the risk-taking and dynamic activity of entrepreneurs. Secondly, 'there is the will to conquer: the impulse to fight, to prove oneself superior to others, to succeed for the sake, not for the fruits of success, but of success itself. From this aspect, economic activity becomes akin to sport'. Finally, the more hedonistic, and less economically related motive, which is 'the joy of creating, of getting things done'. The so-called Schumpeter Mark I (SM I) process therefore envisages that entrepreneurs introduce innovations in market with the expectation of making extra profits, where being the first and only in the market increases this expectation. Adhering to this vision, lack of expected extra profits as a reward for innovative entrepreneurial activity will inevitably result in a low rate of innovation for the economic system, and therefore, the importance of the patent and intellectual property rights (IPRs) system comes into play in order to prevent immediate imitation and stimulate dynamic efficiency.[6] However, over time, the innovator's 'entrepreneurial profit' will inevitably decline due to competition from imitators. Thus, according to Schumpeter, innovation operates in cycles, marked by alternating phases of rapid technological change and entrepreneurial activity followed by periods of relative stability. This cyclical pattern shapes the trajectory of economies, leading to waves of 'creative destruction'. This, in turn, not only propels economic growth but also contributes to increased social mobility, where entrepreneurship can effectively serve as an important social elevator:

> [T]he upper strata of society are like hotels which are indeed always full of people, but people who are forever changing. They consist of persons who are recruited from below to a much greater extent than many of us are willing to admit. (Schumpeter, 1911/1934, p. 156)

Thus, in this initial version, Schumpeter sees the entrepreneur as the primary driver of economic development. The entrepreneur is portrayed as a romantic hero who, motivated by the desire to succeed, brings about increasing levels of prosperity and social dynamism for the entire economic system.

[6] Even if, admittedly, Schumpeter never put a strong emphasis on IPRs protection, provided that other appropriability mechanisms are in place and work effectively (e.g. secrecy).

Schumpeter Mark II

But this Schumpeter's early vision was almost completely overturned by Schumpeter himself in his late book, *Capitalism, Socialism and Democracy*. The so-called Schumpeter Mark II (SM II) explicitly recognizes the obsolescence of the entrepreneurial function in economies characterized more and more by the presence of a few large firms operating in highly concentrated markets and where:

> [I]nnovation itself is being reduced to routine. Technological progress is increasingly becoming the business of teams of trained specialists who turn out what is required and make it work in predictable ways. The romance of earlier commercial adventure is rapidly wearing away, because so many more things can be strictly calculated that had of old to be visualized in a flash of genius. (Schumpeter, 1942, p. 132)

And again:

> Thus, economic progress tends to become depersonalized and automatized. Bureau and committee work tends to replace individual action. […] Rationalized and specialized office work will eventually blot out personality, the calculable result, the 'vision'. The leading man no longer has the opportunity to fling himself into the fray. He is becoming just another office worker – and one who is not always difficult to replace. (Schumpeter, 1942, p. 133)

In the innovation process, therefore, the large technological structures, such as Research & Development (R&D) laboratories, of munificent (monopolistic or oligopolistic) companies play a key role, as size-related R&D economies of scale and scope are important. This circumstance, combined with the likely capital market imperfections suffered by start-ups in finding the necessary resources to truly innovate, in a context where innovation is becoming increasingly scientifically complex, leads him to the conclusion that innovation is primarily an issue for large companies operating in highly concentrated markets:

> As soon as we go into details and inquire into the individual items in which progress was most conspicuous, the trail leads not to the doors of those firms that work under conditions of comparatively free competition but precisely to the doors of the large concerns – which, as in the case of agricultural machinery, also account for much of the progress in the competitive sector – and a shocking suspicion dawns upon us that big business may have had more to do with creating that standard of life than with keeping it down. (Schumpeter, 1942, p. 82)

The two visions of SM I and SM II (Malerba and Orsenigo, 1995), which are antithetical in many ways, have coexisted since their inception and remain today not only the two main institutional archetypes for evaluating the ability of economic systems to generate innovation at different levels (e.g. industry or geography) but also an indispensable benchmark for comparing new perspectives on innovation dynamics (Jacobides et al., 2021; Rikap, 2024a).

Institutions and Entrepreneurship

The relative strength of these two patterns, and the eventual dominance of one over the other, also has to do, we would say to a large extent, with the institutions in place. In his work, North (1990) defines institutions as constraints created by humans that shape how people live and interact.[7] He categorizes institutions into two main groups: *formal* and *informal*. Formal institutions encompass economic, political, and contractual rules, while informal ones include social norms, codes of behaviour, and cultural conventions specific to a geographical context (North, 1990, 1994). Formal institutions should be aligned with cultural values, as rules in politics, economics, and contracts are influenced by societal perceptions of 'how things should be done' (Hofstede et al., 2010); and it is important to note that according to this view, the very same formal institutions in societies with different informal institutions (e.g. cultural values) can lead to very distinct economic outcomes (North, 1990). In simpler terms, both formal and informal institutions significantly shape individual behaviour in a society, either independently or in combination (what North defines as 'the institutional matrix') and lead the economies down different paths (related to this concept, see also Hall and Soskice, 2001 and their notion of 'varieties of capitalism').

Entrepreneurship is clearly not immune from this combined influence. The literature generally agrees that formal and informal institutions in place in a given geographical or temporal context could be more or less favourable for triggering and nurturing ambitious innovative

[7] See North (1990, pp. 3–4, emphasis mine): 'Institutions reduce uncertainty by providing a structure to everyday life. They are a guide to human interaction, so that when we wish to greet friends on the street, drive an automobile, buy oranges, borrow money, *form a business*, bury our dead, or whatever, we know (or can learn easily) how to perform these tasks'.

entrepreneurial phenomena (Baumol, 1990; Acs et al., 2008; Sobel, 2008; Parker, 2009; Li and Zahra, 2012). The list is long as to these 'enabling institutions', and it can be sufficient here to recall among others, property rights protection, savings policies, taxation, labour and financial market regulations, competition policy (Henrekson, 2007) for what concerns formal institutions. But also informal institutions are found to play a significant role in entrepreneurial dynamics, reflecting the extent to which a society values and supports entrepreneurial behaviour (Beugelsdijk, 2007). Cultures that value risk-taking, innovation, and individual success, while avoiding stigmatizing failure which is seen as an unfortunate but often natural evolution of entrepreneurialism, tend to foster successful entrepreneurial activities. Correspondingly, societies that display high levels of trust among their members can diminish transaction costs and facilitate cooperation among entrepreneurs and other stakeholders. Furthermore, with regard to informal institutions that matter for entrepreneurship, we cannot overlook the potential impact of role models and entrepreneurial icons: a community in which entrepreneurship is perceived as attractive and successful entrepreneurship as even more desirable can serve as a great source of inspiration for individuals, motivating them to consider embarking on their own entrepreneurial journey.

Thus, we do not know who will be the successful entrepreneurs of the newly created innovative start-ups that will challenge incumbents and bring a new wave of 'creative destruction' to markets, but at the same time, the institutional view provides interesting guidelines for determining which macro and meso characteristics of an economic system are more likely to allow them to flourish.

The Innovative Entrepreneurial Process

But what do we know about the microfoundations of successful innovative start-ups? First, it is appropriate to identify the classic three distinct stages with which the entrepreneurship literature describes the entrepreneurial process (Shane and Venkataraman, 2000): recognition or creation of the opportunity, evaluation and exploitation.

As to the first phase, economists and management scholars have long debated, and still continue to debate, on whether entrepreneurial opportunities are simply discovered, which means that they potentially existed before discovery; or, conversely, they are created

(see Alvarez et al., 2013 for a discussion). Admittedly, there is also a growing stream questioning the conceptual usefulness of building the dichotomy (see McMullen et al., 2007; Davidsson 2023), and also the scientific value of the opportunity construct itself (Foss and Klein, 2020). However, others in the past have maintained the validity of the concept (e.g. Shane, 2003; Ardichvili et al., 2003; Baron and Ensley, 2006), and contemporary research still supports today this assertion (Nambisan et al., 2018; Suddaby et al., 2015).

Our position here is that entrepreneurial opportunities are obviously different from one another, where (regardless of their profitability) one opportunity might be more accessible and another more esoteric.[8] However, a certain degree of foresight and creativity in predicting or imagining something that is not visible at the moment (even if it is potentially there) is absolutely necessary. Consequently, some imagination or predictive capability on the part of the prospective entrepreneur is a key precondition for starting any type of entrepreneurial process. In other words, it is legitimate to expect that innovative entrepreneurial processes – those involving the creation of innovative start-ups (potentially backed by some invention), which are the type of entrepreneurial processes we are concerned with here – should require some creativity skills. Secondly, we refrain from engaging in the discourse surrounding the notion of entrepreneurial opportunity, the one about discovery or creation of this opportunity. We find this (mostly academic) debate not particularly fruitful (as others in the field, e.g. Ramoglou and McMullen, 2024) and, even lexically, we want to position ourselves as far away from the dispute as possible. For this reason, we adopt a concept that, in some ways, we do not entirely deny is somewhat analogous, but undoubtedly less academically connoted and therefore, much more peaceful: that of the *idea*. Moreover, the term 'idea' is by far the one most commonly

[8] Clearly, boundaries are difficult to draw, and probably the identification and creative elements simply coexist. To express this objective difficulty, let us borrow the words of Shane et al. (2003, pp. 262–263) describing the birth of the personal computer industry: 'Before Apple Computer, there was not much of a PC market outside of a few schools and businesses and a few thousand hobbyists using primitive computers. So what was the "opportunity?" There was mainly a potential market – people who might buy a more sophisticated PC if offered. The success of the PC would depend on Steve Jobs' approach, including the quality of the design features, the software, the price, and Apple's ability to market the product'.

used by entrepreneurs themselves, who rarely refer to 'opportunity' when describing the genesis of their start-up. In their view, which we therefore fully embrace, it all starts with an idea. A business idea. Of course, it is an idea that can change, that can move along the process in the early stages or be stopped, that can be modified ('or be pivoted' according to the Lean Start-up framework's lingo, Ries, 2011) in a sort of back and forth process between the recognition and evaluation (or even exploitation) phases, but the very genesis of the innovative entrepreneurial process is still represented by the identification or creation of an initial original idea.[9] Thus, creativity is a naturally important antecedent of this process.

Creativity

Among the many different angles from which the phenomenon has been studied (Kaufman and Sternberg, 2021), a widely accepted perspective on the nature of (human) creativity is that it is a combination of divergent and convergent thinking (Grilli and Pedota, 2024). On the one hand, divergent thinking involves generating a variety of ideas, exploring different possibilities, and thinking expansively without strict limitations (Guilford, 1984). On the other hand, convergent thinking involves analyzing and evaluating those ideas to arrive at a focused and practical solution (Cropley, 2006). This duality is often considered essential in the creative process. The ability to think divergently allows individuals to break away from traditional patterns and come up with innovative ideas, while convergent thinking helps in refining and transforming these ideas into meaningful outcomes. Thus, the two types of thinking highlight the dynamic interplay between exploration and refinement in the creative thought process, and their synthesis is believed to foster a balanced and effective creative mindset, enabling individuals to navigate the complex and dynamic nature of creative problem-solving.

[9] Contrary to what one might think, an innovative (successful) entrepreneurial idea is rarely the result of a sudden 'eureka' moment or an unexpected flash of inspiration. Instead, it is usually the outcome of an ongoing process of trial and error, involving adjustments aimed at finding the right balance between feasibility and market appeal. It is no coincidence that the term 'minimum viable product' is frequently discussed in the literature, especially due to the influence of the Lean approach.

Creativity can clearly also exert an important influence on the other two phases of the entrepreneurial process, that is, evaluation and exploitation, and generally on the entire life of an innovative start-up. Suffice it to mention here the role of convergent thinking in the evaluation phase, aimed at selecting the preferred option by the innovative entrepreneur. Or, following the entrepreneurial bricolage (e.g. Baker and Nelson, 2005) or effectuation theory (e.g. Sarasvathy, 2001), the entrepreneur's ability to combine the few available resources to reach the market and/or to redefine the objectives of the nascent entrepreneurial venture. All activities that necessarily rely on some creative element of the entrepreneur, and in their ability to 'find new solutions' to the difficulties they are bound to encounter along the way. In fact, the innovative entrepreneurship is typically a very resource-constrained process (Botelho et al., 2021) and, among other things, is the art of choosing between different options (e.g. Gans et al., 2019). Creativity, with its ability to 'circumvent' obstacles and/or find unexpected solutions to problems, reduces the number of choices and/or widens the range of options available for selection. It thus becomes a fundamental feature of the innovative entrepreneurial process, a sort of 'conditio sine qua non', which of course requires other elements as well. More specifically, besides creativity (but also linked to that, see *infra*), in this continuous struggle for the innovative entrepreneurial idea to emerge and eventually flourish in markets, the entrepreneurship literature typically identifies three important types of 'capital' that should support the entrepreneurial process and are considered important. Of course, they are not an insurance of success, but at the same time they are almost universally recognized as important drivers for the success of an innovative start-up: they are *human*, *social* and *financial* capital.

Human Capital

Human capital, and in particular founders' human capital, shapes the destiny of the innovative start-up, from its inception to its later stages. As early as 1977, in one of the early, later to become seminal works, searching for the microeconomic roots of entrepreneurial success in innovative sectors, Cooper and Bruno wrote:

> In many ways a new firm is an extension of its founder or founders. If the founder is strong in engineering, weak in marketing, and completely lacking in financial skills, then the new company must struggle with that

combination of strengths and weaknesses. If a team of founders is carefully assembled, skills can be complementary and the new firm is more likely to have the full range of capabilities needed to survive and prosper. The team can examine proposals from more than one point of view and members can cross-check each other's thinking. They also can exchange psychological support. One entrepreneur commented, 'When you wake up in the middle of the night and begin to wonder if you were crazy to have started a company, it helps to know that others are in it with you'. (Cooper and Bruno, 1977, pp. 20–21)

Being a central tenant of the competence-based view of the enterprise (Penrose, 1959; Hodgson, 1998; Freiling et al., 2008), the relationship between the human capital of founders and the performance of innovative start-ups has been the focus of numerous investigations that have been almost unanimous in recognizing the importance of the human factor (Colombo and Grilli, 2005; Gimmon and Levie, 2010; Unger et al., 2011), especially on the founder side, given the difficulty in attracting qualified human capital as employees, in view of the inherent great uncertainty of any entrepreneurial project, if truly innovative (Rocha and Grilli, 2024). In this respect, while maintaining the traditional distinction between *generic* and *specific* human capital proposed by Becker (1975), it is the latter that is often found to have a particularly positive impact on the performance of innovative start-ups (e.g. Cooper et al., 1994; Colombo and Grilli, 2005; see Grilli, 2014 for a discussion). However, the evidence in this field is not entirely unanimous, given the intrinsic nuances and blurred boundaries of the variables involved, how they are empirically measured, and the contingencies to which they are subject (e.g. Grilli et al., 2020; Danneels et al., 2024). Typically, the specific component of human capital refers to the skills and expertise that entrepreneurs possess in both technical and commercial areas, which can directly benefit their start-up. This knowledge is commonly (though not always) acquired through prior professional experience, particularly in the same sector where the innovative start-up will operate. In principle, possessing specific knowledge of the market-technology 'duo' (or at least one of the two) appears to be a crucial component of the innovative entrepreneurial process, based on the observation that successful ideas tend to emerge primarily from those who 'know'. In fact, while heterogeneity in acquired knowledge – even if not directly related to the technological domain of the future start-up – can foster divergent thinking and

creativity, it is well-established that successful innovative ideas often stem from founders' knowledge capital developed in closely related fields (e.g. Shane, 2000 among others). Clearly, the likelihood that this holds true increases significantly in contexts characterized by high knowledge intensity and the technological depth of the innovation involved.

The influential 'Knowledge spillover theory of entrepreneurship' (Acs et al., 2009; see Audretsch et al., 2025a and 2025b for recent surveys, hereafter KSTE) is itself fundamentally based on this stylized fact, and there are numerous anecdotal stories highlighting the fundamental relationship between the specific human capital of the founders and the success of the innovative start-up – where it is not uncommon for this relation to manifest through a spin-off from an established organization, often an enterprise. The paradigmatic example is probably the creation of Fairchild Semiconductor in 1957 by a group of researchers from the Shockley Semiconductor Laboratory ('the traitorous eight', according to William Shockley's definition), where 'over the next two decades, dozens of companies [including Intel and National Semiconductor, author's note] would be spawned by the intellectual talent [or human capital, author's note] at Fairchild Semiconductor, giving rise to Silicon Valley' (Schilling, 2019, p. 235). But there are many more recent examples in the same vein. Both Brian Acton and Jan Koum had previously worked at Yahoo before founding WhatsApp, the popular messaging app. Koum was an engineer at Yahoo, where he became increasingly disillusioned with his professional role, as the company did not offer him the kind of work he was looking for. Disappointed by this lack of support, Koum decided to leave Yahoo in 2007 and in 2009 he teamed up with Acton, another former Yahoo employee, to found WhatsApp. The story of SAP, the renowned Enterprise Resource Planning (ERP) software, is somewhat similar: its five founders were IBM engineers working on the project, when they were informed that the project was no longer required. Rather than abandoning their work, they opted to depart from IBM and establish a new firm. Needless to say, the relationship between founders' specific human capital and the success of their start-up may not only stem from an interrupted professional history, where the incumbent company's lack of foresight towards the innovative ideas generated by its employees drives them to become entrepreneurs in order to bring those ideas to life. It can also emerge in less traumatic ways. For instance, Kevin Systrom worked at Google, and Mike Krieger had an internship

experience at Microsoft before they founded Instagram, the photo and video sharing social media platform, but there is no documented evidence that either the two tried to convince the respective firms to invest in Instagram. The founding and early growth of Instagram were more closely associated with Systrom's relationships with venture capital firms (including Benchmark Capital and Andreessen Horowitz) and individual angel investors (such as Jack Dorsey, co-founder of Twitter, and Adam D'Angelo, co-founder of Quora) rather than by attempts to develop the project internally or secure investment from Google or Microsoft. Moreover, the KSTE (especially in its later formulations) also highlights that the unexploited knowledge which may serve as the foundation for emergent entrepreneurial phenomena may stem from organizational sources other than firms – such as universities or research centres – and may follow different channels, such as the exploitation of patents or the practical application of scientific publications, in line with the more classic mechanisms of technology transfer associated with the academic world's third mission (Audretsch et al., 2025a and 2025b).

In light of these examples and mechanisms, therefore, while it is well-established that successful innovative ideas are based on the knowledge capital of the start-up founders (typically acquired in closely related fields), it is worth mentioning here that founders' (specific) human capital can be crucial at all stages of the process, not only for the identification or creation of the entrepreneurial idea but also and above all for its evaluation and exploitation. In this context, two key points warrant attention.

First, as the view of organizational imprinting emphasizes (Stinchcombe, 1965; Barringer et al., 2005; Grilli et al., 2020), all the influence of the founders codified in formal routines and procedures, but also embedded in something more ethereal (but no less important for that) or in established principles that contribute to forming the so-called corporate culture (see Hewlett-Packard's eleven 'Rules of the Garage', Menke et al., 2007, p. 56), can extend well beyond the founders' actual presence in the company, prolonging its effects for many years after their departure.

Secondly, in this area, it is not just a question of 'know how', that is, knowing what to do and how to do it better than others, but human capital, especially in its specific component, can affect the psychology of entrepreneurs and impact their drive to carry out what is often perceived as a genuine mission. In other words, an individual's

entrepreneurial spirit (or mindset) is only partly innate and partly learned (e.g. Nicolau and Shane, 2009). As such, specific forms of human capital may also influence entrepreneurial motivations – particularly *self-efficacy* and *need for achievement*, which are often considered two key hallmarks of successful innovators and entrepreneurs (Shane et al., 2003; Schindehutte et al., 2006; Schilling, 2019).[10]

The former aspect refers to an individual's (domain-specific) belief that they can successfully complete a particular task or achieve a particular goal (Bandura, 1997). In fact, self-efficacy is not a measure of a person's actual competence or ability, but rather a personal perception of one's own capabilities. While only partly innate, experience in similar tasks or fields – namely, specific human capital – can effectively enhance it, and individuals with high self-efficacy tend to face challenges with confidence and perseverance, viewing setbacks as opportunities to learn and improve. As to the second characteristic, the need for achievement is a motivational trait that represents the absolute desire and devotion shown by some individuals to complete (often difficult) tasks in order to achieve (often high) goals, stubbornly sticking to any problems to the point of overcoming them (see McClelland, 1965 and Schilling, 2019 for a more recent discussion). The concept is associated with the second and third Schumpeterian motivations behind entrepreneurship, which we mentioned earlier, namely the drive to do everything in one's power to achieve success in the chosen endeavour and the joy of working and creating for oneself.[11] And it

[10] Here, we focus on the two psychological characteristics, which we would define as 'motivational', that appear (at least to us) to be the most significant antecedents of entrepreneurial success in innovative contexts according to the literature. However, it is important to state that we do not aim to be exhaustive (see, for example, Parker, 2009, for a more comprehensive list). The scientific literature on psychology and entrepreneurship is vast, investigating the psychological traits that can influence the phenomenon of entrepreneurship – both successful and unsuccessful – starting with the so-called Big Five model. For a review, see Zhao and Seibert (2006), and more recently, Kerr et al. (2018).

[11] This is also well-expressed in the words of (successful) innovative entrepreneurs themselves. Dustin Moskovitz, Facebook co-founder, analyzing all the possible motivations behind the creation of a start-up, at the question 'so what's the best reason?' he provides the simple answer: 'You can't not do it', meaning that 'you're super passionate about this idea. You [think you, author's note] are the right person to do it. You've gotta make it happen' (Altman and Moskovitz, 2014).

is connected with specific human capital through the other interesting notion of *flow* (Csikszentmihalyi, 1990, 1997). This latter term refers to a *state of mind* in which the individual is completely absorbed in the task and derives joy from this complete absorption and detachment from everything unrelated to the task at hand or the chosen challenge. It is likely to happen with the right balance between the skills involved and the difficulty of the challenge ahead, but at the same time, the more experienced one becomes in a particular area, the more difficult the challenge ahead has to become in order to experience the 'flow' (see Schindehutte et al., 2006, p. 353; Schilling, 2019, p. 206). In principle, therefore, it can be argued that entrepreneurs characterized by higher values of specific human capital will engage with more challenging and potentially more disruptive entrepreneurial ideas, having more than others the mastery and tenacity to bring them to fruition.

Social Capital

To paraphrase John Donne and his famous poem, if 'no man is an island entire of itself', so much less is the entrepreneur and the start-up that the entrepreneur creates. First of all, innovative start-ups, especially successful ones, are rarely created by solo founders, but more often by the efforts of a team of individuals. As discussed earlier, the human capital of the founders is of considerable importance, but the strength of a team does not only lie in the accumulation of knowledge of its members, but also in how these competences interpenetrate each other, how effectively they interact within the start-up and how, in summary, one actually complements the other. In other words, logically, there is neither a perfect size of the founding team nor a perfect predetermined mix of these competencies, since human capital is then embedded in individuals with their own character profiles and idiosyncrasies. However, the literature on the subject seems to suggest that the coexistence of members with more technical backgrounds and founders with more commercial and business skills can lead to synergistic benefits (Colombo and Grilli, 2005; Ganotakis, 2012; Protogerou et al., 2017). Thus, from the initial spark of the entrepreneurial idea (which is almost always the result of a 'single' mind), knowing how to involve the 'right' people in evaluating and exploiting the idea, making them co-founders, becomes extremely important for the fate of the fledgling project. According to Schilling (2019), among

others, it is true that Apple would not have been possible if Wozniak's docile and patient character had not been able to accommodate Jobs' often brusque and authoritarian ways, and it is equally true that Apple also thrived because Wozniak's engineering genius in software combined well with Jobs' design skills, leadership and economic vision.

But an innovative start-up, especially if it aspires to success, is typically made up not only of its founders, who become owner-managers, but also and above all of its employees and workers. And it is in this area – the labour market – that the start-up is likely to encounter a first order of problems (a second, certainly not in terms of importance or timing, stems from the financial market, as we will see later). From one hand, attracting human resources in the early stages is considered of crucial importance (e.g. Yusubova et al., 2020), since early recruitment may irreversibly shape the future performance of start-ups (Siepel et al., 2017; Rocha et al., 2018; Gjerløv-Juel and Guenther, 2019). On the other hand, recruiting is objectively difficult for many start-ups for several reasons, on both the supply and the demand sides (see Rocha and Grilli, 2024). For example, the assessment of the potential of a start-up by prospective employees is challenging owing to the significant uncertainty surrounding this type of ventures and its inevitable lack of a track record, exacerbating information asymmetries. Moreover, innovative start-ups have a larger average failure rate compared to established firms, putting their employees at a heightened risk of job loss (Sauermann, 2018) with consequent potential backlashes to their career reputation (Rider and Negro, 2015). Additionally, small start-ups often offer lower compensation compared to larger firms (Haltiwanger et al., 2013; Burton et al., 2018). Other challenges also exist on the demand side of start-ups, as many start-ups may easily lack sufficient proper human resource management systems and policies for conducting recruiting in a professional manner (Cardon and Stevens, 2004).

All these obstacles reduce both the number of potential candidates that might be interested in working for an innovative start-up and the number of valuable candidates that the start-up can credibly attract. In this respect, having a network of contacts that is both extensive and dense – comprising weak and strong ties, respectively, in Granovetter's (1973) classic terminology, or bridging and bonding social capital – can help alleviate these challenges. Such a network supports the search for suitable individuals and, when combined with

legitimacy and reputational capital, enhances the ability to effectively attract qualified employees.[12]

And it is clear that the *network reach* and *legitimacy* of a newly founded start-up, that is, its social capital, can only coincide with that of its founders. Legitimacy, provided by the founders (and their human capital) can profoundly benefit an innovative start-up, 'opening doors' that otherwise could not be opened (Bertoni et al., 2011, p. 1039). Indeed, it can help innovative start-ups gain access to suppliers of complementary goods that are vital for start-ups and their success (Teece, 1986), in a virtuous process where 'legitimacy begets legitimacy' (Rutherford and Buller, 2007, p. 79), and thus new and new doors are opened. Thus, the underlying transactions, which may also take the form of formal alliances, with a variety of actors, including external financiers, input suppliers, R&D partners, distribution channels, sales and marketing forces, may determine the difference between reaching markets and succeeding or not, where all these relationships may simply be out of reach in the absence of sufficient initial endowment of the founders' (and thus their start-up's) social capital. Clearly, social capital is also important for attracting investment from external financiers, in particular specialized financial operators such as business angels and venture capital funds. While a handful of quantitative studies (e.g. Ahlstrom and Bruton, 2006; Bottazzi et al., 2016) provide clear documentation of the relation between social capital and VC, they are not plentiful (see also Grilli et al., 2019). It is therefore noteworthy to mention a few exemplary cases in this regard. The founders of Airbnb joined the prestigious start-up accelerator Y Combinator and were mentored by Paul Graham, the co-founder of Y Combinator, who was a key figure in the development of the start-up certainly because of his coaching skills but also because of his contacts with the investment community.[13] Again, in the early days of

[12] Building on the seminal work of Putnam (2000), it is not uncommon in the entrepreneurship literature to refer to the two types of ties as *bridging* (typically corresponding to weak ties) and *bonding* (typically corresponding to strong ties) forms of social capital (see, for example, Davidsson and Honig, 2003). In a nutshell, despite some nuances, the two dimensions can be considered largely overlapping: one emphasizes the breadth of the network, while the other focuses more on trust and legitimacy among its members.

[13] Harrison Faull (2023) reports an interesting e-mail exchange between Paul Graham and VC Fred Wilson, in which the former tries to convince the latter to invest in Airbnb. Fred Wilson eventually will decide not to invest in the

Facebook, Mark Zuckerberg secured a crucial investment from VC Peter Thiel, where this connection was first made by Reid Hoffman, co-founder of LinkedIn and another influential figure in Silicon Valley (Kirkpatrick, 2010). But emblematic again is the story of Apple, before the launch of Apple II, as Schilling (2019) tells us, where it is evident the importance for an entrepreneur (in this case Steve Jobs) to know and be esteemed by a potential investor (his former employer), who may eventually know and be esteemed by another potential investor, and so on:

> [T]hese design decisions [related to Apple II, author's note] required money, which Woz and Jobs didn't have. They also needed managerial expertise to help run the growing company. Fortunately, being in the heart of Silicon Valley, with its dense network of information technology companies and investors, gave them access to both. Jobs approached his former employer, Atari's Nolan Bushnell. Bushnell declined Job's offer to sell him a one-third stake in the company for $50,000 (Bushnell later remarked, 'It's kind of fun to think about that, when I'm not crying') but referred him to Don Valentine, the founder of Sequoia Capital, a Menlo Park venture capital firm. Valentine also declined to invest, but he gave Jobs a referral to Mike Markkula Jr., a former marketing manager at Fairchild Semiconductor and Intel, who had made millions on stock options and retired at the age of thirty-three. Markkula agreed to join as a partner and provide capital, and in April 1977 they launched the Apple II […]. (Schilling, 2019, pp. 249–250)

Financial Capital

The story of the financing of Apple II also tells us that finding finance is hard for an innovative start-up and its founders, even if your entrepreneurial idea is (extremely) valuable, where what is valuable or not is clearly easy to tell ex-post but it is really difficult to be ascertained ex-ante. While academics discuss many things, and in the social sciences different views on the same issue are often possible and sometimes even consistent with each other,[14] the likely presence of market failures in the provision of funds to innovative start-ups and new

start-up, but after some time Airbnb will secure a significant seed investment from Sequoia Capital, a renowned VC firm.

[14] We make our own the words of Dani Rodrik, in its *Economics Rules* book (p. 16), when he says: 'The correct answer to almost any question in economics is: It depends'.

technology-based firms approaches the 'near universal recognition' (Storey and Tether, 1998, p. 1049). In a nutshell, a large amount of asymmetric information between the two parties of a financing transaction, the investor/lender on the one hand and the start-up on the other, can often lead to problems of adverse selection and moral hazard in the relationship between the two figures, which in turn can prevent the realization of important innovative projects (or prevent them from being realized on a sufficiently large scale), with associated losses in private but also social welfare. The remedies implemented by private operators, such as the use of certain signals by innovative start-ups, are at best only partial and, in any case, not without cost, since any signal in economics has to be costly to be effective (Spence, 1973). Public policy measures implemented to alleviate problems in this area are also far from perfect and, in general, have had mixed fortunes, highlighting the absence of universally valid magic formulas (e.g. Parker, 2009; Audretsch et al., 2020; Bradley et al., 2021).

Venture Capital

In this landscape, considering old and new actors and methods (e.g. crowdfunding, initial coin offerings, etc.) there is still one figure in the domain of finance for innovative start-ups that is often likened to a kind of modern superhero: *venture capital*. We briefly illustrate (some of) the most salient points about how VC funds operate in Box 1.1.

Although the VC segment is mainly developed in the United States (and moreover, heavily concentrated in three areas: Greater Boston, New York, and the San Francisco Bay) and in only a few other global regions, any successful start-up named globally is far more likely to have received some VC funding in the early stages of its life than not. Without stating that VC was the only determining factor in their success, nor claiming that it is always essential for any start-up, it suffices to note that all five MAMAA companies received VC funding in their early stages, when they were genuinely innovative start-ups (Nicholas, 2019). Thus, without minimizing the (at least potential) importance of newer and more innovative forms of start-up financing – such as crowdfunding (Mollick, 2014) – which often serve as a key complement or precursor, it can be stated without fear of contradiction that the lion's share has historically been held by venture

Box 1.1 How VC works

VC funds raise capital from investors and focus on investing in high-potential, early stage innovative start-ups (Gompers and Lerner, 2001). Fund managers (i.e. general partners) actively seek investment opportunities, conduct due diligence, and, upon successful negotiations with founders and entrepreneurs, invest in promising start-ups (i.e. the 'selection phase'), where, ex-post, they often provide strategic guidance and take an active role in portfolio companies (the so called coaching function, Hellmann, 2000). It is also common for VC funds to invest in syndication with other funds, pooling resources and sharing risks, especially in larger or riskier deals. In fact, the ultimate goal is to achieve profitable exits through initial public offerings (IPOs) or acquisitions. Then, profits are shared among investors (i.e. limited partners), where general partners usually receive carried interest. The typical fund's life-span is about 7–10 years, and its lifecycle concludes with the liquidation of remaining assets and the distribution of funds to limited partners (see Alemany and Andreoli, 2018, for a detailed description).

VC investments may target different stages of a start-up's development, ranging from seed funding – intended to support product development and initial market validation – to growth and expansion financing, which aims to scale operations. In the earliest stages, innovative start-ups may also be often supported by (networks of) business angels – wealthy individuals who provide initial capital (and mentorship) – who may play a key bridging role between the informal funding landscape and the more structured, institutional VC environment.

VC environment that indeed is institutionally diverse (e.g. Bottazzi et al., 2008). In fact, the above description primarily refers to what is often called *independent venture capital* (on which we mostly focus, given its central importance), which typically involves private, profit-oriented funds operating autonomously from other institutions. However, the broader VC space encompasses a variety of organizational forms, each potentially with distinct motivations, governance structures, and investment logics. These include 'captive investors' such as corporate VC arms, which may pursue strategic as well as financial objectives; public VC initiatives, often designed to address market failures or foster regional development; bank-affiliated funds; university-based accelerators; and hybrid models. These alternative forms of VC may differ substantially in terms of risk tolerance, investment horizon, and degree of involvement in portfolio companies.

capital. Accordingly, both in the academic literature and in the popular press, the mere attainment of VC backing is frequently interpreted as a significant accomplishment in its own right, not least because of the extreme rarity of the event (according to Wood, 2020, only about 0.05 per cent of start-ups in the US secure VC funding).

Thus, being extremely selective in choosing which start-ups to invest in, VC's business model is very often based on the notion that one win pays for many losses, where the horse-jockey debate is old and still very active (see Kaplan et al., 2009; Gompers et al., 2020; Esen et al., 2023): whether VCs invest more in entrepreneurial ideas and business-related considerations (i.e. 'the horse') or rather in entrepreneurial teams ('the jockey'). Since, as we have mentioned, human capital is often associated with good ideas, it is often difficult to distinguish between the two, and the attempt is probably a bit pleonastic.[15] But the role of VCs certainly does not end with *selection*, as they are active investors. In fact, another central theme in the literature is how much this type of investor contributes to the success of innovative start-ups not only by infusing the money that start-ups need to realize their business idea, but also by offering all those services in terms of *coaching* (including functions as mentoring and networking) in areas where start-ups may lack expertise. Only seemingly paradoxically, these VC skills seem to be even more important for VC-backed start-ups in institutional settings where the VC industry is not as developed (see, e.g. Bertoni et al., 2011), suggesting that in less munificent innovation ecosystems the role for VC could even be more important.

This also explains the reasons behind the attempts made by so many policy makers around the globe to nurture VC and make it grow. Apart from a few notable exceptions (e.g. China), it is legit to say that the VC model for start-up financing has struggled and continues to struggle to fully develop in other institutional contexts (especially in continental Europe) than the US, where it was originally 'invented' right after the Second World War. And this appears true despite the vast array of measures implemented (see Lerner, 2009): almost as if it were difficult to find the right combination of the most functional 'institutional matrix' for VC (Grilli, 2025).

[15] What obviously makes everything more complicated is the fact that founding teams evolve over time, and business ideas may pivot very rapidly in the start-up world.

The 'AI Revolution' and Entrepreneurship

In the previous pages, we started from the now widely accepted fact that innovative entrepreneurship is an essential component that 'good' capitalism should possess (Baumol et al., 2007), and then analyzed the conditions and ingredients that enable innovative start-ups to thrive in markets, often creating them. Continuing our discussion, we will explore the fundamental issue of whether and how entrepreneurship – and consequently capitalism – might change as a result of AI, and what role public policy should play in this area to ensure that social welfare is not only preserved but actually enhanced. Indeed, some of these start-ups, through productive forms of entrepreneurship aimed also at addressing broader societal needs, might eventually challenge established industry paradigms and threaten dominant positions, including those held by Big Tech firms such as MAMAA and other AI giants (though, as we will see later, this is far from simple).

The analysis that follows will touch upon all the key stages that, in our view, are essential to understanding how AI is reshaping the very fabric of innovative entrepreneurship. In particular, after examining the generation and implementation of an entrepreneurial idea in the innovative domain, we will focus on the three fundamental pillars of success discussed in this chapter, in order to analyse whether and how AI technologies are transforming the consolidated framework we have outlined here. Before doing so, however, it is necessary to clarify what AI technologies actually are – or, more precisely, what is meant by them – addressing their technical aspects insofar as these help illuminate the underlying socio-economic processes.

2 *Artificial Intelligence*
Setting the Scene

Prologue

As AI research has evolved, so have the definitions and approaches (Russell and Norvig, 2014). Today, AI encompasses a range of techniques, where terms such as machine learning, deep learning, and neural networks have become increasingly frequent and commonly used in popular media, as well as the subject of growing research across a vast array of scientific fields. At a high level, we can include in this category all technologies that enable the development of intelligent systems capable of performing complex tasks (Rai et al., 2019; Collins et al., 2021; Grilli and Pedota, 2024). In Chapter 1, we point to its *predictive* and *generative* capabilities as particularly distinctive. As to the first salient characteristic, at the heart of AI's learning prowess lies machine learning, which enables systems to discern patterns and make accurate predictions. This learning takes various forms, from supervised learning, where the AI learns from labelled training data; unsupervised learning, which involves finding hidden patterns in unlabelled data; and reinforcement learning, characterized by algorithms learning optimal actions through trial and error interactions with an environment. Deep learning, a subset of machine learning, has further enhanced this *predictive* capability, and it is the dominant force behind *generative* AI. Nowadays, when we talk about AI, we are mostly referring to deep learning, that is, multi-layer neural networks, where AI and deep learning have basically become synonymous.

Belonging to the 'subsymbolic' stream of research in AI rather than the 'symbolic' approach, the initial concept of neural networks dates back to McCulloch and Pitts (1943) and was further advanced by Rosenblatt (1958) through the development of the 'perceptron', one of the earliest neural network algorithms. It is certainly not the aim of our analysis to delve into the intricacies of the history of AI, with

its succession of different winters and springs, marked by overlapping periods of high expectations that were later disappointed, new promises and also concrete technological advances. But it is worth highlighting here that deep learning, the ultimate engine behind the current 'AI revolution', does not really configure itself as a 'new' technology. In fact, for a time, it was surpassed by symbolic AI (commonly referred to as 'good old-fashioned AI', or GOFAI, in the jargon of AI scholars) and the creation of 'expert systems'. In this regard, it is widely acknowledged that the rise of subsymbolic AI as the dominant paradigm in AI from 2010 onward – and what we, as noted, regard as yet another industrial revolution – is largely due to the vast availability of (big) data and the substantial increase in computing power over time (Mitchell, 2019).

Deep Learning

As already noted in the Introduction, from an economic perspective this technology has the potential to trigger an extremely wide spectrum of applications, cutting across industries and functional domains (Cockburn et al., 2019; Bick et al., 2024). By using neural networks with multiple layers, deep learning allows AI to model intricate patterns in data, leading to significant advancements in a great array of different areas, from image and speech recognition to drug discovery and genomics, to only mention a few (LeCun et al., 2015). As pure 'pattern-recognition' engines, the AI's *predictive* ability can also be complemented with a *generative* function. Unlike traditional machine learning models that primarily classify or predict (e.g. k-nearest neighbours, random forests, or k-means clustering) deep learning generative AI can also *create* new data and content, which then can be embodied (or be used as an input for) new products and services, or data instances, such as images, text, music, software code, synthetic information often indistinguishable from real world examples.[1] Consequently, new avenues have opened in creative fields, allowing AI to generate art, music, and textual content.

[1] Synthetic data is artificially created data that simulates real world data, used where actual data is unavailable or unsuitable. It may be essential in training AI models, particularly when real data is scarce, sensitive, or when diverse scenarios are needed (Nikolenko, 2021).

Over time, a distinction between *predictive AI* and *generative AI* has crystallized even in public opinion, although both, in their advanced form, still stem from the same technological root: deep learning. This evolution reflects the specialization of different types of models, but it is important to note that a certain degree of flexibility also exists with the same architecture being adapted for various purposes, whether predictive or generative. For instance, architectures like transformers, which form the basis of generative models such as Large Language Models (LLMs), are also widely used in predictive tasks including automatic translation, natural language recognition, protein structure prediction in biological systems – as demonstrated by the successive versions of AlphaFold over time – or even predictive healthcare, as aimed for by the Foresight AI model initiative in the United Kingdom (UK).[2] In the same vein, generative adversarial networks (GANs), originally designed to generate new data and content, can also be employed to enhance predictive models through data augmentation (and synthetic data) or anomaly detection. The list of (potentially ambivalent) deep learning architectures is quite long. Among the most well-known examples are variational autoencoders (VAEs) and convolutional neural networks (CNNs or ConvNets). Many of these techniques and methodologies can also be combined into more complex and sophisticated systems, such as in Large Multimodal Models (LMMs) or in (so-called) Large Reasoning Models (LRMs, see Mitchell, 2025), further expanding the range of AI technologies and their possible combinations. Box 2.1 provides a brief overview, highlighting how many of these technological advances in this field are often achieved through academia-industry collaborations, and how firms also compete on the scientific front within the AI domain. Both aspects will be revisited later, particularly in Chapters 6 and 7.

Taken together, these developments illustrate just how fast the field evolves. In this light, any list of AI techniques and models risks becoming quickly outdated, as new advancements continue to emerge and, to some extent, overlap. How far technological progress will go and what heights AI will reach in the future are not easy to

[2] LLMs, once fine-tuned with new and domain-specific data, can be employed for predictive tasks across a wide variety of fields – for instance, in domains as diverse as chemistry (Jablonka et al., 2024) and, more relevant to our context, in forecasting start-up success (Maarouf et al., 2025).

Box 2.1 AI breakthroughs: the 'duo' of academy and industry in AI science

Stemmed from the famous Google's article in 2017 (Vaswani et al., 2017), the transformer model is a deep learning architecture that uses a mechanism called self-attention to process input data (such as text) in parallel, rather than sequentially (which was the state of the art at the time, represented by the use of recurrent neural networks in natural language processing). This advancement allowed to capture long-range dependencies in data, making LLMs effective for tasks like long text generation, among (many) other uses. GANs, introduced by Goodfellow et al. (2014), consist of two neural networks contesting with each other in a zero-sum game framework, enabling the generation of new content. VAEs, on the other hand, compress and reconstruct input data into a probabilistic format, efficiently learning and generating data by organizing compressed information into a structured, typically normal distribution (e.g. Kingma and Welling, 2013). Finally, CNNs, introduced by LeCun et al. (1998) and popularized by AlexNet (Krizhevsky et al., 2012), use small filters to detect local features in data – usually images. By scanning these filters across an input, CNNs capture spatial relationships efficiently, making them especially useful in tasks like image classification and object detection. Their focus on local context also proves valuable for certain speech and language applications.

These represent only some of the most prominent deep learning architectures, a few of which we will examine in some detail later in the chapter. At this stage, however, we wish to highlight from the outset several interesting stylized facts that have gradually emerged and consolidated around them.

First, many of these advances (all of those mentioned above) have been made known to the world through scientific publications, often resulting from collaborative efforts between industry – including MAMAA or Big Tech companies more broadly – and academia. This complementarity between the two spheres has been well-documented in the literature (see Lundvall and Rikap, 2022; Liang et al., 2024; Rikap, 2024b). The significance of these scientific outputs for understanding AI's broader economic dynamics is also evidenced by the prominent attention devoted to this topic in the annual *AI Index Reports* produced by Stanford University, which have become an essential point of reference for tracking the latest developments in the field.

Second, many leading scientists have moved from academia to industry or have built enduring collaborations between the two spheres,

often taking on prominent roles within major technology firms. Among the three seminal figures of the deep learning revolution – popularly known as the godfathers of AI and joint recipients of the A.M. Turing Award in 2018 – Yann LeCun has long been affiliated with Meta (until late 2025) as Chief AI Scientist while continuing to serve as a professor at New York University (among other institutions). Yoshua Bengio, while maintaining a strong academic base at the University of Montreal, co-founded Element AI, later acquired by ServiceNow. In some instances, these transitions have not been without tension: Geoffrey Hinton, long based at the University of Toronto and recipient of the 2024 Nobel Prize in Physics, resigned from Google in 2023 to publicly express his concerns about the societal risks of AI, explaining that he wished to speak freely without implicating his former employer.

Third – and perhaps unsurprisingly, given the previous two points – the scientific domain of AI has often become a battleground for industry, with companies competing also through academic publications. A notable example is Apple's paper (Shojaee et al., 2025) on the challenges of LRMs – a field in which the company at that time was not actively involved, unlike several of its competitors.

determine, and our analysis is not intended as a foresight exercise. More specifically, in what follows, we will be taking both steps backwards and steps forwards at the same time as regards the technological aspects.

We will distance ourselves from technological specifics, aiming to maintain technical rigor while avoiding, where possible, the inherent idiosyncrasies of the many technological variants often grouped under the broad category of AI. In other words, we will use a degree of abstraction to raise our economic analysis to a level that can transcend the specifics of individual deep learning applications, embracing different types of multi-layer neural networks (e.g. feed-forward, convolutional, recurrent, adversarial, transformers, and autoencoders) that serve different purposes for different types of information (e.g. text, audio, video, images, and software codes). When we draw on examples based on specific technological products, we will do so to derive more general insights on the economic dimensions involved.

In order to do that, we will focus on the two functions we consider essential and that, to some extent, characterize all the technologies that are placed under the same umbrella term of 'deep

learning': the already invoked predictive and generative powers enabled by AI technologies. Here, we will take a few steps forward. In other words, while examining these two characteristics, we will certainly consider the possibility of some technological advancements in the near future, but always stay within the confines of what we can see on the horizon and, therefore, within the safe bounds of plausibility.

Focusing on the functionalities enabled by technologies, rather than the technologies themselves, will help us avoid getting too caught up in debates that often dominate public and academic discussions. In particular, we are here referring to the long-standing and unresolved debates on 'weak vs. strong' and 'narrow vs. general' AI, and on whether the latter terms of these distinctions are already in existence or might be possible in the more or less distant future. While our approach deliberately prescinds from choosing among these alternative scenarios, Box 2.2 provides a brief overview of what these terms actually mean.

While these distinctions may remain important in many areas, for the purposes of our analysis they introduce a level of complexity that we can afford to set aside when considering their implications for the entrepreneurial process.[3] AI is already powerful enough *today*, and it has the potential to transform a wide range of processes in *today*'s economy. Even an extremely conservative estimate of the progress these technologies will achieve in the near future suggests an amplification of AI's predictive and generative capabilities, and certainly not a regression to the past. This is what primarily matters in our scene. Whether these amplified capabilities will eventually manifest themselves through a variety of AI technologies, each applied to well-defined domains (and potentially communicating and coordinating with one another, as in the case of agentic AI), or through a limited number of machines capable of operating across the same broad range of domains by integrating them, is somehow less critical.

[3] It is important to note that all these are 'social' categories and, as such, potentially mutable over time. Here, we should recall the famous paradox mentioned by the founding father John McCarthy (who not only coined the term 'artificial intelligence' but also (co-)organized the first conference on the subject, namely the (long) summer workshop at Dartmouth College in 1956): 'As soon as it works, no one calls it AI any more' (reported in Meyer, 2011).

Box 2.2 On the elusive taxonomies of AI

Sometimes the two distinctions are considered as one, where weak is equated with narrow, and strong with general (e.g. Mitchell, 2019); but with no general consensus among the academic community either of what is precisely inside (and what not) in each of these four different categories (see, for example, McLean et al., 2023, and their review on 'Artificial General Intelligence', i.e. AGI). Clearly, this vagueness is not only partly due to the adjectives, but also, to a great extent, to the noun to which they refer. In other words, what 'intelligent' means and what form of intelligence one should consider is far from settled, regardless of the scientific field (e.g. neuroscience, psychology, computer science, etc.) from which the subject is approached. The same elusiveness affects the concept of 'consciousness'. That said, and given the intrinsic indeterminacy of any definition one might propose, it may still be useful, for illustrative purposes, to retain the two distinctions.

In particular, following Wooldridge (2020), 'narrow vs. general AI' concerns more the scope of intelligence. Narrow AI refers to artificial intelligence systems designed to perform a specific task or a limited range of tasks; these systems can learn and improve within their domain, but their learning is limited to the specific environment they are trained on. In contrast, general AI encompasses the ability to learn and apply intelligence across a wide range of tasks, mimicking human cognitive abilities. On the other hand, 'weak vs. strong AI' pertains to the nature of intelligence, and the dichotomy became famous after Searle (1980) and his 'Chinese Room' argument (somewhat in contraposition to the Turing test). In this regard, it is fair to say that while 'strong' AI and 'singularity' scenarios are not generally considered highly realistic by the AI community – at least in the short term (see, among many others, Wooldridge, 2020, pp. 239–243) – some notable scholars do not entirely dismiss the possibility of a sentient and conscious AI taking over the world (e.g. Geoffrey Hinton, interviewed by LBC, 2025, or Kurzweil, to whom we owe the concept of singularity in AI and who, in 2024, stated that the event is increasingly imminent). Meanwhile, general AI is admittedly viewed as a more plausible prospect within the AI community. If in fact it is true that some scholars in the current AI research framework consider AGI to be simply impossible (e.g. Roli et al., 2022), others instead argue that partial but concrete forms of general artificial intelligence already exist today (e.g. ChatGPT by OpenAI, cf. Bubeck et al., 2023) and see the arrival of fully equipped 'general AI' (Müller and Bostrom, 2016; McLean et al., 2023; Korinek, 2024) or even a 'superintelligent' form of AI (Bostrom, 2014) as entirely possible (or also likely).

AI: Data and Computing Power

By contrast, what is clearly critical to 'feed the machine' and its two identified capabilities are (big) data. The concept is typically expressed with the cliché 'data is the new oil' taken from an article in The Economist (2017). Yet, like all clichés, it remains fundamentally true, especially when paired with computing power and the associated hardware (HW, e.g. semiconductors and various 'processing units' such as CPUs, GPUs, TPUs, and FPGAs, data centres) and software (SW, e.g. cloud services) layers needed to manage, process, and make data productive.

In fact, data is essential at all key stages in the development of any deep learning technology, whether for predictive or generative applications. First, large and diverse datasets are required to train neural network models. For generative AI, this means using vast amounts of data to learn the underlying distributions and features, enabling it to generate new content that is coherent and contextually relevant. The predictive power of deep learning is manifested in the model's ability to learn (complex) patterns and make (accurate) predictions or classifications when presented with new, unseen data. In a subsequent stage, validation datasets are used to fine-tune model parameters and prevent problems such as overfitting,[4] ensuring that both predictive accuracy and creative outputs are optimized. In the testing phase, again new and separate datasets are used to evaluate the model's performance and reliability and to uncover possible biases and deficiencies. Lastly, even after deployment, data may continue to play an important role (with nuances and differences across different contexts). In fact, models are frequently updated and retrained with new data to adapt to changing environments and incorporate new information. Thus, data is central in every step of deep learning technology development, from

[4] Overfitting is probably one of the main problems in deep learning, much like 'endogeneity' in econometrics. Essentially, a model 'overfits' when it has a low training error but a high validation or test error, indicating that it has learned the specific details of the training data rather than a broader generalization applicable to other data (Goodfellow et al., 2018). As with endogeneity, there are different methods to address overfitting, one typical approach being 'dropout': a technique in which random neurons are 'dropped' or deactivated during training to prevent the model from becoming overly dependent on specific neurons. An example of overfitting in image recognition is nicely illustrated by Mitchell (2019, pp. 120–123).

training and validation to testing and continuous improvement post-deployment. The predictive and generative powers enabled by deep learning's capabilities are intrinsically tied to the quality and quantity of data available, and the computing power necessary to process them. Without ample, high-quality data at each of these stages, deep learning models, and their associated abilities to predict outcomes and generate novel content, would simply not exist. Nor would they exist without computing power, which provides the capacity to process vast datasets and train neural networks with billions or even trillions of parameters. Modern AI depends on specialized computing chips as well as large-scale cloud infrastructures. Yet this computational capacity is highly concentrated in the hands of the AI giants, with firms like NVIDIA and a few major cloud providers among the MAMAA group (Amazon, Microsoft, and Google) controlling a substantial share of the resources on which the entire AI ecosystem relies.

Taken jointly, data and computing power (in both its hardware and software components) represent the first building block of the entire deep learning process, the foundational resources without which the digital 'statistical machine' could not operate. They lead to the definition of the model and, together with it, constitute what we will later refer as the 'inferential infrastructure' of AI systems (see Chapter 6).

Deep Learning in Context: A Socio-Technical Representation

Building on these foundations, information – whatever its possible format – must be organized into proper datasets, which translate the raw material of data into usable input for the model. In Figure 2.1, we exemplify a 'socio-technical' representation of this process.

First, and this is something that is probably not entirely understood in media and popular press, '[n]eural networks require their inputs to be numbers' (Mitchell, p. 237). In practical terms, they operate exclusively on numerical data defined in a metric space. This means that data already expressed in numeric form (such as images represented as arrays of pixel intensities) can, in principle, be fed directly into architectures (as in CNNs). But when the original input is not numeric, it must first be mapped into such a space. For text, for instance, the sequence is tokenized and each token is associated with a (learned) vector representation; this step is often referred to as 'embedding the data' (Scardapane, 2024, p. 54), a term that more generally denotes

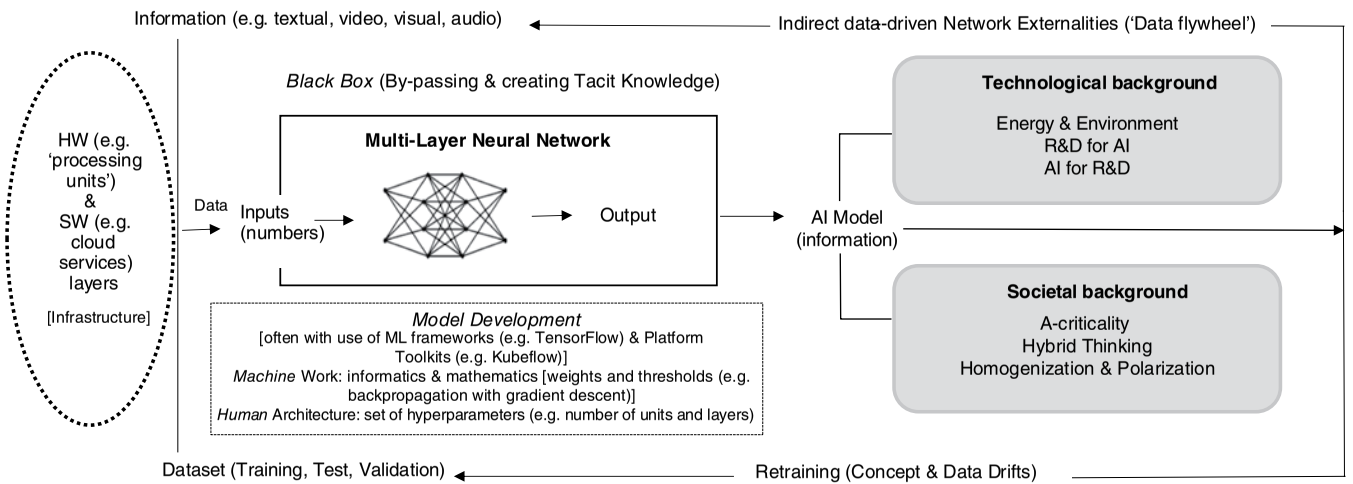


Figure 2.1 Deep learning in context: a socio-technical representation.

any procedure that converts data from a non-metric space into a metric one. Finally, the output layer of a neural network also releases numbers, which are then retransformed into the final output format, converting the numeric result into interpretable information. Thus, in essence, multi-layer neural networks function as digital 'statistical machines', not entirely different from applied econometric models,[5] but with key differences. While econometrics usually assumes a known, predefined functional relationship between inputs and outputs (e.g. errors are assumed to follow certain distributions to ensure diagnostics), neural networks are designed to approximate complex, unknown relationships without requiring a specific form. Each neural network operates as a highly flexible mathematical function, mapping inputs to outputs based on its structure. The arrangement and weights of neurons across layers determine how the network processes information and learns patterns from data, allowing it to handle a wide variety of nonlinear and high-dimensional relationships. The second, and perhaps most notable, difference is that deep learning, by its very nature (and increasingly so as it becomes 'deeper'),[6] is not well-suited for providing meaningful causal inference or offering interpretive analysis, which are the hallmarks of econometric analysis. In other words, apart from possible future advances in so-called explainable AI,[7] neural networks today function very much as very dark 'black boxes' because their internal workings are not easily interpretable. The learned weights (which can be in the order of hundreds of millions or even reach into trillions in the case of the largest models) and

[5] This is certainly not a far-fetched comparison (see, for example, Gans, 2025, pp. 23–24). One could just consider the two major algorithmic types in supervised learning, regression and classification, as essentially extensions of linear and logistic regression, respectively. Furthermore, Igami (2020) shows how three well-known AI applications of their era (including AlphaGo) have clear precursors in the econometric field. Thus, while we are certainly not original in drawing this parallel, it is nevertheless important, in our view, to emphasize once again the connection in this context to somewhat demystify the aura of mystique and supernatural that often surrounds the term 'AI'.

[6] Actually the 'deep' in 'deep learning' is based on the number of 'hidden layers' inserted in the neural network.

[7] Explainable AI (XAI) refers to methods and techniques that provide human-understandable insights into the decision-making processes of AI models. Among these, neuro-symbolic AI figures as a hybrid approach that aims to combine the previously discussed symbolic (logic-based, interpretable) and subsymbolic (neural network-based, data-driven) paradigms.

complex connections in a network make it difficult to explain individual relationships in the data. Thus, 'usage' of the information released by the machine as output, whether it involves predictive or generative power, has, by construction, no big theory behind it, but only big (trillions of) data.

However, the 'blackness' of the box does not mean that humans do not play any role in the functioning of this 'statistical machine' (and here the analogy with applied econometrics may arise again). And we are certainly not just talking about the many Mechanical Turk workers who are often needed to create training or test datasets (e.g. in supervised learning), or for other low-value and repetitive tasks. Indeed, while it is true that the machine *autonomously* finds optimal weights among the many neurons across different layers – using typical techniques such as the backpropagation algorithm with gradient descent over countless iterations of data (often called epochs in the jargon) – specialized experts are still needed to design and implement key aspects of a neural network's architecture. In fact, there are several so-called hyperparameters that are generally set by humans and require some form of human input. The list varies, also depending on the type of multi-layer neural network considered, which is already a human choice among many possible options of different network types and densities. Typically, these hyperparameters include the number of layers, the number of units in each layer, the type of activation function for each layer (e.g. sigmoid, tanh, ReLU, etc.), the choice of loss function (see Terven et al., 2024 for a review of the options in different neural networks).[8] This list may also evolve in the future due to advancements in neural architecture search and meta-learning, which could reduce the need for human expertise in this technical domain. However, it is beyond plausibility to envision a scenario in the near future where human expertise will be *entirely* removed from the equation, in any possible application.[9]

[8] Some of these model's architecture hyperparameters involve some important (economic) trade-offs. For example, *ceteris paribus*, more layers and more units in a neural network raise the possibilities to model the training data with complex functions, but at the same time, increase computational costs (i.e. more weights to be estimated) and increase the risk of overfitting.

[9] Once again, the parallel with applied econometrics can be helpful. Starting from a world composed of linear regressions with few data points, few coefficients, and standard errors that were manually estimated by the

Continuing within our 'socio-technical' representation, the AI's predictive and generative capabilities, embodied in the 'statistical machine', produce digital content information (once retransformed from numbers) that is available for use and action by agents, whether human or not (i.e. robots). Use may take a variety of forms, ranging from predictions that lead to taking certain actions (or even non-actions) to providing full or partial material for the creation of artistic items, as well as serving as inspiration for a new product or service. These *facts* take place in a well-defined (physical or virtual) reality.

And here two different effects can be envisaged.

As to the first, we have already said, that data is the 'new oil that feeds the machine'. As time and usage unfold, new data becomes available and this new data might enable AI systems to perform better, in subsequent future versions of the model. In more technical terms, we can say that in principle, AI-based systems can be strongly characterized by *indirect network externalities* arising from *data-driven learning effects* (Levine and Jain, 2023; Vomberg et al., 2023).[10] The strength of these effects may not be constant over time and across

researcher – where econometric work was done entirely by humans – we have now reached a point where the researcher only needs to press a limited number of 'buttons' to perform complex econometric estimates on a vast amount of data using specialized software. In other words, much of the econometric work is actually done by the 'machine'. However, the researcher still bears the responsibility of knowing which buttons to press, and while these buttons are becoming fewer and fewer, it is difficult to believe in a future where they will disappear altogether. Even admitting a dramatic upsurge in agentic AI (see *infra*) and its eventual diffusion, and excluding scenarios of strong AI, the ultimate responsibility for goal-setting remains a human prerogative.

[10] Beyond the well-recognized virtuous cycle between data and AI models, an equally pivotal dynamic emerges between computing power and AI models, which, in turn, can further amplify the positive feedback loop between data and models. Artificial intelligence is increasingly transforming the design and optimization of chips, driving significant enhancements in their efficiency and performance. These advancements, in turn, enable the development and deployment of even more sophisticated and computationally intensive AI models, creating a self-reinforcing cycle of innovation and progress across both domains. Jensen Huang, (co-founder and) Chief Executive Officer (CEO) of NVIDIA, made this 'computing flywheel' effect very clear in his keynote speech at the 2025 Consumer Electronics Show (CES) in Las Vegas (Huang, 2025). In this regard, it is worth noting that NVIDIA has strengthened its position in the AI chip market partly due to the significant indirect network effects surrounding its GPUs and the specialized software framework used with them (i.e. Compute Unified Device Architecture, CUDA; see Hagiu and Wright, 2025).

different AI domains (for instance, it is reasonable to expect a decreasing marginal utility of data in many contexts). It is also important to note that their existence does not always automatically translate into a competitive advantage for an organization that is able to benefit from them, since competitors in the same arena may also be able to do so (this raises the issue of data accessibility, which we will revisit later). From a technical standpoint, however, these indirect data-driven network effects often exist and give rise to a virtuous cycle, which, in AI jargon, is often referred to as the 'data flywheel'.[11]

The second effect concerns the continuous flow of data, which can be important for mitigating *concept* and *data drift* (e.g. Bauer et al., 2024). (Most current) AI models are inherently fixed and based on past data (even if recent), whereas reality is dynamic and constantly evolving. This mismatch may create concept drift, where the relationships between input features and the target output shift (because reality evolves), or data drift, where the distribution of the input data itself changes over time (because the model enables decisions that change a reality that the model is progressively no longer able to capture properly). Both types of drift lead to degraded performance (in jargon, the system becomes 'brittle', e.g. Mitchell, 2021), as the model no longer aligns with the current environment. From this perspective, the use of synthetic data could be helpful in certain situations but may also increase the risk of model degradation (i.e. 'model collapse', see the analysis on LLMs by Shumailov et al., 2024). Therefore, frequent retraining of models is often necessary to ensure they remain accurate and responsive to real world changes, with data from genuine human interactions being preferable when available.[12]

[11] Qualitative analysis and extracts from interviews with managers in Pedota (2024, p. 16) illustrate well the path-dependent nature of this virtuous circle, and also highlight how better algorithms can also help guide the search for new data or the valorization of old data. Quote from the author: 'The accumulation of data and AI-specific knowledge triggers a positive feedback loop: while data enable crafting more powerful algorithms (thereby improving AI), the accumulation of AI-specific knowledge helps employees maximize the informative content of extant data, predict which data to gather next, and dynamically reconsider old data'.

[12] It should be noted that the discussion here remains necessarily general and panoramic, abstracting from domain-specific techniques that can partly mitigate these issues. For instance, in the case of LLMs, intermediate strategies such as retrieval-augmented generation (RAG) can partially reduce the effects of drift by enriching the model's prompt with more recent or contextually relevant information, even though they do not update the model's parameters.

AI and Knowledge Tacitness

One first important remark about the functioning of this digital 'statistical machine' we are describing – not so much from a technical standpoint as from its social dimension – has to do with epistemology, and in particular with the *tacit component* of knowledge (Polanyi, 1958, 1966). Tacit knowledge includes all those 'types of knowledge that cannot be conveyed effectively by means of language' (Winter, 2016, p. 1).[13]

With the advent of the subsymbolic current as the dominant paradigm in AI, which characterizes today's deep learning, at the expense of the previously dominant symbolic approach, AI is increasingly capable of reducing the relevance of this component in the vast domain of human knowledge. Symbolic AI aimed to translate human knowledge into precise instructions that the 'statistical machine' would operate on. As mentioned, expert systems were expressions of this approach. With subsymbolic AI, it is as if the crucial step in the process of transferring knowledge from humans to machines (typical of expert systems) is somehow bypassed. In fact, AI can encapsulate at least a portion of the experts' tacit knowledge, without the need to make it explicit or explain it. In this way, the mere problem of transferring tacit knowledge from humans to machines is not reduced but is somehow circumvented.[14]

However, AI poses another interesting novelty. As we know, contemporary AI systems 'learn' by adjusting their parameters on the basis of data. But how this learning occurs often remains completely

[13] Often summed up in Polanyi's (1966, p. 4) phrase 'we can know more than we can tell', there are several, albeit fairly convergent, definitions of what tacit knowledge is, and also various categorizations (Collins, 2007; 2010) and possible operationalization (Ambrosini and Bowman, 2001) of the concept in the literature. For the sake of brevity, we have chosen here perhaps one of the most concise definitions, but we believe it is as effective as the others.

[14] To paraphrase the philosopher Andy Clark (1996, p. 60), deep learning may be 'bad at logic' (i.e. it is not well-suited for explicit, symbolically articulated reasoning), but it is exceptionally good at 'playing Frisbee'. In other words, it is good at tasks where it is difficult to define clear, written rules for how to perform the task (see also Mitchell, 2019, p. 35 on this point). While this characterization still largely holds, it is nonetheless interesting to observe that, in the specific domain of LLMs, recent models – including emerging LRMs – partly blur this contrast, as they are increasingly able to generate articulated chains of reasoning or at least seemingly plausible explanations of their outputs.

obscure to humans and generates a new type of tacit knowledge. Thus, in some way, if AI eases the problem of transferring tacit knowledge from humans to machines, it originates the one of transferring tacit knowledge from machines to humans. The point is made clear by Jarrahi et al. (2022):

> Although AI may get things right and provide outputs with high efficiency, AI lacks the ability to explain how the process works – what can be called 'machine tacit knowledge'. This is the knowledge that the machine develops through self-learning but cannot necessarily codify and transfer to humans or other machines (Burrell, 2016; Felzmann et al., 2019). (Jarrahi et al., 2022, pp. 2–3)

This indeterminacy on the part of human beings (even from the AI designer side) in knowing exactly how and why a particular result given by AI occurs (and this, in principle, applies to both predictive and generative AI), raises quite naturally the question of the degree of trust that human beings place in AI solutions (Lam, 2022; Pedota, 2024).

This level of (perceived) trust may vary subjectively between individuals. It is closely linked to the concept of algorithmic, and more generally, AI aversion (Burton et al., 2020; Castelo and Ward, 2021), where something is disliked simply because it is perceived as too 'mechanical', lacking the presence of 'flesh and bones' behind it. Conversely, this may also relate to an opposing tendency of 'hypersuasion' (Floridi, 2024a, 2024b), or 'automation bias' (e.g. Narayanan and Kapoor, 2024; Celi, 2025), where AI is viewed as a kind of infallible digital 'Delphi oracle' (Grilli and Pedota, 2024). Trust may also increase over time as a result of progress in 'explainable AI' research and its applications, even if a small (yet 'only perceived', and not necessarily objectively real) tacit knowledge component of AI is likely to always remain.[15]

[15] A degree of opacity or 'black box' aspect of AI systems could never be completely eliminated due to the inherent complexity of deep learning models. It is widely believed that even with significant advances in XAI, it may not be possible to completely eradicate this opacity (which is why some refer to these methods also as *Greybox* XAI, see e.g. Bennetot et al., 2022). In fact, it is likely that there will always be a trade-off between explainability and performance: making powerful AI models more intuitive to human reasoning by simplifying their operations for the sake of explainability could potentially reduce their effectiveness or accuracy (see, for example, Minh et al., 2022,

Clearly, the relationship between the tacitness of AI knowledge and human trust in it may ultimately have an impact on the generation and pursuit of entrepreneurial ideas, and we will return to this point.

Broader Socio-Technical Background Elements

The analysis of tacit and 'machine tacit' knowledge highlights how AI may not only transform the way knowledge is produced and transmitted, but it may also redefine the boundaries between the material and human agencies. Yet, the epistemological dimension we have just examined represents only one aspect of this broader socio-technical shift. Other factors – both technological and societal – are equally essential to understanding how AI operates and interacts with human systems. While many additional aspects could undoubtedly be identified, those discussed here are selected for their particular relevance to the analysis of innovative entrepreneurship developed in the following chapters. These elements can be initially ascribed to either the technical phase or the social sphere, although, as we shall see, the two levels are deeply intertwined.

Starting with the former, despite its apparent intangibility, it is essential to keep in mind the high computing costs associated with any AI solution and, consequently, the substantial physical infrastructure required to support AI technologies. Beyond the sheer demand for electricity, AI systems also require vast amounts of water for the cooling of data centres and processing units – a resource that is often overlooked in public debates but increasingly critical as global water scarcity intensifies. While future R&D may improve both energy and water efficiency, the overall environmental footprint of AI remains a significant concern (International Energy Agency, 2024) – one whose resolution through technological progress alone remains uncertain (e.g. Bourzac, 2024). In the context of climate change, the balance between costs and benefits ultimately depends on whether AI can itself

their figure 13, at p. 3552). But admittedly, this conservative view is not unanimously shared (see Hassija et al., 2024). In this context, the field has also given rise to a number of interesting and increasingly visible entrepreneurial initiatives, such as the start-ups Bondo AI and Fiddler, which aim to sell explainable AI solutions to existing organizations, or Harmonic, which aims at building mathematical superintelligence for application in science and engineering.

contribute to developing effective mitigation and adaptation solutions. In other words, the energy-intensive nature of AI may prove justifiable only insofar as the technology enables progress towards sustainability goals that outweigh its environmental burden. In this respect, as with other grand societal challenges, AI can indeed be of great help, activating a sort of virtuous cycle. In fact, advances in R&D aimed at improving AI's efficiency and capabilities can make the technology increasingly reliable; in turn, AI can enhance our capacity to innovate across virtually all sectors of the economy, from healthcare and energy to education and manufacturing. Yet this virtuous circle ultimately rests on the assumption that a substantial share of such AI-driven innovations will internalize socially relevant objectives, including environmental ones.

Even without looking too far, and without necessarily contemplating new revolutionary AI solutions, it is a commonly accepted belief that frontier AI systems will experience fewer and fewer occurrences of AI 'hallucinations' or predictive errors over time – that is, instances where AI systems generate false or nonsensical information – alongside improvements in mitigating (human-induced) biases (National Academies of Sciences, Engineering and Medicine, 2024, p. 17).[16] In fact, advancements in data quality and diversity, by enabling training

[16] The issue of errors and hallucinations, in principle, applies to both generative and, to some extent, predictive AI. However, it is undoubtedly more prominently debated in the context of generative AI, whereas in predictive AI it often remains 'under the surface', being less visible or only implicitly acknowledged (e.g. Narayanan and Kapoor, 2024). The greater salience of this issue in the generative domain is likely attributable to the fact that, as will be discussed in more detail later, generative AI, from a technical perspective, does not necessarily prioritize a close alignment with factual reality. Rather, its generative process is fundamentally based on randomization or the probabilistic selection among plausible alternatives. This has led to the emergence of a growing body of literature – often characterized by markedly critical language, particularly in relation to LLMs (e.g. Hicks et al., 2024; Hannigan et al., 2024) – that underscores the risks (also highlighted in the continuation of this section) of excessive user reliance or uncritical acceptance of the system's outputs, also in light of their documented tendency towards *sycophancy* (Naddaf, 2025). As is often the case in the socio-technical research domain, however, some scholars have questioned the more dualistic or overly alarmist tones adopted in parts of this debate (e.g. Gunkel and Coghlan, 2024). In any case, our position here is that it is entirely reasonable to expect that ongoing technological advancements will progressively mitigate many of these concerns across a wide range of applications, both in predictive and generative AI.

on more comprehensive and varied datasets, are continually improving AI models' ability to distinguish between factual information and anomalies or outliers (Gong et al., 2019). AI is becoming increasingly accurate (and will, in all probability, continue to improve in the future) in assessing the reliability of the information it generates or encounters and in reducing biases, where historically, many AI biases have stemmed from unrepresentative or skewed training data (Barocas and Selbst, 2016).[17] Thus, by also increasing the diversity of the data used to train AI systems, the probability of the propagation of biases in AI-generated outcomes is constantly reducing. Additionally, the development of more advanced algorithms and architectures compared to the recent past increasingly incorporates improved error detection and correction mechanisms. These enhancements enable AI systems to self-audit their outputs and make necessary adjustments (Hojjati et al., 2024; Ryu et al., 2024). In this context, agentic AI (which we will discuss in more detail later) can also play an increasingly significant role in mitigating hallucinations (Gosmar and Dahl, 2025). Thus, in summary, with more representative data and more sophisticated techniques – potentially enhanced by future advances in architectures (e.g. liquid and spiking neural networks) and computing power (e.g. quantum computing) – AI is not only becoming more reliable and accurate but will continue to do so in the near future.[18] Thus, the simplistic

[17] As obvious as it may seem, it is nonetheless worth remembering that if biases exist within the society that the data aims to represent, then data representative of that society will produce biased AI models and algorithms (Bauer et al., 2024; Puranam, 2024; Resnik, 2025).

[18] This reflects what we already have, what we are currently experimenting with, and what we will most likely see in the future in the world of LLMs: increasingly 'grounded models' in reality and enhanced thought capabilities with techniques like RAG, often combined with specific prompting techniques (e.g. 'chain of thought'), as in the retrieval-augmented thought (RAT) approach (Wang et al., 2024). Broadly speaking, RAG combines the power of generative models with external knowledge retrieval, generating a response that incorporates the newly retrieved information. It originated from the Meta AI Research group in collaboration with University College London and New York University and was presented in a paper by Lewis et al. (2020). Microsoft enhanced RAG technology by integrating it with knowledge graphs (GraphRAG), enabling it to navigate complex semantic relationships and improve generative AI applications (Edge et al., 2025). Always with Meta among the most active players, we are also witnessing the emergence of so-called Large Concept Models (LCMs), which aim to extend current language models by incorporating more explicit and structured

argument that AI is unreliable due to frequent errors or 'AI hallucinations' can, in our view, be easily dismissed. This, of course, does not mean that AI systems will never make errors (National Academies of Sciences, Engineering and Medicine, 2024) or that their development frontier won't remain somewhat 'jagged' (Mollick, 2024), but errors will certainly become rarer, and the 'jags' smoother.

This, which is undoubtedly good news (and we believe not difficult to foresee), nonetheless entails certain social risks. First, we have already mentioned the a-criticality risk when evaluating AI models' output. Simply put, this risk may increase as the technology becomes more reliable: as AI systems grow more advanced and integrated into various aspects of daily life, people may begin to see them as indisputable digital 'Delphi oracles' across many more domains, leading individuals to become increasingly 'hypersuadable' over time (Floridi, 2024a, 2024b) and pushing our societies towards a form of 'gradual human disempowerment' (Kulveit et al., 2025). This tendency may not only persist but even be reinforced by endogenous economic forces – firms, in particular – since the 'aura of magic' surrounding AI, and its perceived mystique, can significantly enhance adoption and consumer receptivity (Tully et al., 2025). However, 'fewer hallucinations' does not mean 'no hallucinations' (Farquhar et al., 2024). Improvements in the predictive capacity of models encounter physical limits and will never enable AI to shift from statistics to mathematical determinism (Narayanan and Kapoor, 2024). At the same time, several scholars and observers argue that we have long entered an era in which truth is increasingly under threat (Floridi, 2024a, 2024b), with some even referring to it as the 'end of truth' (e.g. Kakutani, 2018; Brill, 2024).[19] Against this backdrop, issues related to deepfakes – already pressing today, and likely to intensify in the absence of robust watermarking technologies and adequate regulatory interventions – further erode individuals' ability to distinguish genuine

conceptual representations, potentially enabling even more effective reasoning, abstraction, and error detection. These developments are closely related to the emergence of LRMs (see Box 2.1), which, despite the documented limitations of early versions (Shojaee et al., 2025), have already shown improvements (Gans, 2025).

[19] The expression has earlier roots, notably in Hayek's 1944 chapter 'The End of Truth' in *The Road to Serfdom*, where he warned about the erosion of shared factual standards – in a very different context, yet with telling resonances for today's AI-driven information environment.

from fabricated information, making the risk of a-critical reliance even more salient. Thus, if users over-rely on AI without critical thinking, they may still make misguided decisions based on AI output (Grilli and Pedota, 2024).

But, this is not the only risk. Another potential significant risk is that a wide range of decisions may be delegated to AI, leading to an impoverishment of critical thinking itself (or 'System 2', to use Kahneman's (2011) terminology). Intelligence, however one chooses to define it, is not something given and immutable over time; it evolves. Our intelligence is not the same as that of humans from the Palaeolithic or the Middle Ages. And all discussions about the possibility of machine intelligence surpassing human intelligence (which, as we know, falls outside the scope of our analysis here) are based on the assumption (often implicit) that machine intelligence will surpass our current human intelligence, without considering that this overtaking might occur, so to speak, by default, because as we become accustomed to using AI, our intellectual abilities may diminish over time (Kulveit et al., 2025). Scientific signs of the 'cognitive debt' resulting from the use of AI technologies – particularly generative ones – are already emerging (e.g. Kosmyna et al., 2025; Melumad and Yun, 2025), and as early as 2012, in reference to a broader context, neuroscientist and psychiatrist Manfred Spitzer was already warning of a growing risk of societal dementia caused by the excessive use of digital technologies from a very young age. Diminished human intellectual abilities would call for even more supplementary functions to be performed by AI, creating a vicious cycle that could lead to potentially bleak outcomes. However, this is far from inevitable. Throughout the history of technology, advancements have generally enhanced our intellectual capabilities, as humans can be understood as 'hybrid thinking systems' composed of the biological brain and the non-biological resources they use as tools (Clark, 2025, p. 1). While AI introduces unique characteristics as a 'tool', particularly in the generative and agentic components, there is no reason to believe it must be an exception. On the contrary, important forms of co-intelligence, as well as hybrid or symbiotic intelligence, could emerge (Zhou et al., 2021; Jarrahi et al., 2022; Mollick, 2024).

Finally, heavy reliance on (potentially the same or similar) AI models, amplified by the network effects mentioned in the previous

section, poses a risk of both increased societal *polarization* and, paradoxically, excessive *homogenization* (Nechushtai et al., 2024) as manifestations of the rising populism in modern societies (Guriev and Papaioannou, 2022). Although the two concepts may seem contradictory, they can coexist (for a similar perspective, albeit in a less general context, see Puranam, 2024, p. 161). Polarization arises as individuals are increasingly funnelled into distinct ideological or interest-based groups, while homogenization occurs within these silos, where exposure to the same content reinforces shared preferences and ideas. This dynamic is already evident in AI systems used for content recommendation (e.g. streaming services), where dominant platforms suggest content based on prior user behaviour and popular trends, creating a feedback loop (Chaney et al., 2018). People are repeatedly exposed to similar ideas and opinions within their group, narrowing perspectives and standardizing consumer tastes, thus fostering both polarization across groups and homogenization within them.

The AI Supply Chains

Following the preceding socio-technical analysis, we now turn to an economic perspective, focusing on the AI supply chain – or, more precisely, on the two main supply chains – that may arise from it: predictive AI and generative AI. Although, as we have seen, both types share the same technological foundation and are central to the socio-technical dynamics and challenges previously discussed, they nonetheless represent distinct economic domains (see Obschonka et al., 2025, on the need for an identical distinction in entrepreneurship studies).

On the AI Predictive Capabilities

Predictive AI refers to the ability to classify data and forecast outcomes based on learned patterns and relationships across large datasets. Practical applications are endless, but certainly one of the most notable examples is undoubtedly the previously mentioned AlphaFold (and its subsequent versions), which predicts the 3D structure of proteins from their amino acid sequences. Since the shape of a protein determines its biological function, such predictions have revolutionized research in medicine, disease analysis, and molecular biology,

providing a far more efficient and cost-effective alternative to traditional experimental methods. For AlphaFold2, Demis Hassabis and John M. Jumper were awarded the Nobel Prize in Chemistry in 2024.

Other important examples align with the broader understanding of 'prediction' as 'forecasting future outcomes'. A key area of application, relevant in our context, involves predicting tastes, moods, and consumer preferences (e.g. Sohrabpour et al., 2021). In this regard, Agrawal et al. (2017, 2018) envisioned a potential transformation of online commerce from the traditional *shopping-then-shipping* model to a *shipping-then-shopping* paradigm. Using Amazon as an example, they describe how increasingly accurate AI-driven predictions could make it more profitable to send consumers what they are expected to want rather than waiting for an explicit order. Even if this scenario may remain futuristic, Amazon's use of big data and deep learning techniques for product recommendations – though not without debate (Bajari et al., 2019) – is well-documented (Brusset et al., 2022; Hagiu and Wright, 2025) and reflected in patents such as the *persona-based data mining system* (US Patent No. 10,157,351 B). Similar capabilities, of course, extend to other AI giants (Ezrachi and Stucke, 2022; Hagiu and Wright, 2025).

In a nutshell, predictive deep learning technologies make production processes more efficient and consumers' (latent) needs, trends, or behaviours more discoverable. Here, we are being deliberately general, as any market and process in which data can be produced, stored and analyzed could be affected by the great predictive power of AI. Given the historical shift in the importance of services relative to manufacturing in modern economies, which is likely to accelerate due to the trend of increasing 'servitization' (e.g. Friesenbichler and Kügler, 2022), it is highly probable that individuals will increasingly purchase fewer manufactured goods, but rather services (possibly in combination with these tangible goods), and that these services will often be digital. Even without referring to the aforementioned growth of services, it is clear that the digitalization of both production processes and the purchase (and consumption) of goods makes the applications of AI models potentially infinite. In fact, we can say that there is no phenomenon that an AI model cannot simulate, provided there is sufficient data to describe it. And from this perspective, given the advancements in all types of sensor technologies, including the IoT sector, it is possible to obtain (real time) data from basically all physical

and social phenomena.[20] This process of *datafication*[21] and the consequent vast availability of data, the 'new oil', will increasingly be able to feed (more and more sophisticated) digital 'statistical machines' in the most diverse fields.[22] Ergo, the socio-technical representation of AI of the previous section will cover more and more economic activities, carrying out its cyclical and predictive action on an ever-expanding variety of sectors.

On the AI Generative Power

Generative AI refers to the ability to create new content by learning and synthesizing patterns from existing information. As already stressed, most contemporary forms of generative AI are based on the very same digital 'statistical machine' we have discussed so far, that is, multi-layer neural networks. In many generative architectures, some form of 'controlled randomness' can be introduced during the process, although this is not strictly necessary (models can also generate outputs deterministically, depending on the decoding strategy). Where and how such randomness is injected, when present, depends on each specific architecture.

In GANs, controlled randomness is embedded in the architecture itself. Comprising two competing neural networks – the generator

[20] As early as 2012, McAfee and Brynjolfsson write, '[e]ach of us is now a walking data generator' (p. 63). This trend has only strengthened since then. As Puranam (2024) puts it, '[d]igital representation means that everything that can be observed is (potentially) data' (p. 23). This logic can even be extended beyond observable phenomena. Where real world sensors fall short, simulation models in virtual reality can step in, producing synthetic data useful for training any neural network (without, however, overlooking the degenerative risks previously mentioned when using data not generated by humans).

[21] *Datafication* is defined as '[t]he imperative to capture all data, from all sources, by any means possible' (Sadowski, 2019, p. 1). On pervasiveness of data in modern life and its ancient roots, see also Alaimo and Kallinikos (2024).

[22] We likely agree with Pope Francis when, in his encyclical *Dilexit Nos*, he states that artificial intelligence will never be able to recreate the nostalgia felt by each individual in remembering the moment as a child when they helped their mother bake a cake. However, it is also true that AI is already capable – and will increasingly become so – of perfectly reconstructing the smells and tastes of every culinary experience (see the *smelling machines* in Smith, 2024; and the *graphine tongues* in Pannone et al., 2024).

and the discriminator, often based on CNNs for image-related tasks – GANs operate through a dynamic adversarial process functioning as a zero-sum game. The generator starts from a random noise vector sampled from a simple probability distribution (typically Gaussian) and learns to transform it into outputs that increasingly resemble the training data. The discriminator, trained on both real and generated data, acts as a critic, distinguishing genuine from synthetic examples. Through iterative feedback and parameter adjustment across multiple epochs, both networks co-evolve: the generator improves its realism, while the discriminator refines its detection accuracy. This continuous interplay – where each network's success depends on the other's failure – drives the system towards producing novel, often highly realistic, content without deviating excessively from the original data distribution.

But this is just one example. Taking as a reference perhaps the other most emblematic architecture in this domain, namely the transformer, controlled randomness, if any, is instead introduced during the token selection process in the output layer. The core of the transformer is the 'self-attention' mechanism, which considers the relevance of each token relative to the entire context. Extending slightly on the title of Vaswani et al.'s seminal paper (2017), *Attention is All You Need*, we might say that 'attention to the accompanying context is all you need'. After processing input data through multiple layers, the model reaches the output layer, where it assigns probabilities to each possible next token based on patterns it has learned. Instead of always choosing the single most likely token (as with 'greedy decoding'), controlled randomness is applied through sampling techniques like top-k sampling, which restricts the choice to the top few tokens, and nucleus sampling, which dynamically selects from tokens within a probability threshold. Temperature scaling further adjusts randomness – lower values yield more deterministic outputs, while higher values foster diversity.

A Representation of the AI Vertical Technology Stacks

In principle, starting with predictive AI, we can say that the supply chain characterizing this capability is generally, by its nature (and in comparative terms), much more 'vertical' and less 'data-hungry'. While, broadly speaking, generative AI is likely more fragmented

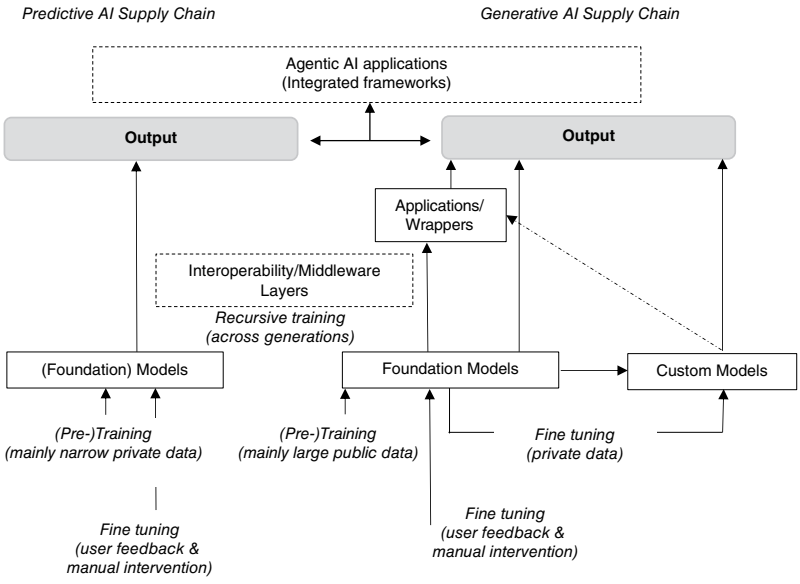


Figure 2.2 AI supply chains. A 'bottom-to-top representation' of the two supply chains is presented here to convey the progressive ingredients and phases involved in building the final AI output.

and has more interconnections, and by nature is much more 'horizontal' and far more 'data-hungry' to achieve acceptable functionality. In Figure 2.2, we provide a stylized representation of the two supply chains, along with some details regarding their functioning.

On the predictive side, the supply chain unfolds in a relatively linear manner, from model pretraining on specific, often proprietary datasets to fine-tuning through user feedback and manual intervention, before reaching domain-specific applications. An interoperability layer may be necessary to connect models to other operational systems (e.g., in credit scoring, where a model estimating risk is embedded within a bank's information technology (IT) infrastructure to support real time lending decisions). Overall, the process relies on comparatively narrow data inputs and a more contained flow of feedback effects.

On the generative side, the chain is broader and more interconnected: foundation models trained on large, often public datasets may subsequently be adapted with private data and deployed through application programme interfaces (APIs), licensing frameworks, and customizable

wrappers.[23] This structure supports many post-deployment feedback loops and indirect network externalities arising at multiple levels – for instance, between foundation/custom models and applications, as in a classical hardware–software paradigm. On the one hand, this makes the generative supply chain more complex, data-intensive, and interdependent, with the interoperability/middleware layer increasingly serving as the enabling ground for agentic AI applications. On the other hand, the relatively greater presence of indirect network externalities stemming from the complementarity between applications and underlying models makes overall *winner-takes-all* (or *most*) configurations certainly not implausible.

LLMs, LMMs, and the Like

A very special category of foundation models in the generative AI stack is certainly represented by LLMs. Rooted in the scientific AI subfield of natural language processing, LLMs are deep learning architectures (based on the transformer) trained on massive corpora of textual data to predict and generate language.[24] The emergence of LLMs, and later of LMMs,[25] marked a fundamental leap in AI's generative capabilities. They can rightly be regarded as the generative backbone of today's AI ecosystem. The best-known examples (all based in the US) include OpenAI's GPT series, Anthropic's Claude family, Google DeepMind's Gemini models, Meta's LLaMA series, and Mistral's Mixtral. Each reflects slightly different training philosophies and degrees of openness: for instance, OpenAI and Anthropic adopt

[23] The term foundation model was popularized by Stanford University, through its extensive monitoring of the AI landscape, to describe a class of large-scale AI systems characterized by broad pre-training and subsequent adaptability across multiple downstream uses. While the concept does not rely on sharply delimited boundaries, it is most readily applicable to the generative-AI stack; nonetheless, it does not, in principle, preclude predictive systems when these function as general-purpose representational infrastructures (see the Stanford *AI Index Reports* for emblematic examples in this respect).

[24] For an insightful yet entirely accessible overview of how they work, see United States Copyright Office (2025).

[25] LMMs represent the natural evolution of LLMs, extending their functionality beyond text to encompass images, video, audio, and code within a unified transformer-based framework. For simplicity and consistency with prevailing usage, in what follows we will refer to LLMs as potentially encompassing also these extended multimodal capabilities.

more closed, API-based access models; while Meta and Mistral follow more open-weight approaches that encourage external fine-tuning and community development. This diversity reiterates the classical tension between proprietary and open innovation models within the generative AI industry – a point to which we will return later in the chapter.

Outside the Western sphere, China has developed a rapidly advancing ecosystem of LLMs, supported by major technology firms and State-driven research initiatives. Among others, prominent examples include Baidu's Ernie Bot, Alibaba's Qwen series, Tencent's Hunyuan, and especially DeepSeek (e.g. R1), developed by DeepSeek AI, which has recently gained international attention for its strong performance, cost efficiency, full open-weight availability, and complete peer review transparency (Nature, 2025). These models reflect China's strategic push for AI self-reliance (i.e. *AI sovereignty*) and its pursuit of domestic alternatives to Western foundation models. While technical details and degrees of openness vary across providers in this case as well, the rise of DeepSeek and its peers underscores the geopolitical dimension of AI innovation – a competition primarily between the US and China, in which the European Union (EU) strives to catch up (see its 'AI Continent' Action Plan of 2025) but still occupies a clearly peripheral position.

This global race for technological leadership might well shape innovation trajectories, and in doing so, the industrial organization of the AI sector. As is typically the case, technological characteristics exert a decisive influence on market structure, and LLMs are no exception. From a technical standpoint, LLMs follow scaling laws – relationships between model size, training data, and performance – that have guided much of the field's rapid progress (Kaplan et al., 2020). Economically, they reinforce the data- and compute-intensive nature of generative AI, consolidating the dominance of a limited number of actors capable of financing large-scale training and maintaining the required infrastructure.[26] The resulting concentration (further amplified by the indirect

[26] On this basis, some analysts have argued that large language and other foundation models may exhibit the economic features of a *natural monopoly*. As noted by Schmid et al. (2024), the combination of very high fixed and sunk costs, strong economies of scale and scope, and limited product differentiation suggests that a single provider could, in principle, supply such models more efficiently than multiple competing firms. The authors, however, caution that despite the presence of such characteristics, the case for regulatory intervention remains, in their view, premature.

network externalities discussed earlier) at the foundation-model level renders smaller AI-based start-ups increasingly dependent on access to these systems, which are likely to remain very few in number (particularly under *winner-takes-all* or *most* dynamics) and accessible primarily through APIs and licensing. In this sense, LLMs function as critical gatekeepers to the broader generative AI supply chain – interfaces where research capabilities, infrastructure, and market power converge.

Concurrently, this concentration of technological and economic power is accompanied by a qualitative evolution in model capabilities. The progressive shift from large language to large multimodal and, more recently, agentic architectures marks a new phase in AI's developmental trajectory. The evolution of LLMs towards greater autonomy and multimodal integration has laid the foundations for the emergence of *agentic AI*, where models are not merely generative but increasingly capable of performing sequences of actions in pursuit of defined goals.

Agentic AI

Foundation models (often LLMs, though not necessarily always) that enable 'reasoning', decision-making, and contextual understanding constitute the technological substrate of agentic AI. This latter category of AI should not be regarded as a subclass of either generative or predictive AI, but rather as a higher-order integration that leverages generative and/or predictive capabilities to autonomously pursue goals. More often than not, it is the orchestration of diverse AI tools – geared towards taking action ('doing') – rather than generation or prediction alone, that defines the distinctive nature of agentic AI. In this regard, architectural frameworks (the middleware/interoperability layer in Figure 2.2) such as Anthropic's Model Context Protocol (MCP), LangGraph, or AutoGPT play a central role in coordinating the internal operations of agentic systems by structuring how goals, memory, planning, and execution modules interact. These frameworks support both vertical AI agents (i.e. domain-specialized agents optimized for specific tasks) and AI agent swarms, where multiple autonomous agents collaborate (or compete) to perform actions.[27]

[27] World Economic Forum (2024) provides a systematic classification of AI agents, tracing their evolution from simple rule-based systems to goal-based, utility-based, and advanced agents integrated with large language

Together, they enable scalable, composable agentic systems that adaptively respond to evolving environments.

Unsurprisingly, given the nascent nature of the phenomenon, multiple definitions of agentic AI exist, each with distinct nuances, including technical ones. A sufficiently general, yet at the same time precise, definition – which we therefore adopt – is the one proposed by Hughes et al. (2025, p. 493): 'AI agents are autonomous technologies, designed to perform tasks, make decisions, and interact with their environment to achieve specific predefined goals, with minimal human intervention'.[28]

Access, Ownership, and Openness in the AI Supply Chains

From the description of the two AI supply chains, it should by now be sufficiently clear how both can easily exhibit significant asymmetries in accessibility and rights among different actors. Both can, in a sense, be represented as successive *layers of accessibility regimes*: who owns or can reach the data (i.e. 'access'), who holds or shares the resulting intellectual artefacts (i.e. 'ownership'), and who can experiment with or adapt the models once they are trained (i.e. 'openness'). Each layer reproduces or amplifies asymmetries between incumbents and new entrants, shaping both market structure and, as far as the scope of our enquiry is concerned, the conditions for entrepreneurial entry.

Data Accessibility and the Dynamics of Competitive Advantage in AI

Besides computing power, we have already discussed extensively the centrality of data in AI. Data to which economic agents may or

and multimodal models. This 'white paper' illustrates this taxonomy through concrete examples – from basic spam filters and smart thermostats to autonomous driving systems – thereby offering a rather comprehensive overview of the different typologies currently available.

[28] Perhaps equally unsurprisingly, it must be noted that the debate over the risks and benefits of agentic AI is already well underway, with positions ranging from enthusiastic to deeply critical – including, paradoxically, the view that such systems should not be developed at all (Mitchell et al., 2025). Analyzing this literature lies beyond the scope of this work; instead, we will limit our discussion to the role of agentic AI within our field of innovative entrepreneurship – fully aware, from our own 'control tower', that this form of AI is already here and likely to stay and develop (barring unforeseeable regulatory interventions).

may not have equal access. This is the first point at which AI-related knowledge becomes distinctive and a potential source of competitive advantage (Bessen et al., 2022; Vomberg et al., 2023). A competitive advantage that by definition is highly *sustainable* over time, because it is continuously alimented throughout time by the action of indirect network externalities arising from data-driven learning effects. In the jargon of the resource-based view of the firm, these are likely to be 'hard-to-imitate resources' from outside the market, and for this reason they can easily present themselves as important barriers to entry.

The issue of data and its accessibility, with all that this entails, is clearly highly relevant across a wide range of fields, from industrial organization (Varian, 2019) to strategic management (Krakowski et al., 2023; Felin and Holweg, 2024), to name just a few. For the reasons outlined earlier, it is also fundamental in our entrepreneurship context. In all these domains, an important factor for understanding competitive dynamics in AI-driven markets is not so much the standard dichotomy between public and proprietary data used to fuel AI models. Rather, the distinction is more nuanced: data may be public in theory but not easily accessible in practice, while proprietary data may range from being purchasable or replicable to being entirely exclusive. Thus, and this should be placed at the forefront of any reasoning about data, it is the *irreplicability* of data derived from *unique* customers or users that can greatly advantage an incumbent or any established organization (see, among others, Mihet et al., 2026). Bessen et al. (2022) make this point particularly well (although their empirical analysis is somewhat more nuanced; see our Chapter 5):

> In the case of AI-enabled products, firms use their initial datasets to train their algorithms. These algorithms produce data as output which becomes an input in the next iteration in the algorithm training process, potentially exacerbating the importance of data inputs on competitive outcomes. Firms with better initial products attract more customers, creating more user data to use as inputs in future production. […] Unlike large technology firms, AI start-ups do not have user-based platforms or other business lines that enable them to collect large amounts of data. Furthermore, even if they did, they might not have the complementary assets necessary to benefit from this additional data (Brynjolfsson and Hitt, 2000). (Bessen et al., 2022, p. 2)

In this context, Hagiu and Wright (2023) and (2025) distinguish two types of data feedback loops: (i) across-user learning and (ii) within-user

learning. In the first case, more data enable a better product for everyone, whereas in the second case, more data improve the product specifically for the individual user generating them. This dynamic makes potential customers more attracted to the incumbent with the highest market share in the first case and/or makes existing users less likely to switch to other providers in the second case. Thus, the combination of both loops can make it extremely difficult for any newcomer to compete with an incumbent that is already deploying its digital 'statistical machine'. At best, an entrant firm – by accessing some hypothetical dataset – would have to outperform the incumbent in (i) to compensate for the user switching costs associated with (ii). As we will see in Chapter 3, even under the most favourable conditions for an innovative venture, this remains a highly challenging scenario in the predictive AI realm.

Ownership and Openness: Intellectual Property and Model Accessibility

Once data are translated into AI models and products, questions of ownership and openness become central to understanding the underlying economic dynamics of the two AI supply chains. Accordingly, IPRs and, more broadly, 'appropriability mechanisms' come into play and must be taken into account.

In predictive AI, most of the economic value is embedded in proprietary datasets, trained model weights, and the knowledge accumulated through iterative experimentation. These are only partially covered by traditional instruments such as copyright or patents, leading firms to rely heavily on *trade secrets* and *contractual protection* to ensure effective 'appropriability'. Openness has historically been limited: firms tend to guard their models closely, since competitive advantage arises from proprietary learning feedback loops and the unique (and irreplicable) data that sustain them.

In generative AI, however, the issue of IPRs becomes more complex on the one hand, and, on the other, the application of 'appropriability mechanisms' considerably more nuanced. Regarding the former aspect, two intertwined problems emerge. Models are typically trained on vast corpora of human-generated content, much of which is protected by copyright. This raises difficult questions about data provenance and the legitimacy of using copyrighted material for training

purposes (see United States Copyright Office, 2025).[29] At the same time, the outputs of generative systems blur the boundary between human and machine authorship, compelling legislators and courts to reconsider what constitutes a creative act and who – or what – can be recognized as its legitimate creator. As for the latter aspect, the classical dichotomy between open and closed innovation models has become increasingly blurred. Along a continuum where full closure and full openness represent only the two extremes, we observe a wide range of intermediate solutions. Firms often adopt hybrid regimes that combine different degrees of openness across several dimensions (e.g. code libraries, APIs, data-sharing practices, etc.). For instance, with regard to LLMs, we have already noted the diverse strategies pursued by major players, which vary across the three key dimensions of model weights, training code, and training data. It is not our purpose here to examine these in detail – many recent contributions already provide systematic analyses (e.g. Jacobides et al., 2021; Srnicek, 2025) – nor to suggest that such strategies are immutable over time or cannot be combined within the same organization,[30] although a certain degree of path dependence clearly exists. What remains constant, and is therefore worth emphasizing in our context, is that these possible graded openness strategies, in the hands of AI giants, appear designed not merely to manage risk but, above all, to shape and protect their respective ecosystems (Wu et al., 2024; Wu and Yang, 2025). In this sense, openness becomes less a normative stance and more a strategic variable in firms' competitive positioning – one that directly affects the degree of exposure of entrepreneurs, and potential entrepreneurs, to third-party control and dependency (see also Fossen et al., 2024).

Thus, seen in the broader perspective of 'access', 'ownership', and 'openness', it becomes evident that the two AI supply chains can each generate specific forms of asymmetry. In the case of predictive AI, the issue of access may easily translate into an *incumbent advantage*,

[29] Emblematic in this regard is The New York Times's 2023 lawsuit against OpenAI and Microsoft, which was soon followed by many others. See also Wu and Yang (2025) for a list of these main lawsuits (updated as of August 2025) and a concise discussion of the nuances underlying the 'fair use' doctrine of copyrighted material.

[30] Google, ranging from Gemini (closed) to Gemma (a 'light' open-weight model family), provides a clear example of this type of mixed strategy for AI models within the same large company.

reinforcing the market position of those already endowed with large datasets and cumulative data-driven learning effects. In the generative AI supply chain, by contrast, the dynamics of ownership and openness may easily cause entrepreneurial activity to *gravitate around the world of the AI giants*. Indeed, both aspects will be thoroughly discussed in what follows. For now, however, we delve into the microfoundations of the AI entrepreneur. Quite naturally, this analysis must begin at the very outset of the entrepreneurial journey.

PART II

The Encounter

3 *Artificial Intelligence and the Entrepreneurial Idea*

Prologue

As we have already said, it all starts from an *entrepreneurial idea*. Possibly the result of a sudden and, moreover, unexpected stroke of genius, or, conversely (and more probably), the result of long, incessant meditation in search of 'some enlightenment', the innovative entrepreneurial idea, if it is to be *truly innovative*, by definition must predict or imagine something that is not there. It can be more 'demand-pull' or 'technology-push', using typical categorizations from innovation economics, and more likely an indistinguishable blend of the two (Mowery and Rosenberg, 1979), yet it is undoubtedly in the ability to 'predict or imagine something that is not already here' that we identify the *prodrome* of the entire entrepreneurial process. Trivial as it may seem, this ability has historically been a purely human prerogative. Moreover, in the pre-AI revolution era, this human capacity was generally amplified by individual human capital. In the domain of our analysis, greater entrepreneurial ideas were on average generated by those who knew something *different*, better or more than others. In other words, at least looking at the past and with a certain (but not excessive) degree of simplification, an innovative entrepreneurial idea, if minimally (ex-ante) meaningful, was rooted in human knowledge.

Does AI change this stylised fact? And if yes, how?

To properly address these and related questions, it is useful to start with the underlying question at the core of all others.

Can AI Be (Truly) Creative?

Briefly Skimming over a Very Recent Body of Literature

How does AI relate to generation of novel material, that is, creativity? To this question, a nascent but already considerable literature

in economics and management, mainly of an experimental nature, attempts to provide an answer. Most of these recent studies take an exploratory approach and rely more or less on one or more variants of comparing (in an experimental setting) the creativity of a given output (e.g. a short story or an artistic object) depending on whether it was produced by generative AI (often abbreviated in jargon as 'gen-AI') alone, by humans with the help of AI, or without the help of AI. The perspective is thus primarily to investigate the extent to which generative AI instruments can be a substitute or a complement to human creativity. Leaving aside the obvious limitations that one or the other study may have, and taking a somewhat panoramic view of the main findings obtained, what seems to emerge quite clearly from these studies is that yes, AI can indeed help the *average* individual to become more creative. But with interesting nuances and distinctions. Without aiming to be exhaustive, and focusing primarily on the earliest contributions in this emerging literature, one of the first studies on the topic is Koivisto and Grassini (2023). In their experiment, participants were asked to generate uncommon and creative uses for everyday objects, and the authors find that, on average, AI chatbots outperform human participants. But at the same time, they find that the best human ideas are comparable to, or even better than, the quality of ideas produced by the chatbots. In a similar vein, but with some notable differences, the findings of Girotra et al. (2023) point perhaps even more decisively to the creative potential of AI. These authors explore the use of LLMs, specifically OpenAI's ChatGPT, in the generation of ideas for new consumer products. The study compares the productivity and quality of ideas generated by ChatGPT with those generated by students at an elite university. The results show that ChatGPT not only generates ideas faster and more cost-effectively, but also generates ideas of higher average quality and with greater variance in quality. Significantly, the majority of high-quality ideas came from ChatGPT, especially when prompted with examples of high-quality ideas. Furthermore, Doshi and Hauser (2024), by analyzing the specific context of short story writing, find that AI enhances creativity, making stories more enjoyable and better written, but at the same time it also reduces the diversity of content, resulting in more similarities between stories, when human creators have access to the same multiple AI-generated ideas (similar findings are also found in the experiment of Anderson et al., 2024, while Nave et al., 2024 reach the very same conclusions by reanalyzing the results

of the experiments conducted by Lee and Chung 2024). The resulting trade-off in the use of AI, on which we will return later on, between the benefits of improved ideas quality and the risk of diminishing their diversity,[1] and hence novelty, cannot, however, be taken for granted, as Lahoti et al. (2023) suggest that LLM-induced homogenization (at various levels, including the creative one) may be easily mitigated with targeted prompting techniques.

In summary, this very recent and still nascent stream of literature seems to suggest an increase in the creative potential of the average individual thanks to AI, but not necessarily a substantial number of AI-generated ideas that would appear to surpass those produced by the most creative human agents. Taken together, these findings seem to indicate a possible rise in the population's average level of creativity enabled by AI technologies, and, at least in some settings, a tendency towards a compression of creative performance across individuals. There is, nevertheless, experimental evidence (e.g. Luchini et al., 2025) indicating that generative-AI tools do not necessarily erase pre-existing differences in creativity, so that the overall impact on the variance of creative ability across individuals remains an open empirical question. This literature is clearly in its infancy, closely aligned with the development of what we consider the current wave of generative AI. Over time, further analysis will clarify whether the results mentioned earlier will be scientifically consolidated and in which direction, and whether new findings will emerge. However, it is worth noting that valuable insights into human-AI interaction in the context of creativity can also be drawn from deep learning technologies that are not necessarily labelled as 'gen-AI'. Put differently, an important point we want to reiterate is that even a predictive AI model can, in certain contexts, display genuinely creative

[1] Hunt and Kurdoglu (2025), drawing also on this body of literature and tracing an insightful parallel, define this potential effect of reduced diversity in entrepreneurial ideas – caused by the massive use of AI tools (in their case, generative ones) – as 'algorithmic deforestation'. Just as the deforestation of a rainforest risks compromising the biodiversity of a biological ecosystem, the widespread use of identical AI tools by prospective entrepreneurs may endanger the biodiversity of ideas that an entrepreneurial ecosystem can express. Although the authors do not present this outcome as inevitable (but rather as a scenario to be avoided), and although we do not see this risk as particularly pressing, given the importance of hybrid intelligence in fostering variety, we nonetheless acknowledge the relevance of the issue, especially from a policy perspective, and will return to it in Chapter 8.

behaviour (and, in doing so, help humans become even more creative). And to illustrate this, there is perhaps no better example than the evidence emerging from comparisons between AI and human performance in intellectually demanding activities such as board games.

AI and Games

There is a long tradition linking the development of a technology to the passing of certain symbolic tests, very often of a sporting nature. In fact, this tradition not only benchmarks progress but also captivates public interest, steering the direction of technological development. One of the first emblematic examples is probably the car industry. In the nascent stages of the automotive industry, the internal combustion engine (ICE) faced stiff competition from electric and steam counterparts. It was also through the crucible of early car races, such as the Paris-Rouen race in 1894 and the Paris-Bordeaux-Paris race in 1895, that the ICE demonstrated its superior reliability, efficiency, and performance. The results of ICE vehicles in these competitions played a pivotal role in cementing their dominance in the automotive sector. Moving to our area of interest, the realm of intellectual games, particularly chess and Go, offers a parallel narrative where symbolic tests have undoubtedly shaped the development of AI.[2] The 1997 match between IBM's Deep Blue and World Chess Champion Garry Kasparov stands as a landmark moment in AI research. Similarly, the 2016 victory of Google DeepMind's AlphaGo over the human Go champion Lee Sedol from South Korea marked a further significant leap in AI capabilities. Given Go's complexity (superior to chess) and the supposedly intuitive grasp required to excel at it, AlphaGo's success underscored the advanced strategic thinking and learning capabilities of AI systems. This achievement not only demonstrated the progress in deep learning and AI but also inspired further innovation and exploration in the field. The story is well narrated by 2017's multi-award-winning documentary *AlphaGo*, directed by Greg Kohs.

[2] There are, however, other notable examples in the history of AI. For instance, AlexNet's victory by Krizhevsky et al. (2012) in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) in 2012 demonstrated the effectiveness of CNNs for image classification. Similarly, the remarkable success of AlphaFold's first version in 2018 in the Critical Assessment of protein Structure Prediction (CASP) initiative (Senior et al., 2020) highlighted AI's potential in structural biology.

In this short film describing the stunning, albeit unexpected, victory of the machine over the human, 4 games to 1, we focus on two aspects. The first is the AI's never-before-seen patterns, which on the one hand confuse the expert player, and later prove to be winning in their total originality. The second relates to the fact that, stimulated by this originality, which appeared naive at first sight, the human being's only victory came about mainly because he decided to 'break the mould' established by the game of Go. The two games above have gone down in history because of 'move 37', which was the machine's victory over the human, and 'move 78', which was the human's partial revenge on the AI, where it is widely believed that 'move 78' would never have arrived if 'move 37' had not come first. There are very interesting excerpts by Metz (2016) in this regard:

> With the 37th move in the match's second game, AlphaGo landed a surprise on the right-hand side of the 19-by-19 board that flummoxed even the world's best Go players, including Lee Sedol. 'That's a very strange move', said one commentator, himself a nine dan Go player, the highest rank there is. 'I thought it was a mistake', said the other. Lee Sedol, after leaving the match room, took nearly fifteen minutes to formulate a response. Fan Hui – the three-time European Go champion who played AlphaGo during a closed-door match in October, losing five games to none – reacted with incredulity. But then, drawing on his experience with AlphaGo – he has played the machine time and again in the five months since October – Fan Hui saw the beauty in this rather unusual move. Indeed, the move turned the course of the game. […]
>
> In Game Four, he [Lee Sedol, author's note] was intent on regaining some pride for himself and the tens of millions who watched the match across the globe. But midway through the game, the Korean's prospects didn't look good. 'Lee Sedol needs to do something special', said one commentator. 'Otherwise, it's just not going to be enough'. But after considering his next move for a good 30 minutes, he delivered something special. It was Move 78, a 'wedge' play in the middle of the board, and it immediately turned the game around. […] Among Go players, the move was dubbed 'God's Touch'. […]
>
> Hassabis [who at the time oversaw the DeepMind Lab, author's note] told me that AlphaGo was unprepared for Lee Sedol's Move 78 because it didn't think that a human would ever play it. Drawing on its months and months of training, it decided there was a one-in-ten-thousand chance of that happening. In the other words: *exactly the same tiny chance that a human would have played AlphaGo's Move 37 in Game Two*. The symmetry of these two moves IS more beautiful than anything else. *One-in-ten-thousand and one-in-ten-thousand*. This is what we should all take away from these astounding seven days. (Metz, 2016)

There are a number of elements that stand out in this emblematic story that are useful in our context.

First, yes, if we follow a fairly general and common sense definition of creativity, *AI can indeed be creative*. Secondly, this creative capacity may depend to a large extent on *being completely free* from (even unconscious) humanly consolidated (unwritten) rules on 'how to do things'. Third, the *human is not 'lost'* in all this. Let us analyse them one by one.

The standard definition of creativity points to the ability to generate *original* and *effective* ideas, where both requirements are necessary (Runco and Jaeger, 2012), and where some authors prefer the expression 'task-appropriate' for the second characteristic (Kaufman and Glăveanu, 2021, p. 17), but basically mean the same thing. Our previous discussion on what AI can perform now and in the next future, the experimental literature mentioned earlier, and obviously also the aforementioned example of AlphaGo, make it clear how AI, even adhering to its (empowered) narrow connotation, can fulfil, let us say with ease, both requirements. To put it simply, 'move 37' was undoubtedly new and equally undoubtedly proved effective or appropriate for its intended purpose.[3]

The creative capacity of AI stems from its ability to break established human patterns and heuristics that are not necessarily conscious and that in some way limit human creativity. Again, the case of Go can help us, in the more recent analysis that Choi et al. (2022) make

[3] Following the 'AlphaGo challenge' story, one might be tempted to conclude that humans can always fully explain the choices made by AI statistical machines, and that the opacity of the 'black box' can, if not always, at least very often, be turned into transparency – that is, made explainable. In this regard, it must be fairly acknowledged that this assumption might appear somewhat valid, but only in well-defined and fixed contexts (albeit highly complex), such as the game of Go presented here, and only recognizable ex post (and never 'in the moment') and only after extensive study. Yet, even in the case at hand, this claim remains highly questionable. Consider, for example, Chouard (2016, p. 3, emphasis mine): 'As DeepMind scientist David Silver explained to the press this January, even though AlphaGo has effectively rediscovered the most subtle concepts of Go (such as sente, moyo, aji or, as we now know, ko fights), its knowledge of these is *implicit* [i.e. tacit, author's note]. The computer doesn't explicitly parse out these concepts – they simply emerge from its statistical comparisons of types of winning board positions at Go. In effect, AlphaGo has a kind of *digital intuition*'. On this latter concept of digital intuition which necessitates 'faith' by humans rather than simple 'trust', see also the editorial of Nature (2016).

of it. In this experimental study, authors focus on the role of AI as an instructor. By analyzing over 750,000 moves in more than 25,000 Go games, they assess the improvement in human gameplay after the debut of an AI-powered Go program (APG). Their findings indicate that training with AI notably enhances the quality of human moves, particularly during the initial phases of a game where uncertainty and unpredictability are at their peak. One of the invoked reasons for this AI superiority in the early stage (p. 25), is that humans often respond to a great bulk of uncertainty by adopting *heuristics* or conventional (opening) strategies, which are not necessarily optimal but merely follow a satisfying logic (Simon, 1956). Provided that computational limits are not severe (and it is likely that, if they exist in a particular context, they will become less and less severe in the next future), AI is freed from a satisficing logic to adopt a maximization logic, which may lead to solutions that may even sound naïve at first sight but, then, turn out to be optimal.

AI, Divergent vs. Convergent Thinking

We have already mentioned in Chapter 1 how creativity often results from the interplay and the right balance of two types of thinking, that is, divergent and convergent thinking. In a recent study of mine (Grilli and Pedota, 2024, pp. 5–6), we argue that AI could alter this balance. In particular, echoing the foundational work of Carr (2003, 2004) and subsequent studies on the commoditization of information technology (e.g. Abonamah et al., 2021; Bronkhorst et al., 2019; Neirotti and Paolucci, 2007), AI could turn convergent thinking into something of a commodity, or at least something that reduces its inherent ability to offer much differentiation between users, thereby increasing the relative importance of divergent thinking by default. As the use of (big) data and algorithmic machine learning becomes widespread in modern economic systems (see the 'datafication' process already recalled), AI systems could provide a process of exploration of 'solutions', free from traditional heuristics and where (as we have seen) humans are often outperformed,[4] more available to everyone.

[4] To some extent, the theoretical analysis of Raisch and Fomina (2025) aligns with this view, asserting that companies may discover more 'distant' solutions to problems when relying solely on autonomous AI search, as opposed to alternative methods (sequential or interactive search) that involve humans

While, if it is true that AI systems could still perform some forms of divergent thinking by their own (think LLMs), working for 'analogies' and 'metaphors' without any human input whatsoever has always proved to be hard in artificial intelligence (Mitchell, 2019). These difficulties somewhat stem from the very essence of AI in relation to the different possible typologies of 'reasoning': mastering *deductive* and *inductive* reasoning while struggling with *abduction* (e.g. Larson, 2021). Unlike deductive reasoning (which derives specific conclusions from general principles) or inductive reasoning (which generalizes from specific observations), abductive reasoning requires proposing novel explanations that extend beyond existing knowledge. Since AI technologies, even generative ones, by definition operate within the boundaries of their training, they cannot truly generate explanations that go too far way beyond their data. That said, recent advances – particularly those discussed in Chapter 2, such as the emergence of LRM-like capabilities – seem to indicate that AI systems are making progress in tasks that approximate abductive reasoning, even if a clear performance gap remains when compared with their inductive and deductive strengths.

In this sense, one can still reasonably affirm that, *de minimis*, humans are not necessarily outperformed in the capability to *spontaneously* integrate knowledge from *vastly* different and previously unrelated domains, and it is precisely this unpredictability and serendipity – often arising from the fusion of disparate and divergent fields – that can be crucial to many creative breakthroughs.[5] Indeed, IT

to varying degrees. It is also worth emphasizing that the perspective on the impact of AI on convergent and divergent thinking presented here – primarily drawn from Grilli and Pedota (2024) – is not at odds with theirs. Raisch and Fomina (2025) define 'distant' solutions in terms of computational breadth, highlighting AI's capacity to explore vast, high-dimensional latent spaces within a given problem domain. In contrast, Grilli and Pedota (2024) emphasize a more radical, cross-domain form of divergence, grounded in human intuition, serendipity, and meaning-making – dimensions that remain beyond the reach of AI's finite representational space. Thus, while AI may easily outperform humans in structured exploration (i.e. identifying 'distant' solutions within a defined domain), humans may retain a comparative advantage – or at least not be entirely outpaced – when it comes to unstructured, associative leaps.

[5] From this perspective, many of the experimental studies we previously cited, which compare AI and human abilities in creative tasks, assume that the comparison takes place at a specific moment and in a given location (*hic et nunc*). While this, of course, poses no problem for generative AI tools,

and AI can certainly support divergent thinking, as seen in studies of serendipity in recommender systems (Fu et al., 2023) and in the proposed 'AI-augmented double diamond framework' (Bouschery et al., 2023). However, humans are not out of the equation, serving both as input providers for AI's divergent thinking and as individuals capable of performing extreme forms of divergent thinking, with or without the help of AI.

Thus, the human being and their creative abilities are certainly not completely obscured by AI, and indeed the former can learn from the latter to free itself from certain established patterns, and like Seedol with the 'God's Touch' (move 78), become more creative. In this regard, there are those who speak of hybrid or symbiotic intelligence to signify the synergetic and complementary aspects that can result from the interaction between humans and AI also in creative terms (Jarrahi et al., 2022; Zhou et al., 2021). At the same time, however, as noted in Chapter 2, literature on algorithmic and AI aversion suggests that these advantages are far from certain and, in general, highlights the complex nature of the interrelation between humans and AI, which may also depend on many contextual factors (Castelo and Ward, 2021; Guerreiro and Loureiro, 2023). In this regard, what we consider legitimate to say is that AI modifies the creativity potential of any economic system, with some creative skills potentially more within the reach of any organizations and individuals.

However, in making this point, it is essential to dispel any potential misunderstanding. What appears to become more widely accessible in an economy permeated by AI is above all the *function* of convergent thinking rather than the ultimate *quality* of the outcomes it generates. The growing diffusion of AI, together with its enhanced predictive and generative capabilities, may indeed broaden access to certain cognitive tools. Yet the actual results produced by these tools remain deeply contingent on contextual elements, including the availability of

which are unaffected by the *hic et nunc*, it is a very different matter for humans, whose 'creativity' cannot be activated on command. Rather, it typically manifests in the most unpredictable ways, often emerging from the most unusual and serendipitous circumstances. Thus, to be fair, it must be acknowledged that in such comparisons (among other potential experimental caveats), the creative performance of humans *vis-à-vis* AI – if biased at all – is certainly downwardly biased.

complementary human capital, the strategic behaviour of the major players shaping the AI ecosystem, and the availability and quality of the underlying data on which these systems depend.

Kirznerian Entrepreneurial Ideas: Under Pressure?

A massive increase in the number of innovative start-ups associated with a technological revolution is certainly not new in economic history; it is, in fact, somewhat expected. Schumpeter (1939) himself based the dynamics of the business cycle precisely on 'clusters of innovations', which often arrive at the same time and often are incorporated into new firms, leading the business cycle into a turbulent period, which in turn leads the economic system into an expansionary phase. What perhaps appears to be new to this AI revolution, compared to previous technological waves, is the potentially limited 'creative destruction' effect inherited with it, at least in terms of the actors involved. This is certainly due to the intrinsic economic characteristics of the technology at hand (e.g. network externalities, economies of scale and sunk nature of many AI investments), the dominant role of AI giants and their corporate strategies (see Chapter 7), but also, to some extent, to the nature of what we may label as 'AI(d) offering' start-ups: new innovative ventures which are created with the primary mission of providing existing businesses with AI-based solutions to enhance their processes. This type of start-up, which we will revisit in Chapter 6, often seems, based on their premises and stated missions, more inclined to reinforce the dominance of large firms in markets rather than challenge it. Clearly, related to this 'incumbency reinforce' action there is also a shrinking of (certain types of) entrepreneurial ideas across all sectors (or in all sectors that make heavy use of AI, which, as we know, will continue to grow). There are a number of ways in which this, in principle, can happen. Basically, all of them can be traced back to the (alleged) ability of AI to learn and, consequently, *predict* and eventually *act*. Here, we now focus primarily on the first aspect, *prediction*, postponing the second one, related to *action*, to Chapter 4, where we will discuss in greater detail the relationship between *agentic AI* and entrepreneurship.

To shed light on this mechanism, it is useful to introduce the distinction commonly found in the literature (e.g. de Jong and Marsili, 2015) between 'Kirznerian' (Kirzner, 1973) and 'Schumpeterian' forms of

entrepreneurship. The former involves entrepreneurial ideas emerging from market disequilibrium, aiming to adjust supply and demand, correct price discrepancies, and bring markets closer to equilibrium. Conversely, the latter refers to entrepreneurial ideas with disruptive potential, where innovation is fundamental and shifts the market away from any existing equilibrium rather than adjusting current conditions. Thus, while the former may originate from pure arbitrage and might involve innovation, though typically of a more incremental nature, the essence of the latter is creation, leading to radical change.

For the reasons outlined earlier, it is fair to say that the former type of entrepreneurial ideas – 'Kirznerian' – could be placed under significant stress by AI. Trivial as it may sound, the better incumbents can anticipate what consumers want or do not want (through for example, unsupervised learning), the smaller the space for new firms. In other words, identifying unmet market needs or new incipient trends as 'outsiders' based on small-scale anecdotal evidence (without the support of big data) may become increasingly rare. In fact, AI may offer an unparalleled advantage in the *systemizing mechanism*.[6] This advantage mainly accrues to incumbents, who have greater opportunities to collect data, greater resources to analyze and systematize it and find patterns (with the help of AI) that are simply beyond the reach of a newcomer (consider, for instance, Amazon's *shipping-then-shopping* model mentioned earlier; see also Li, 2026). The data collected by the incumbent in the sales space (physical or virtual) are, by definition, of a private nature (thus exclusively benefiting the incumbent). The words of Peter Skillman, Global Head of Design at Philips, are emblematic in this respect (emphasis mine):

> We use AI constantly at Philips, I mean *every day*, like we use it to take hundreds of hours of interviews and observations of customers and then we create text-based summaries of this stuff and then we create structured queries to unpack the insights inside of those. (Skillman, 2025)

In contrast, existing public data (e.g. from social media) can be difficult to probe and interpret without an established presence in the market, as such data often come in vast and complex volumes,

[6] Many inventions and creative acts in human history are based on what psychologist Baron-Cohen (2023) defines as the 'systemizing mechanism': the cognitive ability to analyze and construct systems by identifying patterns, rules, and regularities. In this respect, it should sound straightforward the beneficial role of AI in 'systemizing'.

requiring specific expertise and advanced tools for analysis. Moreover, their use may not always be fully legal; for example, data scraping can be explicitly prohibited by platform providers. Although enforcement of such prohibitions is often challenging in practice, the mere presence of a prohibition may serve as a strong deterrent, discouraging potential users from engaging in such practices altogether. In any case, for many aspiring entrepreneurs, the required level of analysis is likely to involve significant upfront investment during the search phase, with evident sunk cost characteristics. This is likely to pose a considerable barrier and placing them at a competitive disadvantage compared to incumbents, who already have established resources and infrastructure. The formal modelling analysis by Norbäck and Persson (2023) strongly supports this idea. In their model, historical (big) data, coupled with AI, enables incumbents to refine operations and better anticipate consumer demands, thereby reducing 'business-stealing' entrepreneurial opportunities (of the Kirznerian type in our framework). The proprietary nature of such data creates a closed-loop advantage for incumbents, allowing them to continuously leverage past data to optimize products, personalize customer experiences, and respond more effectively to market shifts – advantages that are challenging for newcomers to replicate. The financial and logistical challenges of building comparable datasets from scratch place new entrants at a severe disadvantage, particularly in fast-moving markets where speed to market is critical.

As a consequence, we think, start-ups like Dell in the 1980s or Netflix in the 1990s would probably be much more unlikely to emerge in the next future. Michael Dell's story is a perfect example of how (even occasional) interactions with customers and others interested in computers could uncover unmet needs that existing businesses at that time were not addressing (Dell, 1999). Also through direct engagements with customers, and certainly facing some bulk of uncertainty, Dell saw a space to sell computers directly to consumers, bypassing the retail mark-up and providing customization options that weren't widely available in the market at that time. Here, his own words:

> I knew, based on my own experience as a user – and my limited experience with customers – the customers would become even more knowledgeable and demanding every year. I started the business with a simple question: How can we make the process of buying a computer better? The answer

was: sell computers directly to the end customer. Eliminate the reseller's markup and pass those savings on to the customer. It didn't occur to me that others hadn't figured this out. I thought it was pretty obvious. (Dell, 1999, p. 12)

The point is that with the vast amount of data about consumers, the ability of AI to learn from that data and use it to experiment with new solutions and offerings, the likelihood that someone from the outside with 'limited experience with customers' (to use Dell's own words) could get there before the incumbents to identify a latent consumer need will be much lower than in the past. What was obvious to Dell, but not at all obvious to the incumbents, will in all likelihood be reversed: new potential Dells will remain potential because they will not understand, while it will be the incumbents, thanks to big data and AI, who will 'see the obvious'.

Netflix is a different story, perhaps more controversial in its veracity (e.g. Keating, 2014), but it is plausible and illustrates the point equally well (e.g. Tahmincioglu, 2004). In the late 1990s, Reed Hastings was charged a $40 late fee for renting a film (*Apollo 13*) from a traditional video rental store, Blockbuster. Frustrated by this experience, Hastings began to think about the rigid and customer-unfriendly policies of the video rental stores that were the norm at the time. Probably realizing (he did, not the video rental stores) that this frustration was shared by many others, he co-founded Netflix with Marc Randolph in 1997 with a more flexible, customer-centric approach to movie rental (Walker et al., 2017). Our point here, and what connects the two narratives, is that AI will allow companies already operating in the market to learn from their customers, becoming the first to notice what's working, what's not, and what needs improvement, thereby magnifying their 'alertness' (the main entrepreneurial trait identified by Kirzner).

More Efficiency of the Incumbent = Less Entrepreneurial Ideas

There is a long tradition in the history of technology of significant innovations by start-ups (which later became highly successful) that emerged to advance technology and, in doing so, addressed existing inefficiencies in the processes of incumbent companies

and industry leaders of their time. For instance, the founding of Fairchild Semiconductor, already discussed in Chapter 1, played a crucial role in supplying silicon transistors to companies like IBM.[7] While not founded solely to correct IBM's inefficiencies with vacuum tube technology, Fairchild's innovations enabled IBM and others to transition to more efficient and reliable semiconductor technology, profoundly impacting the computing industry. Thus, the equivalence between incumbent inefficiencies and greater entry opportunities for start-ups is certainly not new or unprecedented, and it has been so widely documented and studied that it would take too long to list all the studies on the topic without failing to acknowledge some contributions. But, somewhat surprisingly, the consequential flip side – greater incumbent efficiency resulting in fewer entry opportunities for newcomers – has received relatively less attention and emphasis.

In this respect, our view is that, beyond the enhanced ability of incumbents to recognize latent needs in the market, a second important channel through which AI's predictive power may significantly narrow the space for entrepreneurial ideas lies in the streamlining of incumbents' production processes and, more generally, in the improved management of supply chains. Again, this is not a new phenomenon: it is rooted in the digitalization of companies and their production processes, such as the introduction of supply chain management (SCM) software and enterprise resource planning (ERP) systems. However, with AI, a further leap forward is expected in this regard, with increasingly precise systems for controlling a variety of activities and, more broadly, better management of inventory, capital employed, human resources, and material inputs along the supply chain. How does this affect innovative entrepreneurship?

A concrete example helps me to clarify the concept and also to open up the certainly complex topic of how much this reduction in the domain of entrepreneurial ideas is a good or a bad thing from the point of view of social welfare. Suppose a hospital requests a leading technological university (or if you wish an 'AI(d) offering' start-up) to develop a multi-layer neural network capable of predicting, with high

[7] From the Computer History Museum: 'In early 1958 Fairchild Semiconductor procured its first order, for 100 transistors at $150 apiece from IBM's Federal Systems Division'.

accuracy, the number of occupied beds over a reasonably long time horizon, segmented by department. Let's also assume that the commissioned party succeeds in building such a high-performing AI model – an achievement that cannot be entirely taken for granted (see Narayanan and Kapoor, 2024), but which, given sufficiently large and representative data, an appropriate model design, and the absence of intervening factors causing 'concept drift', is certainly not implausible (e.g. Piccialli et al., 2020). This predictive model will have great benefits for the management of the hospital itself, the scheduling of medical and non-medical staff, the use of diagnostic equipment, and many other activities. In fact the system will be able to predict, with increasingly small margins of error (achieved by subsequent versions), the number of patients segmented by ward (and perhaps by pathology within each ward). As an ancillary benefit, these enhanced predictive possibilities will also allow a better management of meals for patients and staff. In other words, through AI, the hospital will be able to determine exactly how many meals of a particular type will be needed on a given day, minimizing or even eliminating *tout court* food surpluses. And now what is unequivocally a good thing, that is, the increased efficiency of an important medical service hotspot, can become a little less so. In fact, the enhanced predictive power of the hospital will probably turn out as bad news for the local network of soup kitchens which, perhaps through the use of an app created for this purpose by a start-up (a concrete example in this area, among others possible, could be platform Replate), used to obtain supplies from the hospital. In other words, the correction of a productive inefficiency, which as such can only be seen as a positive thing, may eliminate the space for the birth of some prosocial innovative start-ups aimed at exploiting this inefficiency for the social good, leaving the social problem (i.e. feeding the hungry) unresolved. Clearly, one might think that the financial resources freed up as a result of the improvement in the production process (and the resulting cost savings net of the AI investments required to achieve them) could simply be transferred to soup kitchens as donations, with the net effect of having both an improved production process and the same amount of hot meals served to those in need. However, putting aside textbook discussions, even if such an adjustment were to occur (which is far from certain), it would not be immediate and would require time, with the risk that the network might dissipate. That is to say, what an ultra-orthodox neoclassical economist would only judge as fully

positive, that is, a greater production efficiency, could become a little bit trickier and more controversial in a 'general economic equilibrium perspective' applied to the real world, which is often made up of many frictions, at market and government levels.

Broadening the discussion, the increased efficiency of incumbents through AI solutions from 'AI(d) offering' start-ups encourages this type of entrepreneurship but also narrows the space for other (Kirznerian-type) entrepreneurial ideas in the same areas, making it hard to establish a clear net balance, even from a social welfare perspective.

(AI-Predictive Claimed) Efficiency of the Incumbent = (Still) Less Entrepreneurial Ideas

The interesting thing, however, is to note that achieving this reduction in Kirznerian entrepreneurial ideas does not necessarily require the predictive AI technology to be ultra-performant and accurate as in the earlier example. It is sufficient for potential entrants (i.e. start-ups) to merely 'perceive' that this technology is accurate (without being able to actually verify its accuracy) to discourage their entry (even if it would otherwise be profitable). And clearly, incumbent firms have every interest in generating this type of expectation, regardless of the actual quality of the predictive AI tools at their disposal. This is facilitated by the fact that (proprietary and closed-source) predictive models are unverifiable from the outside, and even without bad faith, the metrics used to evaluate their effectiveness can easily create misleading expectations (Narayanan and Kapoor, 2024, pp. 239–241). Clearly, start-ups, not just incumbents, can also exploit these factors to enhance their appeal to stakeholders, for example, potential investors. However, the key difference is that incumbents, already operating in the market, just because of this incumbency, may often have access to private data unavailable to others. Thus, they can leverage this data to build an AI predictive model that provides a significant advantage in deterring market entry. And, as said, this deterrence effect may arise *whether the built-in AI predictive model of the incumbent works effectively or not* (and without even considering potential switching costs for the incumbent's customers, as noted in Chapter 2). In Box 3.1, we illustrate a formal example within a game-theoretical framework, using plausible figures to show how this can happen.

Box 3.1 Entry deterrence of predictive AI

Suppose an incumbent in a given market faces the potential entry of a start-up which is considering the possibility to enter by exploiting unrecognized market needs. However, using its own private data, the incumbent has just invested in AI predictive technologies to counter this threat. The investment is made with the intent of success (i.e. the incumbent aims for it to be effective), but in practice, the model may turn out to have either great accuracy ('good predictive AI') or limited accuracy ('bad predictive AI').

The start-up does not know how effective the incumbent's instrument is; it only has a perceived probability p of its effectiveness. Let E_G and E_B denote the start-up's entry payoffs under good and bad predictive AI, respectively. If the start-up enters and the incumbent's instrument is good, the start-up will incur a negative payoff ($E_G < 0$): it cannot anticipate market needs faster than the incumbent and lacks access to the data needed to develop its own predictive AI model, leaving it unable to compete effectively in this domain. In essence, the start-up becomes trapped in a 'data competitive vicious cycle'. On the other hand, if the incumbent's predictive AI model proves to be of limited usefulness (but this is known with certainty only by the incumbent), the start-up can compete effectively and achieve positive extra-profits ($E_B > 0$). Of course, if the start-up decides not to enter, its payoff will simply be zero. Thus, the decision to enter or not will depend on the expected payoffs the start-up anticipates from the two choices (with the payoff from no entry being certain).

In formulas:

$$\mathbb{E}[\pi_{Enter}] = pE_G + (1-p)E_B$$

$$\pi_{Not\ enter} = 0$$

Thus, a (rational and risk-neutral) start-up will enter only if its perceived probability of the incumbent's AI predictive model being effective is sufficiently low. That is:

$$p \leq \frac{E_B}{|E_G| + E_B}$$

For example, if E_G is equal to – 1 and E_B to 2, the perceived probability of the incumbent's AI predictive model being accurate cannot exceed two-thirds to convince the start-up to enter. Any higher expectation by the start-up that the incumbent can leverage a 'good predictive AI' instrument may rationally deter the start-up from entering the market.

Schumpeterian Entrepreneurial Ideas: Expanding Horizons?

However, the picture provided here remains incomplete, as we must also consider entrepreneurial ideas of a more Schumpeterian nature, and thus, the *creative* power of generative AI. While the predictive function, by definition, identifies something already in the minds of many consumers or embedded within existing production processes, we are now moving towards the capacity of AI to envision something that ontologically does not yet exist, or only exists in the minds of a few. Rather than merely intercepting latent (consumer or business) needs, it is more oriented towards anticipating or even creating them, thereby to some extent inducing new needs that do not yet exist.

In this domain, what appears to us as an indisputable and undeniable stylized fact is that it is difficult, if not impossible, to argue that AI cannot effectively assist humans – whether established organizations or new entrepreneurs – in generating original, innovative business ideas. This applies both to markets with universally trained AIs based on widely available public models and datasets (e.g. LLMs accessible to the general public) and to those with unique, custom-trained AI solutions tailored to specific end users (i.e. specific Apps). In line with the 'productivity paradox' of information technology (Brynjolfsson and Hitt, 1998), one might argue that only the latter has the capacity to generate economic value (see, for example, Felin and Holweg, 2024). While the potential superiority of customized solutions for select users is certainly a possibility, the widespread use of LLMs does not imply, in itself, a uniformity in the perceived benefits or in the differential utility that can be derived from these tools (Ma et al., 2024).

Accordingly, in the context of entrepreneurship, we find it unlikely that universally available AIs will necessarily lead to homogeneous outcomes or be incapable of fostering 'creative' (i.e. Schumpeterian-type) entrepreneurial ideas. A simple experiment, easily replicable by the reader, can illustrate this concept. Innovation, including technological innovation, often has a combinatorial nature, where invention and then innovation emerge from combining and recombining (technological) parts that already exist (the concept dates back to Schumpeter, 1939; see, much more recently, Varian, 2010). Invention itself is also often conceptualized by many as a 'combinatorial search process' (Sorenson and Fleming, 2004, p. 909; see also Weitzman, 1996, among others), a view consistent with the interplay of divergent and convergent thinking in creative endeavours.

Based on this, and grounding on Davidsson and Sufyan (2023), who envisage AI as a general and widespread 'external enabler' of entrepreneurship and entrepreneurial ideas,[8] I thought of three completely different items (literally the first thing that came to my mind) and had the following conversation with ChatGPT, reported in Box 3.2.

Box 3.2 In search of an entrepreneurial idea with ChatGPT

Me: Think of a unique physical product that connects these three things: 1. pencil 2. bread 3. shoes.

ChatGPT: Introducing the 'SoleScribe Snack Shoe' – an inventive, eco-friendly product blending creativity, comfort, and functionality in an unexpected way. Here's how it connects pencils, bread, and shoes:

1. **Pencil Element**: This shoe features a retractable pencil holder cleverly integrated into the side of the sole. With a biodegradable pencil stowed here, wearers can jot down notes or sketches wherever inspiration strikes – perfect for creatives on the go.
2. **Bread-Inspired Comfort**: The interior of the shoe is designed with bio-based, cushiony foam inspired by the softness and resilience of fresh bread. This bread-like padding gives each step a sense of cozy comfort while ensuring breathability, making the shoe comfortable for long wear.
3. **Shoes with a Purpose**: Crafted with eco-friendly materials and durable, slip-resistant soles, these shoes are ideal for both urban adventures and outdoor exploration. Adding a unique sensory touch, they're subtly infused with a fresh, bread-like aroma for warmth and comfort in every step.

The SoleScribe Snack Shoe combines functionality with a playful twist on everyday essentials, making it appealing to those who value creativity and sustainability.

Me: Can you create an image of this shoe?

ChatGPT: Here's the image of the 'SoleScribe Snack Shoe', combining the elements of creativity, comfort, and practicality with a unique design (see Figure 3.1).

[8] See Davidsson and Sufyan (2023, p. 6): '[T]he wealth of insight that AI can provide is impressive. This suggests that in addition to ventures using AI significantly in their products and business models, AI will find widespread use as a "consultant" that aids ideation, evaluation, and decision-making in any kind of venture'.



Figure 3.1 The 'SoleScribe Snack Shoe'. Image generated using DALL·E (OpenAI), 8 November 2024: 'A stylish, eco-friendly shoe called the SoleScribe Snack Shoe, featuring a small retractable pencil holder on the side of the sole […].' A summary of the prompt-based conversation that led to this image is provided in the main text. The grammatical errors and illegible words appear in the original generated image.

Me: Can a start-up based on the sale of this shoe be profitable?

ChatGPT: A start-up focused on the 'SoleScribe Snack Shoe' could potentially be profitable by aligning with market trends and effective business strategies. Key considerations include:

1. **Market Demand:** With its unique design, this product could attract eco-conscious, creative consumers who appreciate novelty. The shoe's sustainability and artistic features set it apart in the crowded footwear market.

2. **Target Market and Pricing:** This concept likely appeals to young professionals, students, and environmentally conscious buyers willing to invest in unique, sustainable products. Testing price sensitivity with these demographics would help determine the optimal price.
3. **Production and Material Costs:** Sustainable materials tend to be costlier, so securing reliable, eco-friendly suppliers and possibly negotiating bulk rates could help manage expenses.
4. **Branding and Marketing:** A compelling brand narrative around sustainability, creativity, and comfort would resonate with the target audience. Crowdfunding could also serve as a tool for raising initial capital and validating consumer interest.
5. **Distribution Strategy:** A direct-to-consumer (DTC) model offers better pricing control, though eco-friendly retail partnerships might expand reach.
6. **Risk Management:** The niche market appeal could present risks due to shifting consumer interests. Testing in select markets and securing sustainable suppliers could help manage these risks.

With strong market positioning and cost management, the 'SoleScribe Snack Shoe' could find profitability by appealing to today's eco-conscious consumers.

Me: Which data would I need to assess the profitability of such a start-up?

ChatGPT: To assess profitability, gather data on:

1. **Market Research:** Demographics, competitor insights, and consumer demand gathered through surveys or focus groups.
2. **Production Costs:** Costs for materials, manufacturing, and testing to ensure quality and feasibility.
3. **Pricing and Revenue Model:** Analysis of price sensitivity and revenue projections based on sales at different price points.
4. **Marketing and Distribution:** Budgets for advertising, promotional activities, and distribution costs across different channels.
5. **Customer Metrics:** Track customer acquisition costs (CAC) and estimate customer lifetime value (CLTV).
6. **Risk and Sensitivity Analysis:** Assess market and supply chain risks and conduct sensitivity analysis to model different sales and cost scenarios.

Analyzing these data points will clarify the potential profitability of the start-up, balancing both optimistic and conservative outlooks.

I chose the model of OpenAI, the most famous and widespread LLM (at least at that time),[9] but it is worth emphasizing the breadth of options in this regard, with many more available LLMs and a wide range of models and applications – built on these foundation models – specialized (with different nuances) in business idea generation (see Legnani et al., 2024 for an analysis of their characteristics).[10]

The product in question (the 'SoleScribe Snack Shoe') may at first glance seem completely implausible,[11] although of course this can never be said with certainty, and after all, in the premises (i.e. the three items chosen very far apart), it was deliberately an 'extreme' case.[12] But the underlying conversation allows us to make some more general considerations.

Looking at the first part of the interaction (i.e. the first and second prompts), the capacity in terms of creativity and convergent thinking (i.e. joining the dots) of AI is evident, which does not necessarily replace

[9] This conversation, originally held on November 8, 2024, with ChatGPT-4o, was 1,118 words long. For editorial clarity, I asked ChatGPT-4o to condense it to approximately 500 words, keeping my prompts and, as closely as possible, the model's original wording. Here is the concise version.

[10] Interestingly, some applications available at the time of writing – such as Startup.AI or Inspiq – also offer the possibility of generating content autonomously without prompts.

[11] But certainly, in my partial defence, it might not even be the most extravagant one, given that the footwear industry lends itself to numerous examples of improbable shoes that (at least online) have enjoyed a certain degree of success. See, for instance, the 'pigeon shoes' (Jones, 2017).

[12] In this regard, we recall the words of the famous Henri Poincaré about mathematical creation, quoted by Weitzman (1996; p. 209): 'Among chosen combinations the most fertile will often be those formed by elements drawn from domains which are far apart. Not that I mean, as sufficing for invention the bringing together of objects as disparate as possible. Most combinations so formed would be entirely sterile. But certain among them, very rare, are the most fruitful of all'. Interestingly, a recent large-scale empirical investigation by Shi and Evans (2023) fully corroborates Poincaré's view, showing that scientific breakthroughs often arise from the convergence of very distant disciplines. Looking at our specific context, we add that it is always easy to identify a successful innovative product ex-post, but it can clearly be very difficult to discern it ex-ante (its eventual success). It is also important to note that LLMs, by their very nature, perform better with more contextualized and longer prompts than the ones I have used here. Using the jargon dear to 'prompt engineering', we could describe it as a 'zero-shot' prompt, very lean and sparse. Although the zero-shot prompt may have its advantages, its brevity is hardly the best prompt technique for an LLM. Thus, the 'extreme' example shown here, which is already interesting (or, if you like, not necessarily stupid) to begin with, could only be improved with a little more effort than I deliberately put into it.

the possibility of the complementarity of human creativity and divergent thinking, where many, if not countless aspects of refining and fine-tuning the business idea are involved (just think of product design, and its meaning, see Verganti, 2017). Thus, the large-scale accessibility and usability of AI technologies (in this case, the LLM of OpenAI, ChatGPT) may commoditize part of the creative process, but it certainly does not exhaust it.[13] Turning our attention to the second part of the interaction (the third and fourth prompts), the generality (if not the vagueness) of the statistical machine's answers emerges somewhat. Though formally indisputable, these answers clearly need to be substantiated to pave the way for entrepreneurial action. This work of substantiation can clearly stem, to a large extent, from the human capital of individuals and the availability of data and analyses in the field, which may certainly not be to everyone's advantage, but which do benefit those – humans or companies – who are already active in the field in some way. To some extent, however, it may also stem from AI technologies, a question we will investigate in Chapter 4. Yet, it is already clear – and can be anticipated – that some level of human intuition (potentially fuelled by specific competencies) appears necessary to gauge the soundness of AI-proposed entrepreneurial ideas and, ultimately, to trust them. In other words, when evaluating the digital content and information produced by the digital 'statistical machine', a certain level of 'absorptive capacity' (Cohen and Levinthal, 1990) on the human side may easily remain essential for effective entrepreneurial action.

A Complex Picture

Apart from the undoubtedly positive and unambiguous impact of AI-based entrepreneurial ideas in supporting existing companies to

[13] Interestingly, after running the example shown here, I discovered that other authors have used ChatGPT to construct examples with more or less similar logic, albeit to make similar but also slightly different points to those I will make here (e.g. Bilgram and Laarman, 2023; Mollick, 2024, pp. 100–101). Similar examples are also increasingly used in class to show students how ChatGPT (and generative AI tools more broadly) can help develop new products (e.g. Ajjan, 2025). This reinforces the idea that the use of LLMs in the way illustrated here is not particularly original; in other words, it is a use that is not particularly 'difficult' to discover, and thus becomes the preserve of many, and may provide many with a useful starting point. But our argument here is that it certainly does not exhaust the creative process inherent in a (viable and potentially successful) entrepreneurial idea.

adopt and leverage this technology (e.g. 'AI(d) offering' entrepreneurship), the picture we have outlined so far is certainly complex. On the one hand, AI's predictive power tends to increase the number of 'AI(d) offering' entrepreneurial ideas, creating opportunities for innovation within AI-based sectors. On the other hand, it may inadvertently suppress entrepreneurial ideas in non-AI sectors (sectors that, however, are or will soon be served by AI, i.e. AI-served sectors, see Chapter 6) by strengthening incumbents through improved demand management and productive efficiency. In turn, AI's creative power is likely to change the relative importance of convergent (i.e. less significant) vs. divergent (i.e. still important) human thinking in creativity, but it is potentially in a complementary (or not necessarily substitutive) relationship with human creative capabilities (for experimental evidence that goes into this direction, see Orwig et al., 2026). Ultimately, when assessing the potential breadth of entrepreneurial ideas that can be generated, a crucial contextual factor comes into play: accessibility of inferential infrastructure (i.e. data, AI hardware and software). This accessibility clearly defines the boundaries of who can effectively participate in leveraging these transformative tools, and clearly institutions may play a central role in shaping this landscape, as we will analyze in later chapters.

As for now, in concluding this chapter, it is essential to highlight that, per se, and hence looking only at their technological core (i.e. treating them for what they are: digital 'statistical machines'), neither predictive nor generative AI models should necessarily diminish the number or quality of business ideas that can be activated in any economic system. On the contrary, we believe it could very well be the opposite. However, *these business ideas may not necessarily appear as entrepreneurial ideas in a strict sense*. At this stage, we can summarize our reasoning in Figure 3.2 by noting that predictive technologies do offer incumbents a certain advantage for Kirznerian-type business ideas, which are likely more incremental in nature. Conversely, when it comes to Schumpeterian-type ideas – those that are potentially more radical and disruptive – the field remains much more open, presenting significant opportunities for both incumbents and entrepreneurs alike, particularly in the realm of generative AI.

But idea generation is just the first (albeit fundamental) step in the entrepreneurial journey, and this brings us to the next crucial question: how could entrepreneurial action be shaped by the involvement of AI

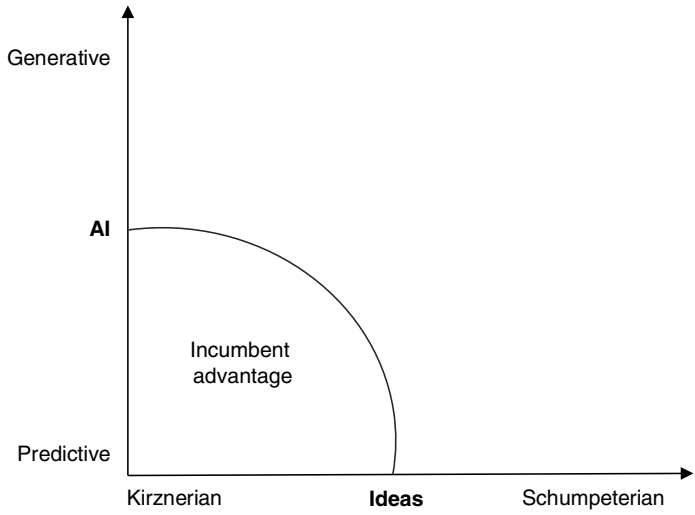


Figure 3.2 Entrepreneurial ideas and AI

in the generation of these underlying entrepreneurial ideas? While AI models are becoming increasingly accessible, potentially making it easier for anyone to formulate innovative ideas (entrepreneurial or otherwise), our simple 'Shoe' experiment above could suggest that humans may remain essential in the implementation phase. But is this really true? Chapter 4 will dig into the realm of realizing the entrepreneurial idea, and what AI implies in this respect. There, we explore what an increase in the materiality of the entrepreneurial idea might mean for its execution, the nuanced interplay between AI and humans during the implementation stage, and how AI can assist entrepreneurs in becoming more entrepreneurial.

4 *Can Artificial Intelligence Be Entrepreneurial?*

Prologue

The area of economics that has, so far, received by far the most attention with respect to the impact of AI is the functioning of the labour market, particularly regarding the potential job losses that the current technological revolution may generate. In this context, it is curious to note that among the professions examined, often broken down into underlying tasks, the *entrepreneur* is never listed (e.g. Felten et al., 2021, 2023; Anelli et al., 2024). Certainly, this may be due to a primary focus on subordinate labour and professional activities, but it is also related to the lack of a specific vocational connotation. The entrepreneur's 'job' is intrinsically elusive and difficult to fit into predefined and standardized job categories. The French word 'entreprende', literally 'to begin something', encompasses far too many aspects.

We have seen from Chapter 3 that it is difficult to say whether AI will lead to more or fewer entrepreneurial ideas in the future. What is certain is that AI is changing not only the 'balance of power' between different types of ideas, but also their very *nature*. In any entrepreneurial idea, there will likely be a greater degree of involvement of the 'material agency' represented by AI than was the case in the recent past. As argued earlier, human agency does not disappear at all in the ideation phase, but will presumably have to share the stage, to some extent, with AI.

Thus, in this chapter, we will investigate how AI interacts with the diverse activities and skills that define the entrepreneur's trade, starting from the premise that the *increased material nature* of the entrepreneurial idea could, in itself, influence the implementation phase. Consequently, our analysis of whether AI might 'erode' the 'entrepreneurial job' must necessarily begin there.

The Disembodiment of the Entrepreneurial Idea

Let us analyze a scenario in which a human can develop a series of potentially interesting entrepreneurial ideas with the assistance of generative artificial intelligence. Now suppose there also exists a multi-layer neural network trained on billions (or even trillions) of data points related to all past start-ups, encompassing their entire lifecycle, including the strategies they pursued and the performance they achieved. This model would thus be capable of prioritizing and accurately classifying entrepreneurial ideas into high-potential and low-potential categories, while also offering a detailed roadmap for transforming high potential into tangible success.

Would such a system fundamentally alter the role of *the human factor* in the entrepreneurial process as we know it? One might argue that the answer hinges on the variety of ideas produced by generative AI and the accuracy of the predictive AI described earlier. While this is certainly true, the key point lies elsewhere. Our point here is that, even if we assume the extreme case of AI's capabilities – encompassing extraordinary variety and near-perfect accuracy – and even admitting the possibility of a reduction in Knightian uncertainty (Knight, 1921) that typically surrounds the entrepreneurial process, *the human entrepreneurial factor* will do remain profoundly important in any case, since Knightian uncertainty is not fully eliminable in the sphere of human activities.

Knightian Uncertainty

Knight's seminal work, *Risk, Uncertainty, and Profit* (1921), introduces a foundational distinction between risk, where the set of relevant future states of nature is known and so, probabilities attached to each possible outcome; and uncertainty, where all the possible states of nature are unknown and/or probabilities are immeasurable. Grounding on this framework, we can identify two distinct forms of Knightian uncertainty (Packard and Clark, 2020).[1] First, we have an *epistemic* uncertainty. Unlike ordinary 'risk', one may

[1] The taxonomy proposed in this study is not the only possible classification of the various types of 'uncertainty' related to the entrepreneurial domain that can be found in the extant literature (for a more general discussion, see Packard et al., 2017).

have a sense of what might happen – such as certain product failures, market responses, or regulatory changes – but have no basis (historical data, stable models, or well-defined distributions) for quantifying how likely any particular scenario is. This is sometimes called 'ambiguity' because the outcomes are recognizable but their likelihood is unclear. Then, we have an *ontological* form of uncertainty, which refers to the inherent unpredictability of the future, driven by the evolution of unknown factors or the emergence of entirely new phenomena. In this case, one does not merely lack probabilities but does not even know what the set of possible outcomes looks like. In other words, the future may hold entirely new events and scenarios that, to put it simply *none* cannot currently conceive. This is a more profound form of Knightian uncertainty – some call it 'fundamental', 'radical' or 'aleatoric' uncertainty. It is easy to say that, if a type of uncertainty can be reduced with the advent of (predictive) AI is more likely to be the former rather than the latter. As a matter of facts, by acquiring more data, improving analytical tools, refining models, or gaining deeper knowledge of underlying processes, it becomes possible to reduce (i.e. *mitigate*, see Packard and Clark, 2020) this (type of) uncertainty: the unknowns can be narrowed down and, if not fully eliminated, at least brought closer to a probabilistic risk scenario. On the other hand, ontological uncertainty stems from the fundamental inability to delineate or even imagine the future space of outcomes. No matter how much data is gathered about the present or the past, it cannot fully prepare decision makers for fundamentally new events that do not fit into existing models. Thus, the very structure of the future environment is unknown, making ontological uncertainty resistant, if not immune, to reduction through traditional means of learning or analysis (i.e. *immitigable* in the words of Packard and Clark, 2020; see also Krause and Hübotter, 2025). Since ontological uncertainty involves futures that transcend current knowledge, even sophisticated AI cannot guarantee accurate foresight into genuinely unprecedented developments. Thus, while it can broaden the imagination and help anticipate a wider range of potential futures, (generative) AI can only partially reduce at its best ontological uncertainty rather than resolving it entirely (Krause and Hübotter, 2025).

This distinction between epistemic and ontological uncertainty has important implications for understanding how AI affects

entrepreneurial dynamics. For the reasons outlined in Chapter 3, incumbent firms, due to their privileged access to private historical data on market functioning, hold a clear advantage over prospective entrepreneurs in reducing epistemic uncertainty. This is again related to the AI-driven reduction in a Kirznerian type of entrepreneurial ideas. However, this advantage does not necessarily extend to ontological uncertainty, which, in principle, can apply regardless of market participation. And it is our deep conviction that as long as (ontological) uncertainty exists within the system, there is room for human entrepreneurship.[2] In other words, there is no other way to navigate (and, in some respects, overcome) that uncertainty except through the action of the entrepreneurial hero, invoking Schumpeter. Thus, not only as we established in the previous chapter, Schumpeterian entrepreneurial ideas are not necessarily diminished by AI, but, in light of the concept of ontological uncertainty, it is also worth emphasizing that human beings, rather than machines, remain the only 'tool' within the capitalist system capable of addressing it. This idea is graphically represented in Figure 4.1. Entrepreneurial action in the area of potentially disruptive innovation should therefore not hinge primarily upon the use of predictive AI, which, by using data, can eliminate or at least reduce uncertainty of an epistemic nature in order to pursue Kirznerian-type ideas. Indeed, this is an area where, all else being equal, incumbents should have a natural advantage, as we also argued in Chapter 3. Instead, the playing field is more likely to be shaped by the triad of generative AI, Schumpeterian ideas, and the need to deal with uncertainty of an ontological nature.

Speaking about generative AI, Townsend et al. (2025) aptly convey this concept when they state (emphasis mine):

> Emerging types of generative AI will play an important role […], but the role of agentic novelty in generating new possibilities involves more than just the random or stochastic 'remixing' of existing data and information. Imagined futures (Ramoglou and McMullen, 2024), entrepreneurial vision (Thiel and Masters, 2014), and profound beliefs (Blank, 2023) all play an important role in entrepreneurial pioneering by human actors. Human

[2] Interestingly, despite the increasing availability of (big) data and advanced computing techniques, available indicators and metrics have shown a rise in global uncertainty rather than its reduction in recent years (e.g. see Ahir et al., 2022).

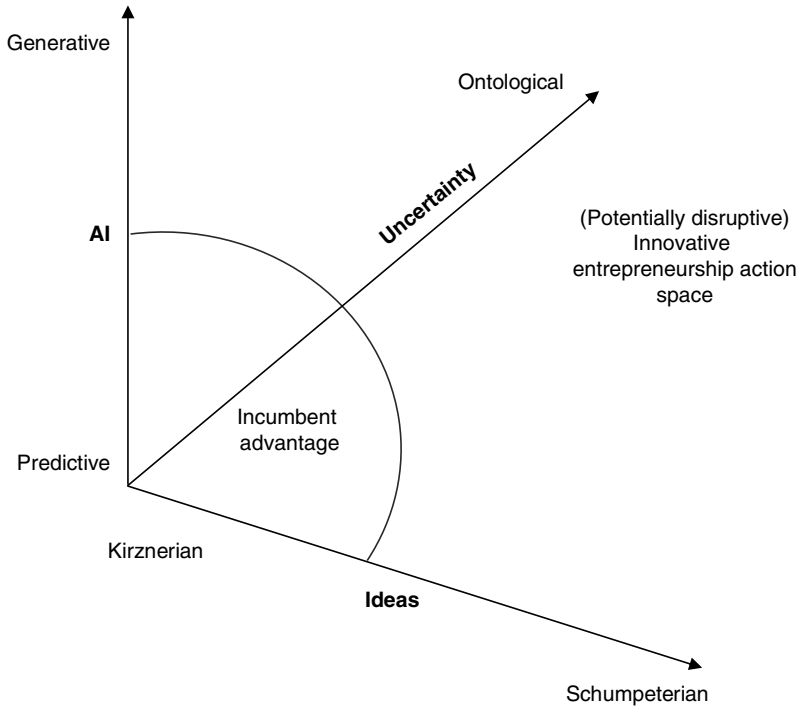


Figure 4.1 Entrepreneurial action and AI

entrepreneurs can imagine and pursue alternatives futures (e.g. Kier and McMullen, 2018), but *existing AI tools are not motivated by a vision of the future*. (Townsend et al., 2025, p. 433)

In other words, even if today – or in the near future – we were to endow AIs with extraordinary capacities for analysis, prediction, foresight and the creation of new possible futures, the motivation to pursue one's vision of the future against all odds, the 'sense of mission' often found in successful entrepreneurs behind revolutionary technological breakthroughs (Dyer et al., 2008; Thiel and Masters, 2014), is something that is unlikely to be transplanted outside the sphere of human agency. Given also the unavoidable presence of errors and hallucinations in the digital 'statistical machine' (see Chapter 2), overcoming what Rady et al. (2026) define as the *ideator's dilemma* – namely, determining how good an AI-generated idea truly is – will always remain, particularly when the idea appears potentially radical, the prerogative of an entrepreneur capable of operating within the realm of ontological uncertainty.

The 'Not Thought Here' Syndrome

If the inability of AI to eliminate Knightian uncertainty preserves a vital space for human entrepreneurship as the driving force of the capitalist system, it is equally true that the increase in material agency in the formulation of an entrepreneurial idea could potentially make its realization in the real world more challenging. First, this difficulty might stem from what could be termed the 'Not Thought Here' syndrome and its impact on the self-efficacy of individuals contemplating entrepreneurship. The inner awareness that one's entrepreneurial idea, no matter how valid and profitable it may seem at the outset, is not entirely 'one's own' can act as a small, initially insignificant nagging thought. The inherent tacitness and inscrutability of AI knowledge may only exacerbate this feeling. Thus, when faced with the first (highly probable, if not inevitable) challenges in developing the entrepreneurial project, this doubt may grow, increasing regret for having 'trusted' AI and started, prompting the search for an exit to abandon the idea. While we lack direct evidence of this phenomenon, extent literature on the formation of start-up founding teams and their subsequent dynamics offers indirect support. The popular press often perpetuates the myth of successful entrepreneurs as lone wolves. However, the reality is that in the realm of innovative entrepreneurship, start-ups are frequently founded by teams of individuals rather than solitary founders. This does not contradict the fact that the initial entrepreneurial idea often (and in our view, almost invariably) has a strong 'solitary' component. In most cases, the initial sparking idea (which can later be redefined and adjusted during the process) originates with a single individual, who then takes on a leadership role and recruits others to join the venture (de Jong et al., 2013; Diakanastasi et al., 2018; Shah et al., 2019; Gray et al., 2024). The 'lead founder' (e.g. de Jong et al., 2013; Shepherd et al., 2021) or 'ringleader' (e.g. Shah et al., 2019), as it is often labelled in the literature:

> [I]s the one who loves it [i.e. the idea, author's note] the most and the one who must inspire the other members of the team. […] Especially in a time of crisis, the leader must be able to stay calm, take decisions, and protect everyone from losing their motivation. If the leader cannot support this role, the most probable thing is that the rest of the team will get really soon demotivated and abandon the endeavour. (Diakanastasi et al., 2018, p. 9)

In other words, a co-founding condition, as opposed to a 'leading condition', should, at least initially, result in less 'emotional' attachment during the process of transforming an idea from conception to exploitation.[3] This, *ceteris paribus*, could lead to a higher likelihood of exiting the entrepreneurial venture when faced with difficulties or obstacles. Similarly, if the 'lead founder' does not perceive the entrepreneurial idea as entirely their own – perhaps due to the sense that the idea was heavily (predictive and/or generative) AI-driven – this relative detachment might also affect the 'lead founder'. To put it simply, due to the role of AI's material agency in shaping and potentially refining the entrepreneurial idea, the 'lead founder' may feel somewhat less like a true 'leader' and a little bit more like a 'co-founder'. In other words, the involvement of AI introduces a certain degree of 'co-foundership' in the leadership position, that would not exist in the absence of AI, along with the associated consequences previously discussed.

AI, the 'Flow' and Founder's Leadership

It is important to note that the detachment process outlined here does not necessarily have to be ontologically true, as it is a purely psychological factor, that is, it occurs by definition in the mind of the entrepreneur. For example, we could consider two entrepreneurial ideas: one where the contribution of AI was objectively greater than the other, but the 'syndrome' only appears in the second case and not in the first, simply because, for some reason, the entrepreneur perceives it that way. This is also strictly related to the concept of *flow* (Csikszentmihalyi, 1990, 1997) and its application in entrepreneurship (Schindehutte et al., 2006; Schilling, 2019) that we introduced in Chapter 1. In this state, external distractions fade away, and the individual experiences a sense of harmony and deep engagement with the challenge at hand. When fully immersed in the intense entrepreneurial activities involved in setting up a start-up – facing fear, fatigue,

[3] The literature on organizational studies identifies several significant cognitive and motivational differences between the two typologies of founders (e.g. Gray et al., 2024), namely *lead founders* and *co-founders*. Ties between these two types can be established following two primary logics: complementary competences and resources in general, and interpersonal attraction. Notably, the combination of these two logics in a dual strategy for forming so-called hybrid ties is considered the optimal approach (Lazar et al., 2022).

and moments of discouragement, but also deriving deep enjoyment and fulfilment from focused concentration on shaping their creation – the perception of AI's materiality can disrupt this flow from the very beginning. Put simply, an entrepreneur might not fully identify with the chosen challenge, perceiving it as something external rather than deeply personal. As Paul Graham (2012) aptly suggests, '[t]he way to get start-up ideas is not to try to think of start-up ideas. It's to look for problems, preferably problems you have yourself'. When a start-up idea becomes more of an abstract or *impersonal* concept rather than a solution to a *personally* meaningful problem – albeit one with broader relevance – the entire implementation process risks being weakened.

This leads us directly to the second reason why an increase in material agency within the entrepreneurial idea can complicate its exploitation. A 'lead founder' who is (even unconsciously) less emotionally attached to their idea and, as a result, less committed, is likely to find it more challenging, *ceteris paribus*, to inspire and emotionally engage other co-founders. This difficulty can cascade, making it harder to motivate employees and other stakeholders critical to the success of the start-up. There are various types and approaches to leadership (e.g. Grint, 2010; Fischer and Sitkin, 2023) that would be too lengthy to detail here and beyond the scope of our analysis. However, it is certain that in the realm of (successful) innovative entrepreneurship, a leader's ability to be fully absorbed in their mission and to be perceived as such by those around them – thereby setting an example – is a critical component for the successful realization of the entrepreneurial project. 'Leading by example' is a characteristic of many leadership styles (i.e. transformational, servant, charismatic), and its importance for involving a team is well-exemplified by the own words of Elon Musk. In an interview with Baron Capital (2022), speaking without the slightest hint of irony (despite the audible chuckles from the audience), he recounts how he slept in the Tesla factory for three years. He explains that, although he had access to a couch, he chose not to use it, opting instead to sleep next to his metal desk so that his team could see him there. Clearly, and needless to say, this decision was not born out of necessity (notably, Musk was already very wealthy at the time) but was a deliberate move to demonstrate his commitment to the mission. Musk wanted his team to witness his full immersion in the work, setting an example during one of the most challenging periods in Tesla's history. While his presence on the factory floor had

practical benefits – ensuring he could address issues immediately – it was, above all, a symbolic act, emphasizing the level of dedication and sacrifice he demanded of himself and, by extension, of others. Without resorting to such extreme forms of example, and this time remaining squarely within the domain of AI, when Michael Truell – co-founder and CEO of the 'unicorn' start-up Cursor, which provides AI-powered integrated development environments for software development – was asked how he maintains the hacker energy within a now well-established company, he replied that he still writes code himself (Y Combinator, 2025).

The Mind of the Entrepreneur

Motivations: Entrepreneurial Self-Efficacy and Need for Achievement

Thus far, we propose that the increased degree of materiality in an AI entrepreneurial idea cannot completely eliminate Knightian uncertainty and, consequently, the relevance of the entrepreneur in capitalist economies. However, we have also noted that the heightened presence of this materiality at the origin of the entrepreneurial process can, in some way, influence individuals' motivations, undermining at the root their self-conviction to become entrepreneurs (or leading them to abandon the path more easily at the first signs of difficulty). Nevertheless, as this is a psychological process and, by its very nature, heterogeneous, it is clear that this effect may characterize some potential entrepreneurs but certainly not all. From this perspective, we believe that the two psychological characteristics, *self-efficacy* and *need for achievement*, already mentioned in Chapter 1, are amplified in their role of countering the 'Not Thought Here' syndrome in this context. Consequently, they are likely to further distinguish successful innovative entrepreneurs in the AI era compared to what has been observed in the past. To understand why, it is worthwhile here to recall Dustin Moskovitz's own words about the best reasons for creating a start-up which we have already mentioned earlier (see again footnote no. 11, Chapter 1). In this case, I emphasize in italics the passages that, in my view, are important to highlight here: 'You can't not do it', meaning that 'you're *super passionate* about this idea. You [think you, author's note] are the *right* person to do it. *You've*

gotta make it happen'. To the extent that the inevitable increase in the materiality of AI-driven entrepreneurial ideas may diminish *human passion* for their realization, the importance of the other two components is magnified. In other words, only if the individual (i.e. a prospective entrepreneur), despite not being the sole architect of the idea, remains deeply convinced that they are the *right* (if not the only) person to make it happen (i.e. self-efficacy) and retains an unrelenting drive to achieve their goal at all costs – '*You've got* (and not others) *to make it happen*' (i.e. need for achievement), can a successful innovative entrepreneurial idea still materialize in markets. Conversely, to the extent that the 'Not Thought Here' syndrome is capable of contaminating the entrepreneur's mind – at the very least diluting or corrupting these two motivational drivers in their essence – the innovative entrepreneurial idea will hardly come to light, and if it does, it will hardly be successful. In this respect, an important distinguishing factor between the two scenarios outlined above – and therefore a key to avoiding this potential 'corruption' of the motivational push – is the presence of a relatively sufficient level of trustworthiness by the prospective entrepreneur in the tacit AI component. In fact, it is clear that only in the presence of a sufficient level of trust in the 'machine' will the entrepreneur refrain from succumbing to doubts about the viability of the entrepreneurial idea in their hand (to which AI has contributed to some extent) and remain steadfast in the face of the inevitable challenges encountered – such as securing the necessary financial and other resources – during the implementation phase of the entrepreneurial project.

It is therefore legitimate to think that AI could impact the nature and motivational identity of innovative entrepreneurs, but it is difficult to imagine institutional contexts where innovative entrepreneurs disappear due to the rise of increased materiality in the entrepreneurial idea – materiality that would harm and not be countered by the motivational aspects of *self-efficacy* and the *need for achievement*. Similarly, however, it is clear that, as a result of this modification, the *changing nature* and *identity* of prospective entrepreneurs could also lead to a change in their *numbers*. And here it is worth emphasizing, in pure Baumol (1990) style, that this number would not be entirely insensitive to the particular institutional context being considered, including the prevailing informal institutions (i.e. cultural aspects) in place.

But there are other, and probably deeper, reasons why the human entrepreneurship component is never destined to disappear completely, especially, we believe, in the innovative context.

Entrepreneurship as Theatre (and Not Cinema): the Entrepreneurial Savoir-Faire

Indeed, in line with the heroic vision of Schumpeter, the activity of an innovative entrepreneur can be understood as an act of engagement, or more precisely, a process of persuasion aimed at securing the buy-in of various stakeholders, from co-founders to investors, from employees to representatives of other organizations, till arriving at even consumers, regarding the validity and potential of the entrepreneurial project (and the embodied product). The entrepreneur, therefore, must also act as a catalyst for resources provided by others (e.g. Botelho et al., 2021). We have already discussed about the importance of entrepreneurs' social capital in Chapter 1 and will further discuss it in the next one in the light of AI advent. This process often involves challenging prevailing opinions or dominant narratives, requiring the entrepreneur to articulate and substantiate their vision in ways that dismantle the Knightian uncertainty we previously discussed. This activity clearly depends on the entrepreneur's capabilities, namely their possession of some human capital (likely specific to the domain of the intended project, especially in scientific and technical fields; see Chapter 5), and that make the others (i.e. potential prospective co-founders, collaborators, investors) think: 'the guy knows what they are talking about'. But this we would say it is a necessary precondition. The 'convincing' process also frequently relies on their 'communicative' ability to inspire confidence, overcome scepticism, dispel doubts and disbelief, and often align diverse interests towards a common goal. In this sense, entrepreneurial activity often goes beyond strictly defined competences (which remain important) to become a profoundly social and communicative act. This communicative dimension embodies a 'theatrical' component which, while not entirely irreplaceable, is highly unlikely to be completely supplanted – despite the significant advancements we might anticipate in digitalization, AI, agentic AI, or the growing trend of remote start-up operations. What we intend with this is that much like the essence of theatre, which relies on the immediacy and authenticity of live performance, entrepreneurship demands *hic et nunc* (i.e.

here and now) moments of interaction that transcend automation or virtual replication. These moments are characterized by an *emotional exchange* – whether between founders, founders and employees, or founders and investors – which, even if potentially reduced in frequency and occurring in specific circumstances, we believe will remain important despite all the aforementioned technological advancements. In fact, we might even argue that it is precisely because of these technological advancements – enabling the real time creation of fully AI-driven multimedia content capable of convincingly 'faking' virtually anything, particularly in a regulatory vacuum and absent robust watermarking technologies – that these moments of human connection will remain (or rebecome) essential. The idea is that one must be physically present in a given place, that is, 'there with their body', in order to *truly understand* something; intelligence is, at times, a matter of 'embodied experience' (Stein and Shollo, 2025), a form of knowledge grounded in a human body moving within a specific physical environment.[4]

And in this regard, it is worthwhile to note that theatre has existed since the dawn of time, surviving every technological transformation in the production and dissemination of content, from the invention of the printing press to the rise of cinema, radio, television, and the internet. Each of these advancements brought new ways to tell stories and engage audiences, yet theatre has maintained its unique appeal by offering something these mediums cannot: the immediacy and intimacy of shared human experience in a physical space. Thus, as we believe, 'theatre' will still perdure in any possible AI world we may envisage (i.e. whatever the future will imply in terms of virtual and/or robotic realities), so will do the 'theatrical' component in entrepreneurship. In-presence *human* interactions and their inherent superior ability to convey passion, enthusiasm,

[4] It is worth noting that we are not the first to draw a parallel between entrepreneurship and theatre. Recently, Chen and Goldstein (2024) explicitly speak of 'the theatre of entrepreneurship' mostly in reference to the communicative rituals and storytelling practices (e.g. pitches) that characterize the entrepreneurial process. When these rituals are merely rhetorical and formal, devoid of substantive content, the theatrical analogy may also carry a negative connotation (see Hartmann et al., 2022). In our context, that of AI, the meaning is precisely the opposite: the very 'theatrical' rituals of entrepreneurship, which only fully make sense when performed live, represent one of the factors that may contribute to keeping entrepreneurial activity alive in the age of AI.

and unwavering confidence will still help entrepreneurs to inspire and persuade others to commit resources to their vision. But clearly, such exchanges hinge on the entrepreneurs' capacity to engage authentically, adapting their delivery to the unique dynamics of the moment and responding to subtle and nonverbal cues. They hinge, at least in part, on what Baron-Cohen (2023) defines as the 'empathy circuit' and Goleman (1995) indicates as 'emotional intelligence'.[5] These latter concepts certainly share some common ground, even if they are not overlapping – at least from an epistemological perspective, with the former rooted in neuroscience and the latter more aligned with the tradition of psychological studies. However, for our purposes, they both converge on the same idea: the ability to navigate and leverage the emotional and social dimensions of human interaction.[6] That is what, inspired by Riggio et al. (2020), we could define as the soft skill of *entrepreneurial savoir-faire*: the abilities *necessary* to the entrepreneur to involve and engage others, act tactfully, and respond successfully in social situations that are critical to the success and growth of their start-up. And in this regard, we claim that it is highly unlikely that AI could ever develop *entrepreneurial savoir-faire* similar to that of humans.

[5] Emotions and intelligence are certainly not at odds with one another, potentially quite the opposite (Minsky, 2006) and rational thinking is by no means necessarily incompatible with emotional responses (Hanoch, 2002). As far as our focus here is concerned, following Baron-Cohen (2023, p. 41), '[t]he Empathy Circuit in the modern human brain has at least two networks: one supporting *cognitive empathy*, defined as the ability to imagine the thoughts and feelings of another person or animal; and the second supporting *affective empathy*, defined as the drive to respond to another's thoughts and feelings with an appropriate emotion'. Emotional intelligence can be defined as 'the capacity to reason about emotions, and of emotions to enhance thinking. It includes the ability to accurately perceive emotions, to access and generate emotions so as to assist thought, to understand emotions and emotional knowledge, and to reflectively regulate emotions so as to promote emotional and intellectual growth' (Mayer et al., 2004, p. 197).

[6] The two concepts we focus on here are in fact only two among many others (e.g., other commonly used ones are 'emotion regulation', 'social intelligence' or, the more general, 'socio-emotional skills'). Indeed, as highlighted by Schmodde and Wehner (2024), there is a multiplicity of concepts and terminologies through which the entrepreneurship literature has analyzed the role of founders' emotions in the entrepreneurial process, resulting in a fragmented account. In this book, we do not aim to cover all the many facets and nuances with which the role of emotions has been addressed in the literature and refer the interested reader to the aforementioned study for a more comprehensive literature review.

There are several reasons to support this claim. One set of reasons relates to what AI experts refer to as the 'lack of theory of mind' in AI (e.g. Mitchell, 2019; Wooldridge, 2020). Admittedly, these reasons are contingent in nature, meaning that while they currently pose obstacles, they represent areas where AI has the potential to improve in the future. In cognitive science, the *theory of mind* (here the seminal contribution is Premack and Woodruff, 1978) can be defined as the ability to attribute mental states – such as beliefs, intents, desires and emotions – to oneself and to others. It encompasses the understanding that others have beliefs, desires, and intentions that differ from one's own perspective. Without the capability of theory of mind, AI cannot genuinely comprehend or predict the motivations, thoughts, and emotions of humans, which are critical for meaningful and empathetic interactions (see again Baron-Cohen, 2023). This limitation significantly hinders AI's effectiveness in contexts that necessitate profound social and emotional engagement, such as entrepreneurial endeavours. As we have already said, the entrepreneurial process inherently involves complex human interactions, where understanding and responding to the emotional states of various stakeholders is paramount. Culturally or socially ambiguous situations, such as innuendos or non-verbal language nuances, are still a minefield for AI, where AI technologies can easily misunderstand or entirely miss these signals. Additionally, AI does not have instinct and 'gut feelings' and the consequent ability to *truly* improvise in response to the unpredictable and often chaotic dynamics of human interactions. Just like in a live theatrical performance, where improvisation and responsiveness to audience stimuli are essential, a (successful) entrepreneur must be able to adapt quickly and creatively to the changing circumstances regarding the *feelings* and *emotions* of the interlocutors.

But we believe there is a deeper reason. At best, AI empathy, no matter how advanced, will always remain, if not 'programmed' in *stricto sensu*, at least *artificial* by definition. Unless one is willing to admit the possibility of AI sentience – a scenario that, as noted earlier (Box 2.2), the scientific community generally considers to be fiction (despite a few prominent voices not entirely dismissing it) – the fact that AI does not genuinely feel or experience what it expresses creates a fundamental misalignment with its human interlocutors. Returning to our example, this disconnect is akin to a hypothetical robotic theatrical performance, where the absence of genuine human presence would undermine the

audience's emotional engagement.[7] A performance featuring only robots pretending to be humans – such as the Hamlet we referenced in Chapter 1 – would likely struggle to evoke the same emotional resonance as the classical Laurence Olivier portrayal. By contrast, in film, the medium allows for the portrayal of characters that are not real people, as the audience can suspend disbelief and engage with purely visual representations. And in this regard, it is worth emphasizing once again, that the entrepreneurial process, at least in some of its most fundamental moments, will remain closer in nature to theatre than to cinema. Thus, despite all possible advancements in the field of so-called 'emotion AI' (often traced back to Picard's seminal 1997 book, *Affective Computing*), an 'artificial' entrepreneur may simulate empathy or emotion (e.g. Goya-Martinez, 2016) or enact 'para-emotions' or 'quasi-emotions' at best (e.g. Stein and Shollo, 2025), but it will always lack *perceived authenticity* and genuine human connection to build trust and inspire meaningful relationships.

The Body of the Entrepreneur

The Anthropomorphization of AI

The introduction into our discourse of robotics brings us to another important consideration: the human, almost inevitable, tendency to 'anthropomorphize' technologies and, more broadly, to attribute human traits to inanimate objects. This inclination, which has ancient

[7] Robotic theatre, defined as the presence of robots on stage interacting either with each other or with human actors, has been a reality well before the current era. It also represents an interesting interdisciplinary field, not only enriching the performing arts but also serving as a fertile ground for research in human-robot interaction (e.g. LePage, 2021). The scientific literature is also not devoid of detailed accounts of theatrical performances performed entirely by robots. For example, Lin et al. (2009) describe a robotic theatre performance that took place at the National Taiwan University of Science and Technology in Taipei in 2008. The article ends with a prophetic statement (p. 6): 'We believe that in the future robot theatres will become one of important intelligent robot applications and great human entertaining activities'. Respectfully, we disagree – at least in not considering it a desirable scenario. However, it must be noted that our discussion does not seek to detract from the field of study we could define as 'emotional robotics' (e.g. Kolling et al., 2016), nor do we intend to imply that humans cannot feel emotions towards robots or that such interactions might not prove useful in certain contexts (e.g. elderly or disability care).

roots, likely arises from various psychological and cultural factors (Guthrie, 2025). Among the prominent explanations is the 'intellectual' one first proposed by David Hume (1739, reprinted in 2000) in *A Treatise of Human Nature*, where he suggests that 'the mind has a great propensity to spread itself on external objects' (p. 112) as part of its effort to understand the nature of these objects. Ultimately, by making unfamiliar things appear more similar to themselves, humans gain the comforting impression of better understanding these objects or phenomena.

In this respect, I think is fair to say that the academic community, particularly within the economic and managerial disciplines, has not been exempt from this tendency to 'anthropomorphize' AI; indeed, it has likely played a significant role in shaping and amplifying it within public discourse. Clearly, the most direct – and to some extent, obvious – way to anthropomorphize AI has been by referring to robots with a human appearance, that is, 'humanoids'. As a purely anecdotal observation, and therefore to be taken with the necessary caution, there has been no scientific conference or academic workshop in the field of economic and managerial studies that I have attended in the last three years (and there have been several, mostly focused on entrepreneurship and AI, for the preparation of this book) where a significant number of scholars, when presenting research related to the *sole* AI, did not include a slide depicting a humanoid. I am sure that everyone can find multiple examples of these 'humanoid forms' in many outlets, ranging from the covers of reports in grey literature to pictures accompanying AI articles in the popular press, and of course, including AI-related content on social media, blogs or newsletters. One of my favourite examples, often (ab)used, is some reinterpretation of Michelangelo's *The Creation of Adam*, focusing on the iconic detail of the two index fingers almost touching. Of course, in all these versions, the hands of God or Adam are not pointing at each other, but at a robotic hand. Although the fresco is undoubtedly so famous that it probably needs no reminder, Figure 4.2 shows an example of what we are talking about, namely the 'AI-robotic' version of Michelangelo's capolavoro generated by artificial intelligence.

There are several possible versions and interpretations that can be given to this image, also depending on how the two hands are positioned, that is, which hand is replaced by the robotic hand, that of God or that of man? One can formulate all possible combinations, and

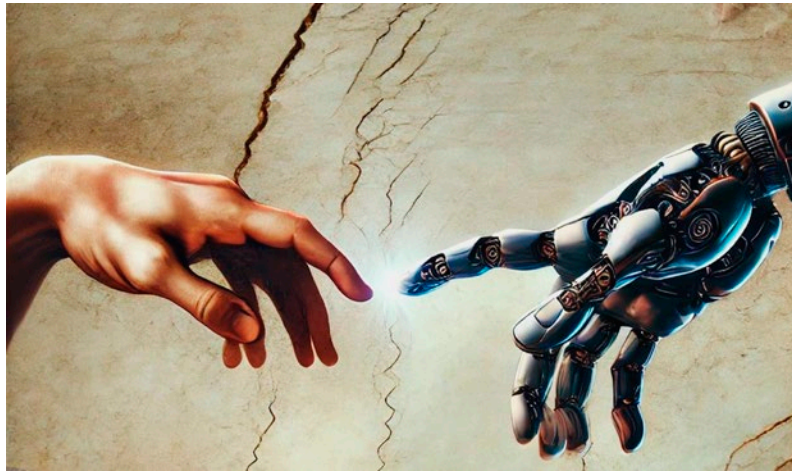


Figure 4.2 The 'AI robotic' version of Michelangelo's *The Creation of Adam*. Image generated using DALL·E (OpenAI), 22 November 2024. The prompt requested 'A photorealistic AI-robotic reinterpretation of Michelangelo's "The Creation of Adam", focusing on the iconic detail of two index fingers almost touching'. The AI tool was left free to choose which hand to replace with the robotic one – God's hand or the human's – and it autonomously chose God's hand.

depending on this, the resulting interpretations can range from more benevolent to less so, but ultimately, and this is our point here, all are dangerous in terms of the misleading message they could convey. For example, the benevolent view is that humans give life to AI (and thus presumably remain in control), while a much more malevolent view sees the hand of a humanized God giving life to AI, just as God did with Adam at the origin of the human species (this was actually the autonomous, i.e. unprompted, choice made by AI in Figure 4.2). But perhaps those who use this image simply want to symbolize humanity's encounter with a robotic AI (to some extent ignoring the origins of the image, which is only part of a more complete fresco in the Sistine Chapel). In any case, the subtle message they may be sending to viewers, even unconsciously, is fundamentally wrong. As we have emphasized in Chapter 2, AI is a digital 'statistical machine'. The robotically revised version of *The Creation of Adam* (and all other possible 'robotic' representations) attributes to AI a *physicality* that AI in principle, does not possess – at least not to the extent that we ascribe

to humans. And while it would certainly be misleading to deny the fact that the two fields, robotics and artificial intelligence, are closely interconnected in many aspects,[8] it is equally misleading, in our opinion, to establish an invariable and unequivocal equivalence between them.

From one side, the translation of the output of an AI process into concrete actions performed by machines in the physical world (the so-called '(generative) Physical AI') is undoubtedly a field that is developing rapidly, it is constantly improving over time (see 'world models'), and it is not implausible to predict that in the future this translation will reach a point where robot movements will perfectly mimic and clone human movements (e.g. Korinek, 2024), possibly even leading, over time, to a mass-market diffusion of robots.[9] However, on the other hand, even acknowledging all this, there are other important factors to consider. First, to put it simply, we are not there yet – robotics is still a long way from replacing humans in many areas (Autor, 2024). The already-mentioned Geoffrey Hinton, the 2024 Nobel Laureate in Physics, referred to 'the plumber' when asked about the 'safest' job in the face of the AI revolution (BBC, 2024), highlighting the challenges robotics faces in keeping up with the pace of AI. Or again, in the various lists provided by the literature on this topic, to which we referred in our prologue to this chapter (e.g. Felten et al., 2021, 2023), occupations that are considered to be less at risk from AI tend to be highly manual occupations or those in which the 'body' is an essential component. Secondly, AI is purely digital, whereas robots, by definition, need to be physically manufactured. This simple (yet significant) distinction implies a completely different cost structure and diffusion pattern.[10] While it is difficult to imagine a world where the

[8] One important interconnection is surely represented by yet another feedback loop between AI and something else – in this case, robotics: AI enables the development and improvement of robotics, whose deployment allows the collection of more and better data, fuelling the virtuous cycle.

[9] Interestingly, these advancements are also driven by (former) start-up initiatives, see, for example, the Polish venture Clone Robotics, the German Agile Robots, the US-based Tangible, or the Norwegian-founded 1X, among (many) others.

[10] Even scholars who believe the arrival of AGI is imminent must concede that, by its very nature, the diffusion of robotics at the societal level will be much more gradual – if not slow (see, for example, Korinek, 2024). According to the same International Federation of Robotics, it could take decades before domestic robots become useful and accepted (BBC, 2026), which seems to us more an antecedent than a consequence of their widespread adoption.

most sophisticated humanoids are ubiquitous (see Autor, 2024), at least in the short to medium run, this scenario is entirely plausible for AI, given its digital nature.

Hikikomori Entrepreneurship: An Oxymoron?

But what does all this reasoning have to do with the innovative entrepreneur and the characteristic of *being entrepreneurial*?

Simple. The profession of the entrepreneur, particularly the innovative one, aided by the increasingly digital nature of their business, has often been erroneously (in our view) 'disembodied' and made almost ethereal in the collective imagination. However, *physicality* often matters in the long, complex, varied, and perilous process that takes the entrepreneur from the entrepreneurial idea to its implementation in the markets. This is probably mostly superfluous to be remarked, if we talk about a nascent innovative business in *manufacturing*, where, to name but one thing, prototypes (and then new products) need to be built, tested and refined and where skilled technicians, engineers, but also – and above all – entrepreneurs must often work side by side. But, to some extent, it can also apply to services, and also to the *digital* world. Certainly, the digital nature of start-ups undoubtedly facilitates a level of flexibility in their location that was unimaginable in past eras (Piva et al., 2011). One can simply consider the fluctuating but always non-negligible share of start-ups emerging from Y Combinator, the world's most prominent accelerator, which are classified as 'remote', meaning they operate without a specific geographic location. However, the fact that innovative start-ups may increasingly transcend physical spaces does not mean that entrepreneurship can completely transcend the need for the entrepreneur to travel, meet people and establish personal connections, despite any possible development in terms of the virtualization of reality that we may envisage for the future in modern societies. The young entrepreneur Pierpaolo D'Odorico, co-founder of Datapizza, a leading Italian community in data science and AI, captures this idea well when, in an interview recounting his (first) experience as an innovative entrepreneur, he basically states that nothing (or very little) in entrepreneurship can be done by 'locking yourself in your room' (Pecora NERA, 2024). We agree, as we have already touched on this theme in the previous pages, both implicitly and quite explicitly, when discussing trust and *entrepreneurial savoir-faire* – activities that

benefit from, or even require, an in-person physical presence. Thus, it seems that we are still far (and fortunately so, I would add) from forms of *hikikomori-style* entrepreneurship. But are we?

Probably yes, especially if we refer to truly ambitious innovative entrepreneurship rather than to 'small-bore' and 'short-haul' ventures. However, broadening the perspective, it is clear that AI can certainly fuel the rise of *solo entrepreneurs* (or in jargon 'solopreneurs') seeking *passive income* – whom we will define as 'passive entrepreneurs' (this may indeed sound as a true oxymoron indeed) – where the active role is played by an artificial avatar, with the help of agentic AI.

Agentic AI Entrepreneurship

As we have seen from Chapter 2, by agentic AI, we refer to artificial intelligence systems that do not merely assist humans in executing predefined tasks but can autonomously plan, strategize, and execute actions in pursuit of specific goals. Unlike traditional AI, which operates reactively based on user input, agentic AI systems exhibit a degree of independence, capable of making decisions, adapting to changing environments, and even optimizing business strategies without continuous human intervention.

In our analysis, agentic AI and innovative entrepreneurship are interconnected on multiple levels. Let's examine them one by one.

The first, almost too obvious to mention, is that agentic AI represents a significant and evident opportunity for entrepreneurship within what we have defined as 'AI(d) offering' entrepreneurship, by intelligently automating key processes that are valuable for both consumers and/or other firms, including incumbents (thus, even in this field, we can say that some start-ups may emerge with the primary goal of strengthening incumbents rather than directly challenging them). These innovative start-ups may leverage agentic AI to optimize operations and introduce intelligent automation within organizations with the ultimate aim of making them more efficient. Innovative ventures such as Aisera, co-founded in 2017 by the serial entrepreneurs Muddu Sudhakar and Christos Tryfonas (both endowed with high levels of specific human capital), have specialized in AI-powered customer service automation. Or Ema ('Enterprise Machine Assistant'), (co-)founded by Surojit Chatterjee, a former executive at Google, has enabled organizations to deploy AI employees – either customized

or pre-built – capable of operating across a wide range of functions, including highly specialized ones. These two represent notable examples, among many others, of this emerging class of agentic 'AI(d) offering' start-ups.

A second level, which we mentioned earlier, presents a more profound and potentially transformative implication: if taken to its extreme, it could challenge – at least seemingly – the very role of the entrepreneur within the capitalist system. To put it simply, if the entrepreneur has just to merely oversee an AI-driven venture that autonomously seeks and exploits market opportunities, is this still entrepreneurship as we know it? At the very least, the seminal distinction (whose importance was among the first to be emphasized by Schumpeter, 1911/1934) between entrepreneurs and investors would become increasingly blurred. Investors, after all, allocate capital to ventures without necessarily engaging in their operational complexities. Passive entrepreneurs, as we have defined them, might operate similarly, but instead of capital, they allocate AI-driven labour – delegating strategic, managerial, and even creative functions to artificial agents. To take this argument further, one could argue that AI-driven entrepreneurship might resemble a new form of franchizing, where instead of hiring employees or subcontracting tasks, passive entrepreneurs deploy AI agents to build and scale businesses with minimal hands-on intervention. Needless to say, this would raise the fundamental question of whether entrepreneurship can be reduced to ownership or whether, instead, it necessarily entails active involvement. Although ownership per se does entail important competencies (Foss et al., 2021) – to the point that we could even, perhaps, speak of entrepreneurship in such a case – we must acknowledge that what we have described above is indeed an extreme scenario. The deployment of such a system, at least to some extent and particularly in its initial phase, would still require human expertise and financial capital. But even if we admit for a moment that such a scenario is somehow realizable (in fact, from the perspective we have adopted, it cannot be fully relegated to the realm of fantasy), it would not eliminate the fundamental role of entrepreneurship in its more traditional, active form. It would be a complementary form, but certainly not a substitutive one, especially regarding the most innovative forms of entrepreneurship.

In fact, this brings us back to the embodied nature of entrepreneurship discussed earlier. Knightian uncertainty will never completely

disappear, and even in an era of increasing digitalization and agentic AI automation, the role of serendipitous networking, intuitive leaps, negotiation dynamics, the ability to inspire and lead human teams, personal interactions, trust-building, and physical presence are very likely to remain critical and very much human prerogatives – for the reasons we have exposed so far. Moreover, AI-driven ventures might face credibility and trust challenges in many domains, where potential customers, partners, and investors may hesitate to engage with a (nascent by definition) business that lacks a clearly identifiable, human decision maker. Thus, even the most advanced AI-driven entrepreneurship models, may still require human entrepreneurs to step in at key moments – reinforcing the argument that entrepreneurship cannot be entirely disembodied. Hence, while an expansion of AI-driven passive entrepreneurship is not excludable, the archetype of the engaged and dynamic entrepreneur – actively navigating complexity and uncertainty – remains central to the entrepreneurial process.

The question, then, is not whether agentic AI will replace entrepreneurs, but rather: what kind of entrepreneurs will thrive in an AI-driven world? This will be explored more thoroughly in Chapter 5, but for now, the issue allows us to introduce our third consideration regarding the relationship between agentic AI and entrepreneurship: the ability of agentic AI to profoundly assist prospective entrepreneurs in a wide range of activities related to the implementation stage. In this context, AI technologies may function as true 'entrepreneurial buddies'.

The AI Entrepreneurial Buddy

Indeed, the vast number and variety of AI-based applications and tools flooding the market, especially those of a generative and agentic nature, not only represent the typical entrepreneurial ferment that accompanies technological revolutions, which we have largely defined as 'AI(d) offerings', but also, beyond their adoption by consumers and established organizations, can provide significant support to nascent ventures, that is, start-ups.

As mentioned in Chapter 1, innovative entrepreneurship is typically a resource-constrained activity (Botelho et al., 2021), where the ability to make the right choices (by founders and active investors) is crucial to avoid dissipating scarce resources in too many directions (e.g. Gans

et al., 2019). Regardless of which entrepreneurial theory we choose to reference – be it bricolage or effectuation, previously cited, among the most prominent – or which perspective we adopt on the entrepreneurial process, from the Lean approach to the scientific method (e.g. Camuffo et al., 2020), AI can provide valuable support according to all these frameworks. AI technologies either alleviate the typical constraints of a start-up (of various kinds, from the entrepreneur's time to financial resources), expand the range of possible actions, reduce the sunk costs associated with decision-making, or a combination of all these factors. In fact, there is virtually no aspect of implementing an entrepreneurial idea that cannot benefit in some way from this vast array of AI tools. It is probably impossible to draw up an exhaustive list of activities, given also the multifaceted nature of the entrepreneur's 'job' that we evoked in the prologue to this chapter. Looking only at the typical dimensions, the benefit could range from product development and prototyping to marketing strategy, from customer engagement and care to operations and workflow automation, from data analytics and business intelligence to patent search and intellectual property protection.[11]

Just to give an example, in Box 4.1, we revisit the entrepreneurial idea of the 'SoleScribe Snack Shoe' presented in Box 3.2 to illustrate more concretely how a wide range of AI tools (among countless possibilities) could help a hypothetical entrepreneur further develop this seemingly bizarre business idea.

This array of functions, which can therefore be outsourced – sometimes even fully autonomously and automatically – to (agentic) AI technologies, represent *true partners*. Some see them as a kind of 'AI workers', while others go so far as to consider them close to being 'co-founders' (e.g. Murtinu and De Massis, 2025) or 'cybernetic

[11] As to this latter aspect, one fundamental, yet perhaps overlooked and certainly underexplored, dimension of entrepreneurship from a scientific perspective is the significant number of innovative start-ups that struggle – sometimes to the point of failure – due to 'contractual incompleteness'. Some time ago, I gathered a testimony from a prominent legal consultant who, with evident regret, pointed out how, in principle, otherwise solid and promising start-up projects were put at risk by the presence of fatal gaps – sometimes surprisingly trivial ones – in their contractual relationships with business customers. These gaps, often involving much larger counterparties, are frequently exploited opportunistically to the great disadvantage of innovative start-ups. On a more general level, a paradigmatic example of the potential contractual power in the domain of IPRs, exerted by a Big Tech company over an innovative start-up (at the time) – to the detriment of the latter – is found in Reback's (2002) account of the relationship between IBM and Sun Microsystems in the early 1980s.

Box 4.1 Implementing the entrepreneurial idea with AI

Like all entrepreneurial ideas (but perhaps even more than most), the implementation of the 'SoleScribe Snack Shoe' will inevitably require a structured approach across multiple dimensions to evaluate the market appeal of the product and a great bulk of experimentation to improve it. In fact, a prospective entrepreneur may want to further refine the idea, its design and experiment with it, in an iterative process.

ChatGPT, in Box 3.2, identified six activities as essential for gauging the potential profitability of the 'SoleScribe Snack Shoe' and six others as necessary for gathering data to assess this profitability, resulting in a total of twelve different dimensions. One could use this framework as a starting point (the analysis that follows was conducted in January 2025).

By outlining the key aspects and selecting the most relevant dimensions among these twelve – particularly those where AI may offer, in our view, the most strategic value – we can focus on the areas that may provide the greatest support at limited costs. For example, in order to better infer market demand and consumer target, one may wish to use some AI-based design applications (e.g. Vizcom) to create sketches of different (3D) versions of the shoe, for then testing their market through large 'in silico' samples of individuals (i.e. potential customers) that could be easily created and surveyed with the help of LLMs. Obviously not without limitations (but equally, of course, subject to improvement over time), LLMs are generally found to provide a sufficiently good proxy for human behaviour in many consolidated experiments (e.g. Argyle et al., 2023; Horton, 2023; Tranchero et al., 2024; Mei et al., 2024). They also offer undeniable advantages in being easily 'modifiable' and 'stratifiable' as needed, as well as in allowing for questions that, in a real world context, would be difficult to ask – either because they are considered inappropriate or impractical. Thus, insights from these surveys analyses could also be triangulated rather easily with AI-powered trend analysis tools (e.g. Exploding Topics, Google Trends, or the relatively more expensive Brandwatch) help assessing whether there is genuine interest in an eco-friendly, multifunctional shoe. All of this is unlikely to serve as a perfect substitute for high-quality information gathered directly in the field (through direct market experience), but it can certainly represent an acceptable surrogate.

Then, production and material sourcing could benefit from AI-based procurement platforms (e.g. Fairmarkit) which may help the entrepreneur (a) identify and evaluate suppliers that meet specific criteria, such as sustainability and cost-effectiveness, which are essential for producing an eco-friendly product; (b) streamline the sourcing process with some forms of automation through agentic AI.

Continuing our entrepreneurial journey, AI-assisted branding and copywriting platforms may help create a convincing narrative around

the 'SoleScribe Snack Shoe' (e.g. Jasper AI or Copy.ai), emphasizing sustainability aspects in the logo they will design (e.g. Looka). Many AI generative tools may help crafting catchy digital ads and marketing campaigns that resonate with the right audience (e.g. AdCreative.ai, see Gao et al., 2023). Additionally, AI-driven crowdfunding tools could enhance early stage fundraising efforts by optimizing campaigns (e.g. The Crowdfunding Formula). Then, if one wants to create an 'app' to better engage customers, AI-powered software coding assistants (e.g. GitHub Copilot) could certainly help entrepreneurs with limited coding skills. Additionally, there are tools that allow 'app' creation using simple language prompts, producing code that (probably, now) may not be highly sophisticated but is still considered acceptable for many use cases (and it is worth emphasizing once again that it is likely to improve over time). Examples in this respect include Adalo, a start-up created in 2018 with the claim 'Build Your Own No Code App', and Replit, (co-)founded in 2016 by Amjad Masad, a former Facebook employee. In a lengthy interview, Masad envisioned a future of citizen-developers, where everyone, without any programming skills, could 'develop' (My First Million, 2024). More generally, the field of 'low code, no-code and AI coding' is rapidly expanding, offering multiple solutions (Carroll et al., 2024) – an emerging trend sometimes informally referred to as 'vibe coding', which emphasizes intuitive, prompt-based, or natural language interfaces for software development.

Finally (though only for the sake of brevity), the prospective entrepreneur could rely on specific platforms that manage the pricing strategies for 'our' shoe, such as Pricen, a Finnish multinational company, whose stated mission is to 'help retail and e-commerce businesses maximize profitability with an AI-driven pricing solution that intelligently optimizes product prices at every stage, from launch to markdown'. Or, they can use Tableau Pulse by Salesforce to gain key insights on key aspects of the nascent 'SoleScribe Snack Shoe' venture, from supply chain optimization to sales performance analysis. More broadly, one could use Shopify Magic – Shopify's integrated suite of AI tools – to generate product descriptions, draft email content, and personalize the store layout, directly supporting the creation and promotion of the online business.

To conclude, the scope and intensity of the support potentially offered by 'AI entrepreneurial buddies' are therefore broad and substantial, respectively. This is especially true considering the fact, already emphasized, that many of these applications either are or have counterparts (or will themselves develop counterparts) based on agentic AI, making it possible to automate such processes (although this does not necessarily have to represent the optimal choice from the start-up's perspective).

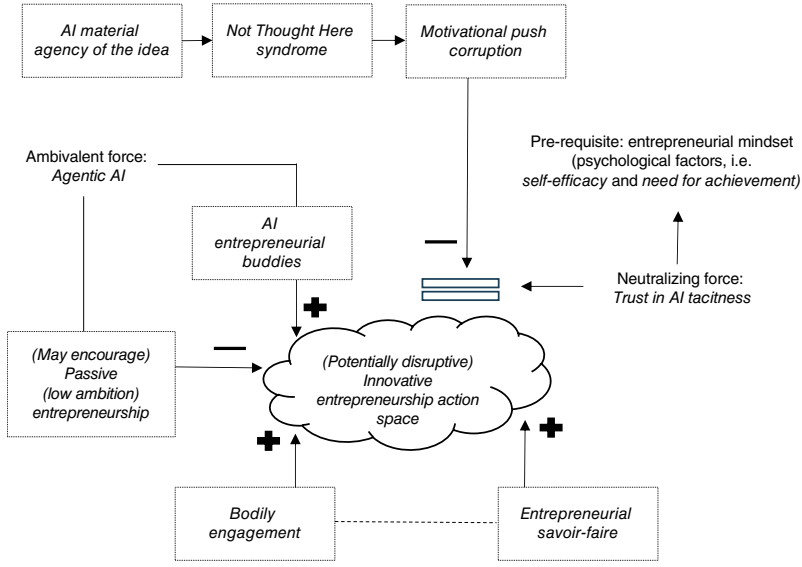


Figure 4.3 The mind and the body of the AI entrepreneur: forces shaping the action space

teammates' (e.g. Dell'Acqua et al., 2025). We prefer to refer to them more simply as AI assistants, or rather, 'AI entrepreneurial buddies'.

In Figure 4.3, we summarize all the themes explored so far, spanning both the increasingly material nature of entrepreneurial ideas and the psychological and physical dimensions of the entrepreneurial act – dimensions that can either narrow or broaden the scope of entrepreneurial action (potentially of a disruptive kind). Thus, within this framework, agentic AI can play an ambivalent and, in some respects, conflicting role. On the one hand, it may enable forms of entrepreneurship that are modest in scope and scale, potentially displacing more ambitious (and, at least in principle, more disruptive) entrepreneurial undertakings. On the other hand, as we have just seen, it can also serve as a powerful enabler in the implementation of any innovative entrepreneurial idea, including the most promising ones (besides the 'SoleScribe Snack Shoe').

Among the various aspects explored in Chapter 5, one concerns how the presence of 'AI entrepreneurial buddies' may influence the relevance of the three 'capital' pillars in innovative entrepreneurship, potentially reshaping their relative importance and altering the ways in which founders deploy them throughout the entrepreneurial journey.

5 *Artificial Intelligence and Success in Innovative Entrepreneurship*

Prologue

We recalled in the Introduction how capitalism 'works well' as an economic system insofar as 'whoever wins does so because of a combination of good skill, hard work, and good luck' (Bowles, 2013, p. 24). Competition, even if not entirely on equal footing, should take place on a playing field that is not too dissimilar among players, ensuring that everyone has a 'fair shot at entry' (Khan, 2017, p. 741), and where those who succeed do so with some degree of merit, according to the three factors above. The weight of each ingredient can, of course, shift over time and across different historical periods and geographical contexts, and as we have already emphasized, the specific institutions in place shape how talent and effort are directed towards productive entrepreneurship (Baumol, 1990). However, a strong presence of the first two factors (and perhaps especially the second) is essential to make the system acceptable to individuals – particularly to those who, even in relative terms, 'lose' – and to ensure social mobility from an SM I perspective (Schumpeter, 1911/1934).[1] After all, even Adam Smith, in *The Wealth of Nations* (1776), identified among the key merits of the market the opportunity for each individual to discover their own talent and then cultivate it to perfection, for their own benefit but also for that of society as a whole. In fact, in the second chapter of his masterpiece, he wrote:

> And thus the certainty of being able to exchange all that surplus part of the produce of his own labour, which is over and above his own consumption, for such parts of the produce of other men's labour as he may have occasion for, encourages every man to apply himself to a particular occupation, and to cultivate and bring to perfection whatever talent or genius he may possess for that particular species of business. (Smith, 1776, p. 23)

[1] But it is important to note here that the role of the third factor, namely luck, remains important and well-documented across various domains in modern advanced economies (see, for example, Andreani et al., 2025).

Do the capitalistic pathways to success remain relevant today, in the AI era?

Having explored the relationships between AI, entrepreneurial idea generation, and implementation in the previous pages, we have now set the stage for analyzing whether, and to what extent, AI may reshape the fundamental drivers of entrepreneurial outcomes. Indeed, we have already outlined the (loose) boundaries of the most promising space for innovative entrepreneurship, and we have analyzed the forces that shape it. This is a space in which the playing field should probably be less uneven between new and established operators. Now let's analyze what the main factors for entrepreneurial success in this space are. In particular, we turn our attention to the three key determinants of (successful) innovative entrepreneurship identified in Chapter 1: the human and social capital of the founders and their access to the capital market, particularly (but not exclusively) venture capital.

Does AI alter the relevance of these three types of 'capital'?

Founders' Human Capital and AI: A Renewed Perspective on the 'Competence-Based View'

In Chapter 1, we established the relationship between founders' human capital and the performance of innovative start-ups as a well-documented phenomenon – what one could label as a stylized fact. Entrepreneurial ventures are typically (though not always, of course) initiated by individuals who possess domain-specific knowledge (i.e. those who 'know'), and the successful development of an innovative start-up commonly requires some form of *knowledge* and expertise that is often deeply embedded in the industry of the nascent firm. This expertise is often acquired through prior experience within existing enterprises, which serve as incubators for entrepreneurial capabilities, as envisaged (among other approaches) by the KSTE.

In a 2005 article, we specifically discussed the competence-based view of the (innovative) start-up to elucidate the close relationship between founders' human capital and firm performance (Colombo and Grilli, 2005). Is the relationship between founders' human capital and start-up performance still strong with the advent of AI?

As a first preliminary answer, for the reasons detailed in the previous chapters, it should be clear that we are rather sceptical on the notion that one could leap directly from an 'AI-generated idea to

wealth' without requiring any competencies whatsoever. To put it simply, despite any foreseeable advancement in passive entrepreneurship potentially driven by agentic AI, the idea that a start-up can generate huge profits without any forms of 'skills' – as some may claim merely to increase the appeal of social media content – is, in our view, bound to remain an illusion in many contexts.[2] However, asserting this does not necessarily imply that AI will not alter the type of knowledge required in the entrepreneurial process. In other terms, 'good (entrepreneurial) skills' do evolve over time. Indeed, the 'knowledge base' on which an innovative start-up relies is, in our view, inevitably subject to change in many contexts. Thus, once again, we align with Rodrik's vision (footnote no. 14, Chapter 1) and respond to the above question with a 'yes and no', meaning 'it depends'. It depends on what exactly is meant by competence, on the type of entrepreneurship under consideration, and ultimately on the specific (formal and informal) institutions in place.

But let's proceed with order. Our focus here in this chapter will be on competences and on the extent to which AI technology, in and of itself, may drive change. More institutionally grounded considerations will be addressed at a later stage (Chapter 7). And certainly, our objective in this respect is not to make bold or definitive predictions; rather, it is to offer some reflections that may, at times, hint at possible future developments – albeit in a more nuanced and speculative way.

AI, Uncertainty, and Entrepreneurship as Social Elevator

As we know, 'while history may not have a tendency to repeat itself, it often rhymes'. This quote, commonly attributed to Mark Twain, aptly captures the fact that one of the most lively debates surrounding AI concerns its impact on the labour market – a pattern that also characterized previous waves of technological change, such as the early stages of the information and communication technology (ICT) revolution (e.g. Fairbairn, 1982). Similarly, the current debate centres on whether workers for the effects of AI will face upskilling, reskilling, or deskilling, as well as on which job tasks are most affected and which

[2] Even when adopting the extreme perspective of the 'passive entrepreneur', it is widely accepted that they must still inevitably rely on some form of human competence to sustain performance (Houston, 2025).

competencies are most in demand. This discussion is heavily fuelled by both academic research and what is commonly referred to as grey literature. We have already mentioned in the prologue of Chapter 4 some examples of the former (e.g. Felten et al., 2021), but major global organizations are also actively engaged in examining how the labour market will evolve in both the near and more distant future. Consider, for instance, the Organisation for Economic Co-operation and Development (OECD), with its periodic *Skills Outlook*, or the World Economic Forum (WEF), with its recurring *Future of Jobs Report*. Although these analyses exhibit a certain degree of heterogeneity – as one might expect – it is far from our intention to suggest that they are not highly relevant. On the contrary, they can still provide valuable guidance for directing human capital investments more efficiently, particularly within the education system. Nevertheless, in our entrepreneurial context, these otherwise noteworthy endeavours appear somewhat less essential. Innovative entrepreneurship – and, by extension, capitalism itself (or at least its 'good' form, as described by Baumol et al., 2007) – thrives on unpredictability. Or, to put it more precisely in the terms of our analysis, it requires a substantial degree of the Knightian (ontological) uncertainty we previously discussed. Programming (and thus predictability) is, in fact, more akin to socialism than to capitalism (Schumpeter, 1942).

In this light, two aspects deserve particular emphasis, as they often go unnoticed in the current debate on AI and the labour market.

In any capitalist society, it is firms – including start-ups – that demand labour and, in doing so, shape the types of jobs available in the market. Accordingly, the kinds of (new) jobs that will exist in the future will also depend to some extent on the new forms of entrepreneurship that will emerge because of the AI. Setting aside institutional considerations (including the power to face AI giants, see Chapter 7), from a purely technological point of view, given the inherent uncertainty surrounding entrepreneurial success – even in the age of AI – it is virtually impossible to predict which firms will thrive. As a result, we can at best formulate only a vague idea of the precise competencies that will be required in the future, and even well-informed opinions on the matter may differ substantially from one another.

In other words, to some extent, it is considerably easier to discuss the kinds of tasks and occupations being affected today, since they are embedded in existing organizations – firms whose presence is certain.

By contrast, reflecting on the jobs of tomorrow proves far more elusive (see Horowitz, 2025 for the same view): at least some of the firms that will create and host these jobs do not yet exist, and our capacity to foresee what they will look like is limited. In this sense, the nature of future entrepreneurship will inevitably shape the future of work – yet it remains remarkably difficult to anticipate how. Simple examples from the recent history of ICT illustrate this challenge. Could anyone, at the dawn of the 2000s and the rise of the Internet, have anticipated the emergence of the role of a social media manager for an online influencer? Or that of a performance coach in the e-sports industry – a professional dedicated to enhancing the mental resilience and focus of competitive e-gamers? And that both professional figures would be in such high demand and require such ad hoc competences as to lead to the creation of dedicated university degree programmes (for further details, see Sagnard, 2024; Woolcock, 2024)? I would argue not, and these are merely two examples among the many we could cite.[3] There is little reason to believe that the challenges of foresight are any less pressing in the context of AI.

But, and this brings us to our second point, it is precisely this ontological uncertainty regarding the skills required in the brave new AI world that may allow entrepreneurship to also fulfil its essential role as a mechanism of social mobility – an aspect particularly dear to Schumpeter (1911/1934) and highlighted in our Chapter 1. If there were no uncertainty about the competencies required in the future, those enjoying privileged socio-economic positions today could easily perpetuate their advantages for themselves or their children by clearly knowing how to guide them or what specific knowledge they should pursue.[4] Therefore, while the effort to predict the skills needed for

[3] We deliberately select here two examples that, at the very least, appear among the most debatable in terms of their potential contribution to collective social welfare – among many others, far more substantial from this standpoint, that could equally be mentioned and were similarly difficult to predict in the 1990s (for instance, telemedicine care coordinator). We shall return, in the concluding chapters of the book, to the broader issue of the 'social welfare significance' of work (and of innovative entrepreneurship), as well as to the non-negligible role that public policy may play in averting the risks of a postmodern drift in advanced societies brought about by AI – including, crucially, through the educational system itself.

[4] The issue is even more pressing today, particularly in the context of current forms of Western capitalism, where an increasingly large share of wealth is inherited (The Economist, 2025a), thereby facilitating the transmission of intergenerational advantages.

tomorrow is clearly justified in terms of preparing the workforce of the future, the concern that such prediction is not entirely possible is surely mitigated by the fact that this very inability feeds the capitalist process with its lifeblood. The greater the uncertainty, the greater the likelihood that innovative entrepreneurship can act as a powerful social elevator, levelling the playing field and ultimately making capitalism more acceptable in the long run. This has occurred with every major technological wave in the past.[5] Considering only its technological characteristics and setting aside other possible institutional factors, there is no reason to sustain that the same cannot also hold true for the AI revolution.

(Specific vs. Generic) Human Capital and the New Dynamic AI-Enabled Orchestration Capability

Besides reaffirming the centrality of competence, some authors (Taddy, 2019; Chalmers et al., 2021) go well beyond this, emphasizing once again the necessity of some *specific* component of knowledge on the part of the successful innovative entrepreneur in the age of AI (emphasis mine):

> That is, for business AI to succeed we need to combine the GPML [i.e. general purpose machine learning, author's note] and Big Data with *people who know the rules of the 'game' in their business domain*. (Taddy, 2019, p. 85)

This is, of course, more than plausible, particularly in highly complex business domains such as knowledge-intensive sectors (i.e. high-tech and science-based start-ups).[6] Even within the AI-based sector,

[5] Starting from the First Industrial Revolution, Josiah Wedgwood, a renowned ceramicist often considered among its most prominent figures, came from circumstances described as characterized by 'moderate means' (Koehn, 1997, p. 21). Similar examples can also be found in the Second and Third Industrial Revolutions, such as Thomas Edison, who rose from humble beginnings to become a central figure in the era of electrification, or Steve Jobs, who – as is well-known – did not emerge from particularly privileged origins either.

[6] In December 2024, Aidan Toner-Rodgers, a Ph.D. student in Economics at MIT, released a working paper titled *Artificial Intelligence, Scientific Discovery, and Product Innovation* on several prominent scientific platforms. The study claimed that in R&D-intensive environments – such as materials science – AI tools significantly enhanced product innovation when used by highly skilled researchers, enabling them to generate more novel compounds, file more

in a constantly evolving landscape where some skills become routine, while others, rarer, are magnified in a veritable war for talent (as we shall discuss in Chapter 6), the *specific* component of founders' human capital seems to remain critical to the most significant entrepreneurial phenomena. While some recent studies – such as those by Brynjolfsson et al. (2025) on workers' productivity in customer services and Otis et al. (2024) on (not necessarily innovative) entrepreneurs in Kenya – seem to suggest that the returns from generative AI can vary widely across skill levels, and do not always favour the most competent individuals (e.g. Brynjolfsson et al., 2025; see also the survey evidence produced by Li, 2026), this latter result appears less true in science-based, highly innovation-intensive contexts.[7] Our analysis of the 'Deep Learning' segment in Box 5.1 is in line with this, and the argument extends naturally to other highly knowledge-intensive contexts. Adhering to this view, especially in scientific and technical fields, AI will act not as a substitute but rather as a *competence amplifier*, enhancing specific human expertise. Admittedly, the relationship between individual human capital, AI, and performance is certainly complex and multifaceted, varying across different contexts (see Li 2026), and a comprehensive review of

patents, and develop more radical inventions. The paper garnered substantial attention and was widely endorsed by leading members of the scientific community (e.g. Lahart, 2024). Its results were even anticipated in the news by Nature (see Castelvecchi, 2024). However, it was later revealed that the study may have been based on fabricated data, and MIT requested that the research be fully retracted (Lahart, 2025). What remains noteworthy, however, is that at least part of the paper's (albeit short-lived) scientific success can be attributed to the fact that it appeared to empirically support a hypothesis that many in the academic community already found intuitively plausible, that is, the complementary role of AI and domain-specific human expertise in knowledge intensive contexts. In other words, the study did not seek to persuade the readers of a counterintuitive claim, but rather to provide empirical validation for a widely held belief for which solid evidence had been (and, to some extent, still is) lacking.

[7] Also, the studies on AI and creativity we reviewed in Chapter 3 point to a possible uplift in average creative potential and, in some cases, to indications that differences across individuals might narrow (although this remains an open question). Needless to say, we maintain this point; but our view here is that when it comes to the (successful) implementation of an entrepreneurial idea – especially in high-tech and knowledge-intensive sectors – this is still likely to depend, at least to some extent, on the founders' domain-specific human capital. This reading is largely in line with the above-cited Otis et al. (2024) and more directly supported by the scholarly references that follow in the main text.

Box 5.1 Deep learning and founders' human capital

As of January 2025, we identified all VC-backed start-ups founded globally since January 2010 that were listed in the PitchBook database at that time and categorized as operating in the 'Deep Learning' sector. Along with Crunchbase (another database we will draw on later), PitchBook represents an important source of information containing 'detailed information on private firms and funding deals with professional investors and is gaining traction in studies on entrepreneurial finance and innovation for growth-oriented firms', and it has already been used in the AI domain for this very purpose (Chattopadhyay et al., 2025, p. 1032). In total, as to aforementioned criteria of selection, PitchBook reports 4,609 deals (funding rounds) across 1,476 companies (i.e. start-ups), involving 5,903 investors. Using the information on the CEO provided by PitchBook itself, taken as a practical proxy of the *lead founder* given the nature of the firms considered, and complementing or filling in missing data through extensive manual research conducted via LinkedIn, we reconstructed the educational and professional background of the identified individuals. Among other aspects, we classified their prior work experience as either *generic* (i.e. not directly related to the sector in which their start-up operates) or *specific* (i.e. closely aligned with that sector). Out of the 1,162 start-ups for which we were able to collect data on this variable, as many as 716 (61.6 per cent) have a CEO with work experience easily classifiable as *specific*. Firms whose CEO possesses this component of specific work experience are also those that, on average, raise a larger amount of investment over their lifetime. Everything else being equal – in particular, controlling for firm age – a start-up whose CEO has specific work experience raises, on average, about slightly less than 30 per cent more investment (in terms of cumulative deal size) than one whose CEO lacks it. Similarly, firms whose CEOs have an educational background consistent with AI or machine learning tend to attract higher levels of funding. Controlling for the different timing of deals, a start-up whose CEO shows a strong match between his or her graduate studies and the field of 'Deep Learning' sees its funding rounds averaging about €15.9 million, compared to €9.7 million for those without such alignment.

Overall, the picture that emerges is certainly one in which *specific* human expertise in a highly knowledge-intensive field such as 'Deep Learning' continues to be extremely relevant.

this nascent literature is beyond the scope of our analysis.[8] However, by narrowing the focus to domains more closely related to the one examined here – that of innovative entrepreneurship – available (albeit still preliminary) evidence does suggest the existence of this amplifying effect. For example, among others, Conti and Messinese (2025) show that even predictive AI tools such as Google Analytics 4 can disproportionately benefit highly skilled founders of innovative software start-ups, who use anomaly detection to refine strategic reasoning and generate novel ideas. Similarly, Asam and Heller (2025) exploit the release of GitHub Copilot as a quasi-natural experiment to show that generative AI tools can significantly accelerate the early development of software start-ups, reducing their time-to-funding by around 20 per cent, with this effect particularly pronounced when founders have prior technological or managerial experience. In fact, *de minimis*, AI can free up time – time that can be reinvested in more value-added activities. But beyond this, AI can also enrich those very activities, especially through its generative capabilities. Both effects, we believe, are likely to be more impactful for entrepreneurs in high-tech and knowledge-intensive sectors possessing strong specific human capital relevant to the task at hand.

In the search for this *competence amplification* effect of AI, as defined earlier, we also argue that start-ups may easily enjoy a structural advantage over established incumbents. In early stage ventures, typical AI capabilities – such as automation of complex tasks, augmented decision-making, and AI-human creative collaboration – are not yet locked into rigid procedures, as is often the case in large corporations (see Hoque et al., 2025). This flexibility allows greater space for valorizing individual competencies, where the optimal mix of AI and human input cannot be defined a priori or standardized across

[8] One of the recent studies that, in our view, best highlights this complexity is that of Agarwal et al. (2023), in the field of radiology. Their experimental evidence shows that while AI predictions are, on average, more accurate than those of human radiologists, human-AI collaboration does not systematically enhance diagnostic performance. This is largely due to systematic deviations from Bayesian belief updating: radiologists tend to underweight AI input relative to their own information and often neglect the correlation between the two sources. The average treatment effect of AI assistance, however, masks important heterogeneity: AI input improves human decisions when the algorithm is confident, but may actually degrade performance when its predictions are uncertain.

employees (e.g. as required in 'humans-in-the-loop' systems within large firms). By contrast, the 'AI-empowered entrepreneur' is not constrained by predefined workflows, and can instead shape AI tools to fit their own knowledge base. It follows that co-intelligence, as well as hybrid or symbiotic forms of intelligence (Mollick, 2024; Jarrahi et al., 2022; Zhou et al., 2021), discussed in Chapter 2, are more likely to emerge and take root through entrepreneurial initiatives than through top-down policies within incumbent firms.[9]

However, while AI can operate as a complement that amplifies entrepreneurial competence (as illustrated earlier), we also claim that it may not only function in this role: in many circumstances it can also act as a substitute for competence, and this may easily depend on the business domain considered, using the aforementioned excerpt from Chalmers et al. (2021). To put it simply, particularly in less tech-intensive sectors, the amplification effect described earlier – stemming from the interaction between specific human capital and the ability to integrate it with AI – could not necessarily require deep domain-specific knowledge. In other words, in many industries (probably unlike deep learning), AI-based entrepreneurial tools make it feasible for new founders to compensate for a wide range of knowledge gaps, convincing prospective entrepreneurs to found the venture, with results not necessarily unsatisfactory.

This is evident in the analysis by Dushnitsky and Stroube (2021) on Shopify and its toolkit of low-code solutions (e.g. Shopify Pay, Hosted, Discounts, Gift Cards). Shopify enables users to launch and manage online stores with minimal technical expertise (see our Box 4.1). Examining a large sample of platform-based and non-platform-based e-commerce start-ups, the authors show that low-code entrepreneurship often follows leaner growth paths but performs comparably to start-ups led by more technically proficient teams – at least according to certain metrics like investor cash-on-cash returns. In their concluding remarks, they assert:

[9] For example, Raisch and Fomina (2025) identify three different ways in which organizations can combine human and artificial intelligence in the search for solutions to business problems (autonomous, sequential, interactive search), each with its own protocols and outcome characteristics. Clearly, no approach is inherently superior to the others; rather, the effectiveness depends on important contingency factors at play. Our point here is precisely that a start-up, by definition, is less constrained by corporate rules and protocols, and is therefore better positioned to find the right fit.

> The impact of a low-code tool such as Shopify is that it alters the resource environment so that the hypothetical entrepreneur described above [i.e. a traditional MBA student lacking technical skills, author's note] now needs to mobilize fewer resources to pursue an opportunity. Put differently, low-code tools can empower non-technical individuals to 'make' a viable attempt at refining and scaling an entrepreneurial opportunity. The cadre of coders is no longer needed (or is needed to a lesser extent) and thus neither is the need to raise substantial amounts of capital to fund them. It follows that low-code tools reduce the need for technical skills in the first place […]. (Dushnitsky and Stroube, 2021, p. 7)

While their conclusion may not apply across all sectors – or outside the digital domain – it certainly makes a valid point, enriching a broader picture in which the technical content of the product likely plays a differentiating role. From a competence-based perspective, it is difficult to conceive of a process or resource as central to a start-up's success if the team does not, in some meaningful way, possess, control, or understand it (also considering the importance of the psychological factors already discussed, such as *self-efficacy*). Yet, what we aim to stress here is that AI may allow a viable entrepreneurial path even with limited prior knowledge, especially in less knowledge-intensive sectors. For example, while starting a biotech company without deep expertise in areas such as molecular biology, biochemistry, or genetics remains implausible – much like, for many observers (though we are fully aware that opinions differ), founding a software company without any knowledge of software architecture[10] – starting

[10] We previously mentioned the inherent difficulty of predicting which specific competencies will be most needed in the future and which will not – and how even well-informed opinions can differ significantly in identifying them. While the evidence reviewed in this chapter (e.g. Asam and Heller, 2025; Conti and Messinese, 2025) shows how AI can amplify the technical capabilities of software-house founders, a parallel public debate has emerged on a different front: whether younger generations should still learn to code in an era increasingly shaped by AI-assisted development. Opinions on this matter vary widely, especially given the growing possibility of using natural language to generate code. Our view is that coding is a language, and as such, it can be translated – with varying degrees of effectiveness depending on the context and the purpose (but destined for continual improvement). However, like any language, it evolves over time, and it is moreover, only one aspect of software development. Most importantly, for the purposes of this discussion, we find it implausible that someone could feel genuinely confident (this is where *self-efficacy* comes into play) in building and selling a core product – one that has to be maintained and renewed over time – without having at least a basic

a fashion business (as in our example of footwear) without in-depth sector knowledge is becoming increasingly feasible with AI support. The story of Gymshark, a Shopify-based unicorn in the sportswear space (Tognini, 2023), may well illustrate this. Therefore, we argue that AI – alongside its 'entrepreneurial buddies' – can lower the entry threshold in the implementation phase, particularly in less technologically demanding sectors. In such contexts, ventures can be successfully launched from diverse knowledge bases, whether technical or commercial, without needing to blend them all within a single founding team.

Thus, in knowledge-intensive businesses, AI may indeed serve as a powerful *enhancer of domain-specific human capital*, but we believe that, in others, less technical, AI technologies can more easily become a substitute for competences.[11] This scenario may also signal a resurgence in the importance of *generic human capital*, at least within sectors characterized by lower knowledge or technological intensity. In fact, at that point, the decisive factor becomes not domain-specific competence itself, but rather the ability of leveraging one's idea and knowledge in synergy with AI systems, and in that respect generic human capital may be helpful as its specific counterpart.

However, both scenarios point to the importance of a new kind of meta-competence, which we call *dynamic AI-enabled orchestration capability* – a perspective that reframes the competence-based view of the innovative start-up in the AI era. This brand new capability

understanding of how it works from a technical point of view. And even if they did, they would probably still be at some (even psychological or merely perceived) disadvantage compared to a founding team with at least some level of knowledge in the field. That said, it is equally evident that in many digital domains – and this is already visible in mobile apps and videogames, where no-code and AI-assisted tools are rapidly lowering the barrier to entry – start-ups can emerge from individuals who possess no technical knowledge in the traditional sense. Whether markets will ultimately be able to absorb this growing number of entrepreneurial attempts enabled by 'creating without knowing' remains an open question, and one that connects to the postmodern forms of entrepreneurship we will discuss in Chapter 7.

[11] More precisely, we could argue that in certain domains, rather than saying that technologies (in this case, AI) substitute for specific human skills, it would be more accurate to speak of a 'transfer of competences' from the human to the machine. A transfer in which the machine may clearly hold an unparalleled advantage: if, for example, one thinks of LLMs, they can draw on billions of years of cumulative experience generated by an innumerable number of human beings (and this wealth of knowledge is, moreover, always available at a relatively affordable cost).

captures the founder's ability to configure, adapt, and extend resources of the innovative start-up in collaboration with intelligent systems. Specifically, this *dynamic AI-enabled orchestration* involves the entrepreneur's ability to select and integrate 'AI entrepreneurial buddies' (AI-to-AI) and effectively coordinate interactions between these tools and human actors, including themselves (AI-to-human and human-to-AI), dynamically reconfiguring workflows and decision-making processes in response to evolving circumstances.[12] In Box 5.2, we

Box 5.2 From tools to orchestration: making the 'SoleScribe Snack Shoe' stack work

To move from an AI-supported ideation process to a fully operational prototype – without requiring significant upfront capital or a large team – the aspiring founders of the 'SoleScribe Snack Shoe', must go beyond simply using AI tools and instead orchestrate them intelligently and flexibly. In other words, what matters is not just the selection of the right 'AI entrepreneurial buddies' for each task, but the way they are dynamically coordinated, recombined, and reconfigured around the venture's evolving needs, always considering the human component. This is the essence of what we intend as *dynamic AI-enabled orchestration capability*.

Just consider the design phase. Using MidJourney, the founder can prompt the generation of dozens of visually striking hybrid concepts – part shoe, part pencil, part 'bread'. But the power of the tool lies not in its output alone: it lies in the human decisions about how to prompt, what to keep, and when to reject a seemingly 'creative' idea that, in practice, would violate basic requirements of comfort or durability. Once an initial silhouette emerges, Vizcom can take over to convert the sketch into a print-ready 3D model. But even this step requires orchestration. Does the render hold up under different lighting conditions? Are the sole dimensions consistent with average foot width? The entrepreneur may choose to bring in a freelance ergonomics expert, or instead re-prompt Vizcom with structural adjustments. The point is: while AI manages the complexity of translating a concept into a model, it is the founder who decides where human intuition or expertise – whether their own or that of an employee – must intervene,

[12] It is clear that, from our perspective, the importance of this meta-competence can only be amplified by the development of agentic AI (for a similar perspective, see Purdy, 2024).

and when speed and fidelity matter more than photorealism. Syntilay's founders report that 70 per cent of their custom-fit shoe's design is generated through AI (Kato, 2025), which means the remaining 30 per cent is shaped by human hands and judgment. But what is the 'right' percentage? And even if one could define it, would it remain valid across different contexts – or over time? Whatever the proportion, the deeper question is: where should the machine take the lead, and where must human insight remain non-substitutable?

We believe that the entrepreneur who excels at answering this type of questions, that is, who excels in the *dynamic AI-enabled orchestration capability*, will hold a competitive advantage.

But this applies only to the initial stages; the importance of orchestration also comes into play once we reach the threshold of physical prototyping. Rather than relying on a traditional supply chain, the entrepreneur can coordinate with a company offering fully 3D-printed footwear manufacturing, for example, Zellerfeld. Each SoleScribe shoe can be printed on demand, without tooling, molds, or batch constraints. The orchestration here means configuring a system in which customer orders, once received, flow directly from design files into Zellerfeld's production queue, with no inventory or warehousing involved. Yet again, human oversight is key: someone must test early batches, provide feedback on fit and print fidelity, and possibly redesign the mesh for improved food-grade integration.

Then, as entrepreneurs, we could be convinced that the success of our offer lies in the use of dynamic sizing and fit logic. Thus, borrowing from Syntilay's approach, the entrepreneur could integrate a foot-scanning app that uses a smartphone camera to generate precise 3D foot models for each customer. These are then processed via slicer software that adjusts the digital shoe file accordingly – meaning every unit produced is technically bespoke. In fact, what truly distinguishes this approach is the learning *feedback loop* that underpins it. Every customer scan, print parameter, or support ticket is a data point – and the system is configured to absorb these into a continuously evolving design logic.

Again, throughout this production process, the ability to switch between AI and human agency can be important. When a new printing anomaly appears, is it better to let the AI-driven slicer resolve it automatically, or pause the pipeline and consult a materials engineer? If customer feedback points to fit issues among runners with high arches, should the AI model be re-trained, or should the product line be diversified manually? These decisions are not fixed in advance – they are made on the fly, context by context.

That's what makes the orchestration 'dynamic'.

continue to draw on our example from the footwear industry to delve a bit deeper into how these orchestration and architectural activities might take shape (inspired by a real case: the start-up Syntilay). This example allows us to highlight how this new capability is highly idiosyncratic: there is no fixed recipe for the human-machine relationship that the AI entrepreneur can follow which would be valid across all contexts – industrial, geographical, or otherwise – and over time. Rather, it is a firm-specific recipe in constant and continuous evolution.

Obviously, it is well-understood that the integration of AI and human capabilities is likely to affect all domains of work. Accordingly, AI literacy is increasingly conceived as a generally essential competence. Among the many signals of this global trend, one notable example is China's recent introduction of mandatory AI education starting from primary school (Reuters, 2025), a clear indication of how seriously some countries are beginning to treat this matter. What should be emphasized here, however, is that innovative entrepreneurship not only does not represent an exception but, on the contrary, due to its already noted complexity, the orchestration activity in this context may prove even more challenging than in more clearly defined and bounded occupational domains. In this respect, developing *dynamic AI-enabled orchestration capabilities* in entrepreneurship is likely to depend on both hands-on experience and dedicated educational programmes that combine technical AI skills with entrepreneurial competencies. Many such programmes are now emerging in leading universities worldwide (at the Master level, for instance, see those offered by Carnegie Mellon University or the Hong Kong University of Science and Technology, among others), reflecting a shared recognition of the hybrid expertise required to navigate AI-powered venture creation. We will return to this in Chapter 8.

On the Size of the Entrepreneurial Team

The above discussion leads us directly to the consideration that, if we focus solely on competences – even without taking into account other factors, such as the social capital that their development may generate (for instance, through education or work experience), or

the financial capital they may be associated with (what Colombo and Grilli, 2005 define as 'the wealth effect' of founders' human capital) – we should expect innovative start-ups founded in the age of AI to be launched by smaller (human) teams, potentially also involving fewer individuals as employees.[13] As mentioned earlier, on the one hand, in more knowledge-intensive sectors, AI technologies – at the very least – may enable entrepreneurs with high levels of specific human capital to devote more time to value-adding activities by freeing up their own time and reducing the need for others to carry out initial entrepreneurial tasks. On the other hand, particularly in less technologically intensive sectors, the diffusion of AI tools may lead to a downsizing of entrepreneurial teams, as 'AI entrepreneurial buddies' increasingly substitute – at least in part – for the human-embodied knowledge and skills that once required a broader founding group. In other words, if the team formation process – frequently observed and previously discussed – involves a lead founder bringing in co-founders based, among other factors, on the need to secure complementary competences to launch the venture, then the ability to rely on AI technologies clearly reduces, in and of itself, the extent of that need. This intrinsic 'substitution' comes with all its inherent pros and cons. Clearly, one loses emotional support, shared commitment, and mutual accountability (see again the final part of the excerpt from Cooper and Bruno, 1977, reported in Chapter 1). However, one certainly gains from lower coordination costs, reduced risk of misalignment in strategic vision, and the avoidance of disagreements and conflicts. The point is made especially clear by serial entrepreneur Zoë Barry in her

[13] Many observers – as well as leading figures within the AI ecosystem – predict that we will soon witness the rise of AI-driven unicorns founded and operated by a single individual, with no employees (Sam Altman, CEO of OpenAI, made such a prediction in an interview with Alexis Ohanian, co-founder of Reddit; see Confino, 2024; see also Sequoia Capital, 2025). This trend, however, is not entirely novel nor exclusively attributable to AI, as it reflects broader dynamics that have long been at play in the digital services sector. A notable example is WhatsApp, which was acquired by Facebook (now Meta) in February 2014 for approximately $19 billion. At the time, the platform had around 450 million monthly active users – a figure that has since grown substantially – while employing only about 50 engineers (Metz, 2015).

insightful article, tellingly titled *In the Age of AI, Do You Really Need a Cofounder?*:

> Really think about how cofounder dynamics have played out historically. Can you name more than a handful of successful companies where all co-founders wound up happy and successful? The vast majority of the time, a company *will* fail because of co-founders disagreements. And if the company does succeed, it usually happens in spite of the disagreements, board struggles and litigation. Cofounder disputes have stunted, or even destroyed, many great companies. Snapchat was embroiled in a messy legal dispute over equity with ousted cofounder Reggie Brown. Zynga struggled with ongoing disagreements over strategy, leading to the departure of several early executives. Uber's Travis Kalanick was forced to step down after severe disagreements with the board. In all these cases, the cofounder conflicts led to financial strain, reputational damage, and significant business disruption. These are the companies whose names we know, that made it through in spite of cofounder dynamics. But, there's a graveyard near Silicon Valley littered with *thousands* of other companies that *might* have made it had cofounder disagreements not led to a failure to be decisive, or a failure to pivot, or a failure to hire great talent due to one founder's ego. (Barry, 2023)

Some social scientists in the field of organizational studies go even further, speaking of the possibility of 'zero-human organizations' (Puranam, 2024, p. 39). We do not consider this a feasible scenario – at least not within our domain of investigation and the foreseeable time horizon we have adopted – if only due to the legal implications such a model would entail. However, without going to such extremes, it is clear that the increasing emphasis in journalistic discourse and narratives (e.g. Ashley, 2025; Strozier, 2025) on AI-augmented 'solopreneurs' – not necessarily in search of passive income, but quite the opposite – may further exacerbate the trend towards downsizing entrepreneurial teams. Emerging evidence on the subject does corroborate the idea that this is already happening, with AI being one of the main drivers behind the phenomenon (Walker and Dowd, 2025).[14]

[14] In its *Founder Ownership Report 2025*, Carta, a US-based fintech company, using data from more than 45,000 start-ups incorporated from 2015 through 2024, documents how the percentage of ventures 'led by a solo founder has more than doubled over the past decade, reaching 35 per cent in 2024' (Walker and Dowd, 2025, p. 1).

Soft vs. Hard Skills

While we may observe a diminishing relevance for entrepreneurship of certain *hard* competences in specific domains, we do believe, for the reasons already exposed, that the importance of *soft* skills – especially in the theatrical elements that entrepreneurship entails – remains intact. These include the performative aspects of pitching, storytelling, negotiating, inspiring teams, and projecting confidence to investors and partners. In Chapter 4, we introduced the concept of *entrepreneurial savoir-faire* to identify the common root in terms of a shared competence across all these activities, allowing us to encompass them. In our view, the *relative* importance of such soft skills could well be rising as some 'hard' tasks may become increasingly mediated or replaced by AI. Thus, adhering to this perspective, the *human* ability to communicate a vision and build trust becomes more – not less – central to successful entrepreneurial practice in the age of AI-powered ventures. As a consequence, although it is not entirely appropriate to draw a strict parallel between hard and soft skills and the two types of human capital considered here, it remains true that *entrepreneurial savoir-faire* can clearly be impacted by specific knowledge (e.g., through the channel of *self-efficacy*), but it also evidently contains elements that are more cross-cutting and broadly applicable across sectors. For this reason, *generic human capital* – at least in relative terms compared to *specific human capital* – is certainly not less important than it was in the pre-AI era, and is probably even more so, at least within our domain.

Table 5.1 presents a synthesis of our arguments on founders' human capital in AI-enabled entrepreneurship, advancing a revised competence-based view of the start-up in the AI era.

In doing so, and following the reasoning outlined earlier, we draw a distinction based on the technological depth and knowledge intensity of the product in question. Clearly, phenomena of potentially disruptive innovative entrepreneurship can, in principle, affect not only high-tech sectors but also those that are less knowledge- and technology-intensive. At the same time, certain considerations apply across both domains and are captured in the ad hoc column on cross-sectoral patterns.

Table 5.1 *Role of founders' human capital in the AI-related innovative entrepreneurship space*

Dimension	High-Tech / Knowledge-Intensive Sectors	Low-Tech / Less Knowledge-Intensive Sectors	Common Patterns Across Sectors
Specific Human Capital	Remains critical and enhanced by AI (competence amplifier). Deep domain knowledge remains necessary (e.g. biotech, software architecture, deep learning).	Less relevant and partially substitutable by AI entrepreneurial buddies with much easiness, lowering entry barriers (e.g. Shopify-based e-commerce start-ups).	Increasingly sector-dependent: remains essential in high-tech sectors but is increasingly substitutable in low-tech sectors where AI can perform key functions previously requiring expert knowledge.
Generic Human Capital	Stable or even increasing in relative importance as AI compensates specific competence gaps.	Stable or even increasing in relative importance as AI compensates specific competence gaps.	Gains relative importance across all sectors, particularly through its association with soft skills and adaptability in leveraging AI tools and orchestrating hybrid human-AI workflows.
Soft vs. Hard skills	Soft skills more irreplaceable than (some) hard skills, that is, entrepreneurial savoir-faire	Soft skills more irreplaceable than (some) hard skills, that is, entrepreneurial savoir-faire	Soft skills gain prominence as AI mediates or automates more easily hard tasks; entrepreneurial savoir-faire becomes pivotal for storytelling, negotiation, trust-building.

Role of AI	Mainly a competence amplifier (AI-human synergy), enhances the productivity and strategic reasoning of founders possessing strong domain-specific skills.	Competence substitute: significantly lowers entry thresholds, enables entrepreneurs with limited technical knowledge to launch successful ventures.	Facilitates entrepreneurial tasks by freeing founders' time, reducing human capital requirements, and enriching strategic activities via generative capabilities.
Emergence of Dynamic AI-enabled Orchestration Capability	Critical. Founders must dynamically integrate domain-specific expertise with AI tools to gain competitive advantage.	Important (and not always necessarily less technically demanding): involves managing (potentially simpler) AI entrepreneurial buddies for business model execution.	Central emerging meta-competence: involves configuring AI-human workflows and coordinating actions across intelligent systems and human actors in both strategic and operational tasks.
Impact on Entrepreneurial Team Size	Team size likely smaller: founders can perform more value-added tasks themselves due to AI efficiencies; however, highly specialized knowledge may still require collaboration.	Clearly smaller: AI substitutes for co-founders' competences, enabling 'solopreneurship' models more frequently.	Trend towards leaner founding teams: AI reduces need for complementary human competences, enabling 'solopreneurship' but at the cost of shared accountability and emotional support.
AI, Uncertainty, and the Social Mobility Role of Entrepreneurship	Uncertainty about which technical skills will dominate reinforces the need for experimentation and entrepreneurial risk-taking; high entry costs may still limit access for some.	Lower entry barriers allow broader participation despite unpredictable skill demands; experimentation can drive upward mobility.	Entrepreneurship thrives on Knightian uncertainty: in the absence of clear future skill maps, AI-enabled ventures can function as social elevators by disrupting status quo pathways and enabling meritocratic access to value creation.

AI and Social Capital

We stated in Chapter 1 that, traditionally, a successful innovative start-up is anything but an isolated monad. The academic literature consistently emphasizes that its performance may largely depend on its ability to establish effective social connections, through a combination of weak and strong ties, and by cultivating both bridging and bonding social capital. This entails developing an extensive, yet in some areas dense, *network* of relationships capable of granting the nascent entrepreneurial venture access to the most appropriate contacts, even if they are 'distant', while also providing the reputational capital and *legitimacy* necessary to be perceived as a credible actor by such stakeholders – and, consequently, to effectively attract them (e.g. Rutherford and Buller, 2007; Bertoni et al., 2011). This line of reasoning applies both *internally*, through the formation of the entrepreneurial team and the capacity to recruit initial employees (Rocha and Grilli, 2024), and *externally*, in the start-up's ability to connect, formally or informally, with a variety of key actors that are critical for its development, market entry, and eventual success (Teece, 1986). How artificial intelligence enters and reshapes these relationships is a theme we have touched on only superficially and implicitly in some of the preceding pages. We now turn to explore and systematize the topic in greater depth.

Small 'Human' Networks

AI, particularly in its agentic form and through the increasing future availability of 'entrepreneurial buddies', suggests a gradual reduction in the 'human' component of the firm, both in terms of the size of the founding team and the overall number of employees, and a decreased need to establish (some) external engagements with (human) complementary asset providers. Thus, by looking both within the boundaries of a start-up and beyond them, it is reasonable to expect that the optimal number of social relationships – whether internal or external – required for the success of an innovative start-up is likely to decrease, if such a number can be meaningfully defined at all. This shrinkage has direct consequences.

Let us assume for a moment – mindful, as we know, that social capital is only one pillar of performance, not the sole determinant – that

the success of an innovative start-up depends (in ways that are neither necessarily linear nor uniformly positive) on the *quantity* and *quality* of the interpersonal ties that the lead founder can forge with human stakeholders, from fellow co-founders and employees to the various providers of complementary assets essential to the start-up's success. If n denotes the number of these stakeholders, the optimal number of interpersonal links the lead founder must cultivate will fall somewhere between n (a star network) and $n(n - 1)/2$ (a fully connected network), depending on the network's topology (which, in turn, is shaped by the start-up's internal organizational model, among other factors). Whether the lead founder orchestrates the start-up's social capital as a star-shaped network or as a densely interwoven web that approximates a fully connected graph, this reduction in network size can lower (the often convex) overall *coordination costs*. This, indeed, leads to another important element to be considered. In fact, although not unanimously (see, for example, Colombo et al., 2009), a number of contributions in the literature on innovative entrepreneurship suggest or document an inverted U-shaped relationship between a generic variable that we may refer to, for simplicity, as '(human) partners to manage (and their inherent heterogeneity)' and *start-up performance* (e.g. innovation, growth, etc.). This relationship emerges across multiple domains, ranging from founding team size (e.g. Backes-Gellner et al., 2015) to alliance portfolios (e.g. Deeds and Hill, 1996; Rothaermel and Deeds, 2006). Thus, combining all arguments, looking at Figure 5.1, AI can have two distinct effects. First, for the reasons already exposed, it may easily bring the peak of the performance curve closer to the origin – that is, it reduces the optimal number of partners associated with maximum performance. Second, as emerging literature in the field of supply chain management suggests (e.g. Li et al., 2024; Wu et al., 2025), it can lower the average unit cost of coordinating each interaction (and this particularly applies to external engagements). Thus, coordination costs not only decrease in absolute terms (fewer relationships to manage), but also in relative terms (lower costs per relationship to manage).

Therefore, in this simplified framework where start-up performance depends solely on social capital, the optimal n shrinks and, one can infer that entrepreneur(s) have more possibilities to explore the availability of partners (whoever they are), and meet the (reduced) target, given the decrease in coordination costs (that also include the

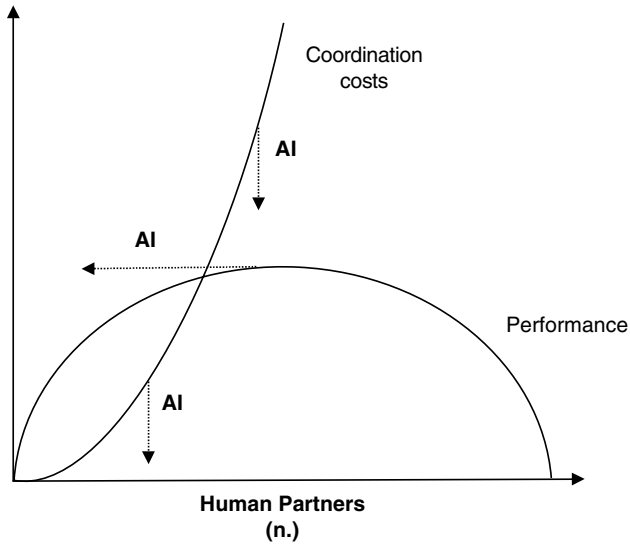


Figure 5.1 AI and start-ups' human partners to manage

cognitive and emotional resources necessary to invest in building any type of human relationships). From this perspective, one can see the reduction in the number of necessary interpersonal ties as good news. If the founder has an increased number of 'shots to fire' (given their reduced cost) – then having fewer 'targets' to hit in order to achieve success (i.e. fewer critical relationships to manage) appears advantageous. In other words, recalling the fundamental dimensions of social capital already introduced in Chapter 1 – namely, *network* and *legitimacy* – the fewer the partners to be involved, the higher the likelihood that even a modest *network reach* may prove sufficient to achieve the intended goal. Moreover, AI and dedicated digital platforms can increasingly support prospective entrepreneurs in efficiently and effectively identifying and connecting with the right stakeholders.[15] This is particularly relevant in the context of the growing digital servitization of economies, which facilitates the crossing of territorial boundaries with unprecedented ease. Looking at the very foundations of

[15] Nowadays, many (if not all) major work and start-up platforms, from Upwork and Toptal to Y Combinator and AngelList, increasingly rely on AI-driven tools to facilitate more precise and efficient matching among entrepreneurs, collaborators, mentors, and investors, thereby reducing traditional frictions in network formation.

entrepreneurial teams – not to mention inter-firm relationships in globalized economies – one clear indicator is the growing number of dispersed founding teams establishing innovative start-ups that are 'remote by design'.[16]

Enhanced Legitimacy?

Clearly, this does not necessarily imply a diminished importance of the other fundamental dimension, that is, *legitimacy*. On the contrary, it is likely the opposite. When two elements are both critical to the success of an innovative start-up, and one – in this case, *network reach* – becomes relatively less decisive, the other is unlikely to decrease in importance and may, in fact, gain relative weight. If the network breadth ceases to be a source of competitive advantage for the reasons discussed earlier, legitimacy may become all the more crucial.

This is particularly relevant given that many of the decisions entrepreneurs make are far from reversible – across various dimensions of the start-up, including, crucially, the network of relationships that define the venture. While AI may indeed expand the reach of entrepreneurs' networks, reduce the number of contacts strictly necessary for effective operations, and lower the coordination costs associated with these interactions, the relatively few relationships that do need to be formed inevitably carry greater relative weight in shaping the company's trajectory (*ceteris paribus*). In this context, relational ties become even more critical, as they often represent non-redeployable investments – sunk costs that cannot be easily recovered or repurposed. This makes 'getting them right' all the more important. Naturally, possessing a certain degree of reputational capital, that is, *legitimacy*, enhances an entrepreneur's attractiveness to external stakeholders, thereby increasing the likelihood of securing the most suitable and strategically valuable partners.

There can, of course, be many sources of irreversibility. Put simply, across all these 'irreversibility domains' in which an innovative start-up is required to operate, the margin for error remains roughly the same:

[16] Since its inception in 2005, Y Combinator – arguably the world's most important and influential start-up accelerator, headquartered in Silicon Valley – has funded over 5,000 start-ups (specifically, 5,235 as of June 5, 2025), with a combined valuation exceeding $800 billion. According to Y Combinator's own classification, more than half of these companies (55 per cent) operate either partially or fully remotely.

an entrepreneur may have fewer targets, and, in principle, more shots to take – yet irreversibility can offset that advantage. This undoubtedly holds true for the dimension we have defined as *internal*: within the entrepreneurial team, but potentially first employees too. As we noted in Chapter 1, early recruitment is often found to define, in a path-dependent manner, the future performance of start-ups (Siepel et al., 2017; Rocha et al., 2018; Gjerløv-Juel and Guenther, 2019). But the same may easily apply also to *external engagements* – just consider the potentially critical role of early investors (as discussed in the next section). In both cases, it is legitimate not to expect any decline in the importance of *legitimacy*. On the contrary, the AI-empowered nature of modern start-ups, along with the growing relevance of the meta-competence of *dynamic AI-enabled orchestration capabilities*, only increases the need to be perceived as legitimate by (the probably few needed) prospective partners, *also (and probably particularly) with regard to this new meta-competence*. This clearly also involves a non-negligible social component. The ability to select and integrate, over time, 'AI entrepreneurial buddies' and to effectively coordinate interactions between these tools and human actors, also depends on the entrepreneur's capacity to emotionally engage other partners, instil human trust in AI processes and solution, and overcome potential algorithm aversion. This dynamic clearly relates to the previously discussed notion of *entrepreneurial savoir-faire*. But the point here concerning legitimacy, is that demonstrated (technological and commercial) competence in this context will act more and more in the future as a powerful 'door-opener', enabling access to high-quality collaborations and resources. Signs in this direction are already apparent, especially among the 'AI(d) offering' start-ups. UiPath (originally a Romanian start-up, then relocated to the area of New York, co-founded by the software engineer Daniel Dines) was the first external platform to receive from Microsoft bi-directional integration with Copilot Studio, explicitly highlighting UiPath capacity to effectively coordinate agents, robots, and humans across complex processes and 'solve unique business challenges' (Brekelmans, 2025). Clearly, as already emphasized in Chapter 1, 'legitimacy begets legitimacy' (Rutherford and Buller, 2007, p. 79), and thus more and more doors can be opened with greater ease (e.g., in the case at hand, one such door was Google; see UiPath, 2025). Another interesting example is Ciroos, a California-based start-up co-founded in 2025 by Ronak Desai (with former intensive *specific* work experience matured at Cisco

and Amazon). Ciroos offers an AI multi-agent system to work alongside employees in site reliability engineering, detecting incidents, diagnosing root causes, and either fixing problems autonomously or guiding humans to a faster resolution. Basically, this innovative start-up aims to provide tailored human-AI teaming model solutions within other firms' operations. It was precisely this proven orchestration capability – their 'cross-domain correlation' between tools and people – that was cited as one of the decisive factors behind their securing of a large seed investment led by the corporate VC firm Energy Impact Partners (Stokes, 2025).

Entrepreneurial Finance: AI, Venture Capital and Beyond

The case of Ciroos and its seed investment leads us directly to the third, and equally important, pillar of success in the field of innovative entrepreneurship: financial capital. Looking retrospectively, VC is often regarded as a fundamental ingredient in the creation of successful innovative start-ups. At the same time, other emerging forms of financing, particularly crowdfunding, are becoming increasingly relevant, helping to complement the broader funding landscape. As introduced in Chapter 1, two are the 'questions': In what ways is the venture capital industry evolving as AI tools become more widespread? Will venture capital continue to play a pivotal role, or could it be increasingly replaced by newer financing mechanisms such as crowdfunding?

Going straight to the core of the nature of venture capital, and with some simplification, the role of this type of investor – beyond the infusion of financial resources – can be summarized in two main components: the *selection* of investments, including deal sourcing, screening, due diligence, and decision-making; and the *coaching* function, namely mentoring, networking, signalling, monitoring, and strategic guidance, aimed at ensuring that the 'promises' of start-ups are realized profitably in the market.[17]

[17] Following the logic of a previous study of mine, Colombo and Grilli (2010), the selection phase involves the VC's ability to 'pick the winners' and provide them solely with the financial resources needed to execute. The coaching phase, on the other hand, encompasses all the additional resources and competencies that the start-up can access through the VC. This guiding principle also explains the inclusion of certain functions in the coaching category, specifically signalling and networking, which might not, at first glance, appear to fall within its scope.

Selection: on the 'Democratization' and 'Commoditization' of VC

Starting from upper part of the funnel, if one were to identify a single stylized fact characterizing its typical functioning, it would arguably be the persistent 'daily flood' of deal flow – that is, the often overwhelming volume of potential investment opportunities that VCs are required to evaluate on a continuous basis. This substantial inflow is, of course, subject to a range of moderating factors, including the broader institutional environment and the reputation of the general partners, among others. However, it inevitably collides, in the vast majority of cases, with the limited, and in any case finite, screening and due diligence capacity available to VC firms. Consequently, and to play on words a little, we could say that the typical selection stage by a VC is indeed selective: the acceptance rate of proposals (leading to investment) usually hovers at just a few percentage points (Andreoli, 2025). To capitalize on their limited 'selection' resources, typically, one of the most important, if not the most important, source of VC deals is 'referrals', investments that originate from 'VCs' networks in some form or another' (Gompers et al., 2020, p. 175). Emblematic, from this point of view, is the example of Apple II we made in Chapter 1.

Thus, two stylized facts stand out. First, as already noted, although venture capital is widely regarded as a critical enabler of entrepreneurial success, in practice only a very limited number of innovative start-ups ultimately secure such funding. Second, there are significant entry barriers even for talented entrepreneurs who lack strong connections to the VC ecosystem. We argue that both of these stylized facts are, to some extent, being challenged by the advent of AI.

And this for a very simple reason: AI technologies allow VCs to lower the costs of selection (see Vismara et al., 2026 and their analysis on the use of LLMs by a German early stage VC fund). This cost reduction surely applies to the stages of screening and due diligence. But, at least in principle, AI can also greatly streamline deal sourcing, enabling both the effective filtering of inbound proposals and the proactive search for promising innovative start-ups. Following a purely economic logic, *ceteris paribus*, this multi-level reduction in the costs of selection should lead the supply side to an increase in the *intensive* margin – meaning that more innovative start-ups can be screened and potentially funded (assuming that sufficiently promising

entrepreneurial projects are available) – as well as an expansion of the *extensive* margin, whereby, at least in theory, more VC firms can profitably enter the market. Even if, admittedly, the use of data science in the VC industry began well before the advent of deep learning (and affects multiple aspects of the entire VC funnel), we may still consider the industry to be at the dawn of this transformation. However, clear indications of this shift are already evident. Bonelli (2026), in his global analysis, shows that VC firms that have hired data-related employees have increased their number of investments per year and assets under management by 30 per cent compared to other VCs. At the same time, a growing number of VC actors can be described as *data-driven* born, highlighting that the shift is not merely quantitative but also qualitative. SignalFire, for instance, has developed a proprietary AI platform, Beacon, to algorithmically source and evaluate start-ups. InReach Ventures uses proprietary AI algorithms for deal sourcing, particularly targeting underexplored regions in Europe. Other prominent examples of this approach could be the US Venture Science and the European-based Fly Ventures (this latter benefited from EU-backed public funding), just to name two more from a long (and growing) list. In the same vein, there is already a considerable number of AI-powered platforms providing incumbent and more traditional VC firms with a range of tools for the evaluation and potential prioritization of investment opportunities – and here too, in some cases, we can speak of 'AI(d) offering' entrepreneurship. These actors range from more established companies (e.g. PitchBook AI) to innovative start-ups such as Harmonic, which uses AI to collect and organize real-time data on millions of start-ups to support more efficient deal sourcing, and VenturePole, which offers a comprehensive suite of AI-driven tools to assist VCs with screening and due diligence.

On the basis of these signals, a growing number of industry practitioners (Chen, 2021; Nikkhoo, 2025), probably less so among scholars, speak of a 'democratization' of venture capital. Whether the term 'democratization' is entirely appropriate remains open to question. What is certain, however, is that – aside from possible cyclical effects – the structural use of AI enhances the investment capacity of VC funds, thereby enabling a greater number of innovative start-ups to receive funding. But then the key question is: could AI also exert a *qualitative*, not just quantitative, impact? In other terms, could the use of AI tools transform the very *identity* of the innovative entrepreneurs who gain access to VC?

The answer is rooted in predictive AI models – or more precisely, all the possible answers to the question depend on considerations concerning the inherent nature of predictive AI.

Clearly, there are all the premises for the industry to become less and less reliant on 'referrals' and to rely instead on more objective criteria in the selection process, certainly allowing even talented individuals without the 'right connections' a greater chance of gaining access to the VC gate. Although this appears to be good news, and to some extent it surely is, it nonetheless conceals certain risks. First, there is a risk of standardizing – if not crystallizing – these identities, a risk that becomes all the greater as VC operators rely on the same training datasets and predictive algorithms, whether developed in-house or supplied by third parties. Moreover, as we emphasized in Chapter 2, if historical data incorporates certain biases, predictive models will perpetuate those biases.[18] However, upon closer examination, it is fair to acknowledge the presence of some mitigating factors for such risks. In fact, there are elements that may discourage VC firms from making strong and unconditional use of AI, without human intuition, due to the persistence of the very same Knightian uncertainty previously discussed, which AI cannot eliminate.

First, the extraordinary success of a start-up is, by its very nature, a rare event – an aspect that underpins the entire VC model, as discussed in Chapter 1, where we noted that 'one win pays for many losses'. Yet, predicting these rare, high-impact outcomes remains challenging, even for advanced AI and deep learning models, while it is relatively easier to predict milder forms of success (like the obtainment of follow-on funding, see Wu and Gnanasambandam, 2017; Chen, 2021; Bonelli, 2026, and partly, Maroof et al., 2025).[19] Secondly, the drivers of success (whether modest or substantial) in innovative entrepreneurship can vary significantly over time, and, as we know, the typical VC cycle spans a long period (generally 7–10 years): using the terminology of Chapter 2, AI models in the

[18] One of the most important biases is certainly related to gender, where historically, female founders appear as highly penalized (e.g. Bittner and Lau, 2021).

[19] The evidence presented in Maarouf et al. (2025) is somewhat more nuanced. Their fused large language model, based on BERT, aims to predict various successful outcomes of innovative start-ups – such as IPOs, acquisitions, and the attainment of external funding, with the first two clearly being much rarer events than the latter. While the model performs reasonably well in predicting IPOs and external funding, its accuracy is significantly lower when it comes to acquisitions.

VC context, and more generally, in the entrepreneurial finance domain (e.g. Ben-David et al., 2025), may easily become 'brittle' due to both *concept drift* and, to some extent, *data drift*.[20] Thirdly, as said in Chapter 2, the unavoidable presence of AI knowledge tacitness immediately invokes the role of human trust, and – depending on its degree – may lead to different attitudes among VC investors towards the investment prospects suggested, always on a probabilistic basis, by the 'machine'. One can therefore imagine the entire spectrum of attitudes, from algorithmic aversion to 'hypersuasion', which – even if not particularly broad within the VC industry – will in any case, result in some heterogeneity of investment policies among VC firms. Fourth, and probably most importantly, we believe that there is a fundamental force pushing against the risk of 'VC commoditization', even in its selection phase: the fact that, as we saw in Chapter 3, after the AI's winning 'move 37', there may always come, sometime later, a human's equally successful 'move 78'. In other words, it may take some time, but human beings can learn to be productively creative also thanks to the use of predictive AI. Indeed, insofar as predictive AI is capable of freeing investors from their personal biases and cognitive heuristics (rather than perpetuating them, as noted in the risk mentioned earlier), it may also encourage them to be even more creative and effective than the machine itself. There are clearly no (yet) in-depth studies on the matter, but the evidence provided by Blohm et al. (2022), upon closer examination, seems to point precisely in this direction.

Is the VC Coaching Function Substitutable?

Thus, the analysis conducted so far reveals two stylized facts – both supported by academic and practitioner-oriented literature. First,

[20] There is also another reason, still related to data, that is often mentioned as an additional mitigating factor with respect to the mechanization of investment decisions. It concerns the traditionally low quality of 'data' that has characterized the VC industry (see Hellmann et al., 2024), where the most important information is often jealously guarded. This, according to some, would make AI models substantially unreliable – even though predictive models of start-up success are already achieving performance that is by no means negligible. In any case, although there is certainly some truth to the opacity of data in the VC sector, we would nonetheless argue that it seems difficult to claim that the process of *datafication* – already discussed in Chapter 2 – does not also affect VC, especially when considering the potential role of synthetic data in improving the predictive capabilities of such models.

while the VC industry has experienced explosive growth in investment volumes in recent years (Janeway et al., 2021; albeit more markedly in some institutional contexts than others), AI makes it possible to further expand the base – that is, the number of innovative start-ups that receive funding, potentially enabling talented but less connected prospective entrepreneurs to gain access. Second, the human component in VC remains essential (Wu and Gnanasambandam, 2017; Chen, 2021; Hellmann et al., 2024; Nikkhoo, 2025). Although fully automated investment decisions may be found in certain segments of the broader and increasingly diversified financial sector (e.g. fintech), such practices remain (and will likely remain) rare in the equity financing of innovative start-ups, due to the specific nature of VC. Importantly, with respect to this second feature, our argument is not only rooted in the selection stage but also pertains to the typical coaching function exercised by VCs.

As we know, the term 'coaching' encompasses a broad range of functions that VCs provide to innovative start-ups. In this respect, we do not exclude the possibility that, for some of these functions, a certain degree of substitution may occur. 'AI entrepreneurial buddies' may be perceived by start-ups not necessarily as perfect substitutes, but at least as imperfect – yet nonetheless acceptable – surrogates. Thus, from the entrepreneurs' perspective, these alternatives may suffice to conclude that 'the game is not worth the candle' when it comes to seeking VC funding, given the associated loss of ownership, loss of control, and the coordination costs (both implicit and explicit) that any founder – VC relationship inevitably entails. But clearly, these possible substitution effects cannot be taken for granted and are probably unlikely in the short to medium term, or limited to 'small-bore' and 'short-haul' entrepreneurial phenomena. First and foremost, it is important to note that just as there are AI buddies for entrepreneurs, the same (or even other) tools and models can act as 'AI VC buddies' – that is, they can support investors in all their coaching activities towards innovative start-ups, in what is often referred to as 'venture building'. In this light, if the superadditive effects brought by the human–machine relationship in the start-up are ultimately determined by the *dynamic AI-enabled orchestration capability*, the substitution effects mentioned earlier become less likely, to the extent that the innovative start-up perceives the added value provided by the VC's *dynamic AI-enabled orchestration capability* as significant, and

potentially superior, to its own. At the same time, some functions can be deemed as simply irreplaceable – at least within the time horizon we are considering – not only because of their intrinsic nature, but also due to the *informal institutions* shaping innovative entrepreneurship. The loss of such functions would require a paradigmatic shift rooted in cultural aspects – one that, as Williamson (2000) reminds us (even without necessarily adopting the timescales he suggests), would undoubtedly take time, even in the face of rapid technological advancement. In fact, securing a seed investment from a VC syndicate led by a prominent firm is widely regarded by all the relevant actors (including entrepreneurs) as a hallmark achievement that other forms of financing – while not necessarily substitutes (e.g. Colombo and Shafi, 2021) – rarely replicate in terms of signalling power. The level of social connections of that very same prominent VC, which can be so important for a start-up to establish relationships with other input providers (Bertoni et al., 2011), are hardly matched by any talented prospective entrepreneur, even those rich in *entrepreneurial savoir-faire*, and by definition are not part of the package offered by more democratic forms of entrepreneurial finance (e.g. crowdfunding).

A broader empirical perspective further reinforces this conclusion: even in the most technically sophisticated segments of AI, venture capital seems to persist as a key source of financing. Box 5.3 provides an illustrative overview of VC investments in deep learning, showing how this type of investor continues to shape frontier innovation sectors.

A further, and in some respects even starker, piece of evidence comes from a simple AI-powered statistical exercise, which again highlights how pervasive the role of VC remains in the AI domain. We asked OpenAI's Deep Research to identify, from the CB Insights list of unicorns, the number of active AI-based unicorns. Among these, we then requested an analysis of the share of AI-based unicorns funded exclusively by VC, those funded by both VC and other forms of financing (such as crowdfunding or initial coin offerings), and finally, the percentage of AI-based unicorns funded solely through crowdfunding or initial coin offerings. We conducted multiple iterations and found that, despite varying the boundaries used to define the AI domain, the results remained consistently stable. As of early 2025, the percentage of AI-based unicorns funded exclusively by VC never drops below 98 per cent; hybrid funding cases stabilize around 1–2 per cent, and there are no AI-based unicorns funded without any VC involvement. The same

Box 5.3 Deep learning and VC

In order to gain a more precise picture of the investments made by various VC actors in the more technical (and likely more highly valued, from the investors' perspective) segment of AI, we can use the same data extracted from PitchBook as presented in Box 5.1. Focusing on the 4,597 deals for which we have the precise geographical location of the invested start-up, 1,678 (36.5%) of these deals are in start-ups located in North America (32.3% in the US); 1,514 (32.9%) in Asia (12.9% in China, 8.2% in South Korea, and 5.3% in India); and 1,054 (22.9%) in Europe (6.2% in the UK, 3.2% in France, and 2.8% in Germany).

The top five urban areas receiving the highest number of deals are the San Francisco Bay Area (473), the Seoul Capital Area (325), Beijing (203), London (180), and New York (161).

As for the deal type reported by PitchBook, out of 4,609 deals, 1,166 (25.3 per cent) are classified as 'early stage'. And although PitchBook is clearly focused primarily on VC investments, the data suggest the very small role of crowdfunding (both equity- and product-based): only thirty deals are classified as such. Unsurprisingly, the primary industrial sector recorded by PitchBook for the 4,609 deals is Information Technology – accounting for 3,323 deals (72.1%) – and nearly 60% of all deals fall within software, followed at a large distance by computer hardware, representing 7.8% of the total. The second most represented industry is Healthcare, with 836 deals (18.1%), of which the 'Pharmaceuticals and Biotechnology' segment accounts for 3.1%.

We have information on deal size for 2,853 deals, which shows an average amount per deal of approximately €10.8 million. Based on the 1,451 VC-backed firms for which we have data on the dimension, during their lifetime (with an average age of about seven years), a start-up receives on average about three funding rounds.

For almost one-third of the 1,476 start-ups, PitchBook reports some patent activity, and for more than 10 per cent of them, the number of patents is equal to or greater than ten. The average number of so-called verticals (definable as 'niche market spaces', see Chattopadhyay et al., 2025, pp. 1076–1077) targeted by the 1,476 start-ups is nearly three, indicating a non-negligible average tendency towards technological versatility – consistent with what has already been highlighted by Chattopadhyay et al. (2025). Besides the obvious Artificial Intelligence & Machine Learning, the most frequently occurring 'verticals' include Big Data, Software as a Service (SaaS), Technology, Media & Telecommunications (TMT), HealthTech, Mobile, Advanced Manufacturing, Internet of Things, Robotics & Drones, and Life Sciences.

exercise, now performed on the entire universe of active unicorn companies and not limited to the AI domain, reveals similar percentages.

In summary, and to put it simply: a world in which the *selection* of investment opportunities is fully democratized through crowdfunding, initial coin offerings, or other future forms; and where the typical VC *coaching* function is completely displaced in favour of entrepreneurs relying solely on their own arsenal of AI tools, is – at least in our view – not just around the corner. AI is indeed a revolution, but not necessarily in this domain. Conversely, while it may be an overstatement for some to claim that the AI revolution is actually driven by venture capital, it is certainly fair to say that the VC actor still figures prominently among its key players – just as it did during the ICT revolution (hence, once again, the aptness of the Twainian aphorism).

'Capitals' Change, Yet Stay

We started this chapter by asking whether the capitalistic pathways to success can still be considered relevant in the AI era. The answer is that skills, effort, and, clearly, some degree of luck remain essential ingredients which, although possibly in different proportions, are unlikely to be overturned by AI. The same line of reasoning applies, in the most promising domain of innovative entrepreneurship, to the three types of 'capital' we have analyzed. In principle, setting aside institutional considerations and focusing solely on technological aspects, there is no indication that these three fundamental drivers will cease to be relevant in the future, perhaps in different proportions than in the past and, as we have discussed, evolving in their internal composition.

Relevant, yes, but for the formation of what kind of innovative start-up? This is the question that naturally follows. Chapter 6 addresses precisely this point, offering an initial phenomenology of start-ups in the AI era. The aim is to move beyond abstract categories and map the potential forms of entrepreneurial activity shaped by the distinctive affordances and constraints of contemporary AI technologies.

6 *Phenomenology of the Start-Up in the AI Era*

A Taxonomy

As with any technological wave, the emergence of start-ups both drives and is driven by specific technological changes, and AI is no exception. Today, we are witnessing the birth and consolidation of countless start-ups whose core offerings are based on AI technologies. While any statistics in this area risk being somewhat inflated by the phenomenon of 'AI washing', driven by the current hype surrounding the term, the underlying trend is undeniably real, and by no means unprecedented in the history of technological innovation. However, the start-ups we typically label as 'AI-based start-ups' – that is, ventures whose products or services are directly built on AI technologies – represent only a part of the broader picture. They are just one among several other ways in which AI can influence the emergence of new entrepreneurial ideas and ventures. In this context, one prediction seems safe: sooner or later (and we are inclined to think sooner rather than later), virtually no start-up will remain untouched by AI; there will be no such thing as an 'AI-free' start-up. All entrepreneurial endeavours – and more generally, all forms of economic activity – will inevitably involve AI in some capacity, much like it is now difficult to imagine a business-related activity that does not, in some form, rely on the Internet.

Yet, without pushing the argument to that extreme, and narrowing the focus somewhat, it is already possible to begin constructing a taxonomy. First by considering the ways in which AI is used, and second by examining the specific capabilities of the AI technologies involved. In fact, AI can be seen as 'an end' where the core offering of the start-up is based on AI, but, AI technologies could also been used as a 'means' to arrive at some entrepreneurial ideas in all economic sectors. Thus, by considering the two AI capabilities – predictive and generative – and whether these types of AI are used merely as tools to create new businesses (even outside AI-related sectors) or instead form the core

of the offering, we can classify all possible AI-related business ideas, including entrepreneurial ones, and their embodied activities, that is, start-ups. Figure 6.1 presents the taxonomy along the two identified dimensions. While the combination of these two dimensions clearly gives rise to four distinct types of businesses, it also allows for the emergence of hybrid forms (e.g. cross-AI function and cross-sectoral configurations). Moreover, insofar as the platform pervasiveness of AI giants continues to expand – along with their capacity to foreclose competition across multiple sectors – the boundaries between these categories are likely to become increasingly blurred, or at least increasingly amenable to being absorbed by the same large firms.

Predictive Analytics

A first quadrant encompasses start-ups offering products or services where the main value derives from the predictive capabilities of AI. These companies typically sell to consumers or businesses, using AI to extract insights from historical or real time data to support decision-making, classification, or forecasting. Of course, they thrive in sectors where data plays a critical role in decision-making, such as finance, healthcare, marketing, and logistics, but, as we know, due to the ongoing intense process of *datafication* across all economic sectors (Sadowski, 2019; Alaimo and Kallinikos, 2024), this list is bound to expand significantly. In the financial and insurance sector, start-ups often focus on credit scoring, fraud detection, and risk modelling. For instance, firms such as Pagaya develop deep learning models that offer dynamic credit evaluations, while others provide real time fraud analytics and predictive pricing tools for insurers based on behavioural data (e.g. Shift Technology). In healthcare, platforms exemplified by Aidoc employ deep learning techniques, often integrated into always-on clinical AI systems, to analyze radiological scans in real time, flag urgent or critical findings, and thereby support early diagnosis and triage, while assisting radiologists and broader clinical workflows. In the life sciences domain, a start-up such as Peptone has capitalized on both the limitations and the widespread acceptance of AlphaFold predictions (Eisenstein, 2023), using deep learning techniques to model intrinsically disordered protein regions (areas largely not addressed by Google's AlphaFold), thereby enabling novel drug discovery applications. Similarly, Cradle, an innovative biotech start-up founded in 2021, employs generative deep

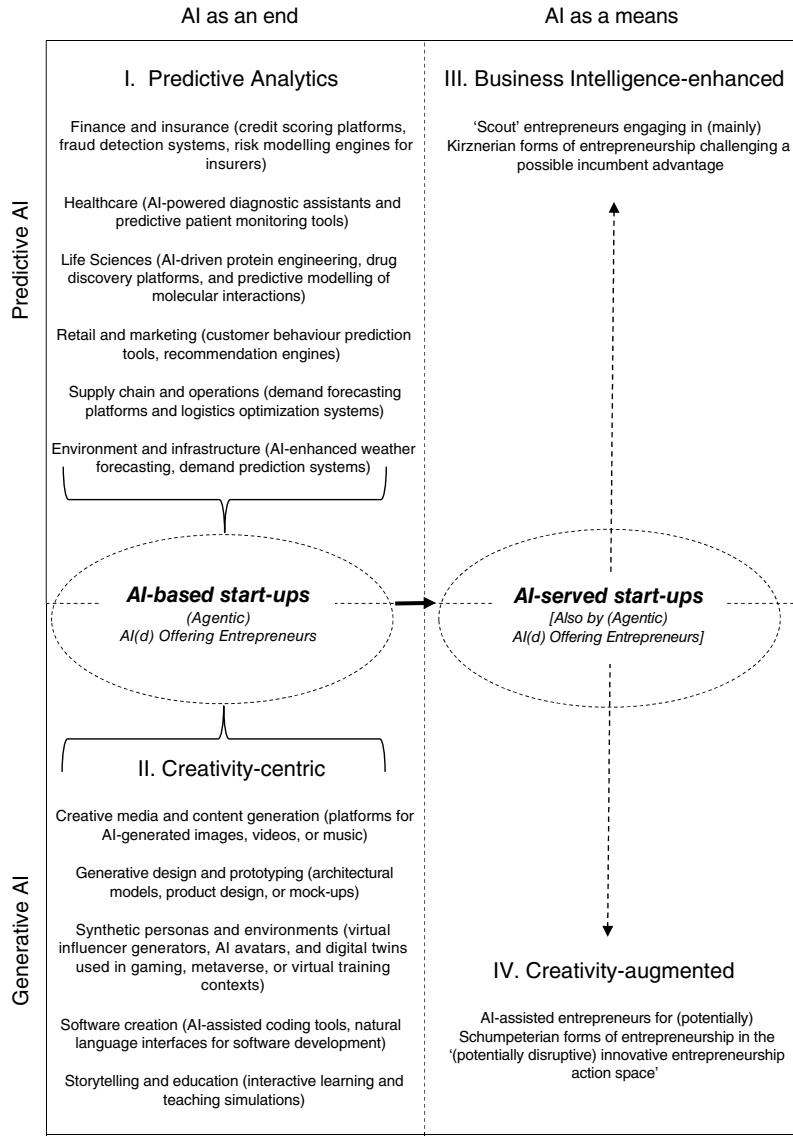


Figure 6.1 AI and start-ups: a taxonomy. Two-by-two matrix classifying start-ups in the AI era by predictive vs. generative AI and AI as an end vs. as a means. Illustrative domain examples for AI-based start-ups provided for clarification; not exhaustive.

learning models in the context of protein engineering, aiming to predict how specific sequence modifications may affect protein properties and functionality. Continuing this overview, and as already noted in our discussion of 'AI entrepreneurial buddies', in retail and marketing many AI-based start-ups offer predictive models to anticipate consumer behaviour – including churn likelihood, purchase intent, or customer lifetime value – as well as systems that optimize promotional timing and pricing based on individual user profiles and real time engagement. In the area of supply chain and operations, start-ups may forecast demand fluctuations, detect bottlenecks, and optimize logistics routes. Then, to conclude what is, in principle, a much longer list (see, among others, Moosa, 2025 and his chapter 3, for a discussion of additional domains), start-ups active in environmental and infrastructural applications represent a growing frontier for predictive analytics: AI-enhanced weather forecasting platforms, such as Climavision, are now being used in agriculture, energy, and aviation to improve planning and resilience. Similarly, utilities and smart grid operators are adopting AI tools offered by more mature innovative start-ups such as Grid4C and AutoGrid, as well as exploring newer solutions from entrants, like Capalo AI, which aim to predict energy consumption patterns, optimize battery and load operations, and detect system anomalies.

Creativity-Centric

Innovative start-ups in this second quadrant leverage generative AI to offer users systems that can (co-)produce (with humans) new content, designs, or media artefacts. In this case, the AI system itself is the core product, positioned as a creative or co-creative tool that generates value through its ability to produce novel outputs across various media, formats and knowledge required (on the human side). Generally, these start-ups operate in the upper layers of our 'bottom-to-top' representation of the generative AI supply chain (see again Figure 2.2) and appeal to creators, marketers, software developers, educators, and other enterprises seeking personalized, scalable, or low-cost alternatives to traditional content or design workflows. Only looking at the very specific and narrow domain of 'business idea generation', as already noted in Chapter 3 (see Legnani et al., 2024), a wide range of AI tools has emerged to support this purpose. It is worth emphasizing here that, in the vast majority of cases, such tools

originate precisely from the world of innovative entrepreneurship, in a sort of 'from start-up to start-up' trajectory.

Beyond this, many indeed are the emblematic application domains for creativity-centric start-ups. In the area of creative media and content generation, a broad set of platforms enables the automated or semi-automated production of images, videos, and music, from AI-generated visuals for advertising campaigns to tools for producing (royalty-free) soundtracks and synthetic video content. In the space of generative design and prototyping, many innovative start-ups were born as tools to transform early sketches into 3D renderings or to accelerate the development of architectural, industrial, and digital mock-ups. Several of these have already been discussed in Box 4.1 as valid 'AI entrepreneurial buddies' in the implementation of the 'SoleScribe Snack Shoe' project – such as tools for visual prototyping (e.g. Vizcom), automated branding (e.g. Looka), or AI-driven marketing asset generation (e.g. AdCreative.ai, Jasper AI, Copy.ai). Other creativity-centric start-ups focus on the creation of synthetic personas and immersive environments. For example, Didimo offers a platform to generate lifelike 3D avatars from a single photo, supporting both entertainment and training contexts. Similarly, Hour One turns text into photorealistic video of AI characters, used in education and corporate communications. In adjacent sectors, Inworld AI develops interactive characters with memory and personality, targeted at gaming studios and virtual world builders. In the field of storytelling and education, generative platforms such as Sanas and ElevenLabs provide synthetic voice and language localization services, while others like Curipod or Khanmigo (developed by Khan Academy in partnership with OpenAI) aim to transform the classroom with interactive, personalized teaching tools based on LLMs and generative content. The intersection with software creation has also become a highly dynamic space. In addition to the platforms already mentioned in Box 4.1 (e.g. Replit and Adalo), new entrants such as Cogram (focused on automating technical documentation and business meeting summaries), Uizard (which turns hand-drawn wireframes into editable digital prototypes), or the already-mentioned Cursor (an AI-native code editor integrating conversational development workflows) exemplify how natural language and sketch-based input can drive co-creative software development. Similarly, Durable enables entrepreneurs to spin up entire websites with just a short description of their business, blending web design and generative copywriting into a single interface.

The list is (again) long. And it is important to emphasize here that, exactly as in the case of the 'predictive analytics' category, the above cases should be understood as indicative rather than exhaustive. In the long run, it is fair to say that the range of domains in which such creativity-centric AI start-ups may emerge will, much like in the previous case, become virtually boundless.

Cross-AI Function, Cross-Sectoral Domain, (Agentic) 'AI(d) Offering' Innovative Start-Ups

Moreover, it is reasonable to anticipate – and we already observe – a progressive blurring of boundaries (and areas of activity) among start-ups originally dedicated to a single AI function, in favour of a multifunctional approach: some ventures combine predictive and generative capabilities, others add early agentic features to one or both, and still others integrate retrieval-based search into this mix. Thus, before moving on with the analysis of the taxonomy, it is worth noting that the line separating the first two quadrants is expected to become increasingly porous. In practice, many start-ups that originated within either the predictive or the generative domain now integrate multiple capabilities, partly because of the underlying overlap in deep learning techniques (and their progress, as discussed in Chapter 2) and partly because organizations increasingly demand solutions that combine prediction, generation, retrieval, and interaction within a single offering. This progressive blurring of lines and the consequent growing multifunctionality help explain the rise of a number of 'AI(d) offering' innovative start-ups whose products operate across sectors and across AI functions, in some cases already exhibiting strong agentic features.

For example, DataRobot, founded in 2012 by two data scientists, Jeremy Achin and Tom de Godoy, is an AI platform that by relying on multiple foundation models automates the building, deployment, and management of machine learning models for a variety of organizations. Specifically, companies in the energy sector use DataRobot for load forecasting, grid optimization and predictive maintenance, while e-commerce platforms use DataRobot for inventory management, demand forecasting and price optimization. In the same vein, C3.ai, founded in 2009 by Thomas Siebel, provides a suite of enterprise AI applications that integrate both predictive and generative capabilities to serve clients across manufacturing, defence, utilities,

and healthcare. Its models are used for fraud detection, energy management, predictive maintenance, and for generative document summarization and natural language interaction with complex industrial data. More recently, Adept AI, founded in 2022 by a group of former Google and OpenAI researchers (including Vaswani, first author of the 'transformer' paper), has developed ACT-1, a model designed to interpret natural-language goals, navigate software interfaces, and execute multistep workflows across enterprise applications, thereby combining different functions with the ability to perform actions autonomously and displaying agentic capabilities in enterprise-oriented AI systems.

These three cases – DataRobot, C3.ai, and Adept AI – illustrate how 'AI(d) offering' start-ups (or former start-ups) are increasingly serving businesses through modular, sector-agnostic platforms capable of handling distinct tasks.[1] Yet similar logics are also being applied in user-facing contexts. For instance, Abacus.AI combines predictive models with generative tools that help both organizations and individual users build and deploy recommendation systems, forecast outcomes, and interact with model outputs through natural language explanations. An equivalent illustration of emerging cross-functionality with early agentic features is Inflection AI and its product Pi, an 'emotionally intelligent' general-purpose personal assistant spanning a series of different activities (with the start-up subsequently entering into a major licensing and 'acqui-hire'-style deal with Microsoft). Similarly, Perplexity.ai, launched in 2022, blends retrieval-based search with generative AI to offer individuals an AI-powered alternative to traditional search engines – integrating search, summarization, and conversational interaction into a single interface and, increasingly exhibiting agentic features.

Business Intelligence-Enhanced and Creativity-Augmented AI-Served Start-Ups

Coming to the third and fourth quadrants of our taxonomy, we turn our attention to AI-served start-ups, that is, innovative ventures in which AI acts primarily as an input – a key enabler of internal

[1] Another example that can be seen as falling within this trajectory is UiPath, an 'AI(d) offering' start-up already discussed in Chapter 5.

processes – rather than as a market-facing output or product. These start-ups leverage AI technologies, which may be developed partially or entirely in-house and/or acquired from the market (and, if necessary, customized), with the primary aim of enhancing their internal capabilities in areas such as business intelligence and creative production.

As argued earlier in the book, in the third quadrant, associated with predictive AI, the strategic advantage tends to rest with incumbent firms. In fact, as we has already said, established players may typically enjoy privileged access to large-scale, often proprietary, datasets that are continuously generated through their ongoing market operations. This steady and almost endogenous flow of data provides a structural advantage in the development and refinement of predictive models, reinforcing what we have referred to be the *incumbent advantage*. Therefore, forms of Kirznerian entrepreneurship in this domain, while not necessarily precluded, must nonetheless contend with the challenges outlined earlier.

Conversely, in the fourth quadrant, which corresponds to the use of generative AI, the focus shifts towards the enhancement of creative capabilities, offering a broader opportunity space for new ventures to experiment, differentiate, and innovate. As we have already emphasized on several occasions, the lowered barriers to content creation and design enabled by generative AI may open what we have defined as the '(potentially disruptive) innovative entrepreneurship action space' – a zone where innovative start-ups can more credibly challenge existing market paradigms established by incumbents.

One sector (among others) in which both dynamics are likely to manifest themselves – that is, the challenges of entrepreneurship in the former case, and the significant opportunities for innovative venturing in the latter – is the fashion-design segment, where both faces of the AI-served archetype are clearly visible.

On the predictive side, an innovative start-up such as the Paris-based Heuritech, founded in 2013 by Tony Pinville and Charles Ollion (both holding PhDs in AI), was created to build a computer-vision pipeline capable of analyzing more than three million social media images daily, detecting approximately 2,000 style attributes to forecast the direction and velocity of upcoming fashion trends. The firm had to rely primarily on publicly available imagery, lacking access to key proprietary data streams – such as sell-through data, radio-frequency identification (RFID) signals, and supply chain tracking – that, typically,

incumbent fast fashion giants like Zara can integrate in real time across thousands of stores worldwide. The wide availability of such points of sale on a near-continuous basis provides a structural advantage in training and fine-tuning any possible demand-forecasting model. In other words, this resulting steady and endogenous flow of data echoes Skillman's observation quoted earlier (see Chapter 3) and encapsulates the *incumbent advantage*. Thus, while Heuritech can arguably be ascribed to the list of successful innovative start-ups (at least by current criteria, see *infra*, given that it was acquired in 2024), it tended to position itself as an analytics supplier to the very incumbents it could otherwise have aspired to disrupt. Indeed, many of its clients were well-known incumbent brands (e.g. Adidas, Dior, and Louis Vuitton, among others), whose own strategic capabilities may have been further reinforced by Heuritech's 'AI(d) offering' services.

Conversely, on the generative side, we have already introduced several examples in Box 4.1 of 'AI entrepreneurial buddies' of a generative nature, which clearly serve as powerful means of augmented creativity. We have also already introduced Syntilay, on which our case illustrating the relevance of *dynamic AI-enabled orchestration capability* for the implementation of the 'SoleScribe Snack Shoe' is largely based (see Box 5.2). Recently launched (i.e. 2025) by Ben Weiss (with Joe Foster, co-founder of Reebok, as strategic adviser), independently of the success it may achieve, Syntilay in our view well exemplifies the point about the grounding of Schumpeterian entrepreneurial ideas – (at least potentially) more disruptive – that the use of generative AI can allow even for innovative start-ups outside the AI domain. The start-up combines generative AI tools (e.g. MidJourney and Vizcom) to generate print-ready files very rapidly, which are then outsourced to the 3D-printing specialist (Zellerfeld), while a smartphone foot-scan ensures each individual fit. The ultimate result is a fully customized shoe, made-to-order and delivered without the need for a traditional large-scale manufacturing infrastructure.

Thus, taken together, Heuritech and Syntilay illustrate the bottom line in our taxonomy: when predictive AI relies on privileged, firm-specific datasets and/or (closed-source) models, incumbents may easily retain the upper hand; whereas to the extent that generative AI becomes more accessible to all – increasingly so, thanks to intuitive and designer-led prompt engineering – and remains affordable (as it generally is), the 'innovative entrepreneurship action space' opens up,

allowing innovative start-ups to experiment more freely and challenge incumbents more credibly (or at least from a less structurally disadvantaged position).

AI-Served Start-Ups and the Risk of AI-Vertical Foreclosure

More fundamentally, upon closer inspection and at a slightly higher level of analysis, what truly influences innovative entrepreneurial processes is not merely the type of AI input, but rather its accessibility. That is, the difficulty of engaging in innovative entrepreneurship within the 'predictive AI input and Kirznerian entrepreneurial ideas' framework persists only insofar as there is a deliberate (and clearly comprehensible) choice by the incumbent to exploit the *irreplicability* and *uniqueness* of (its own) data, as well as the *closeness* of the underlying models (see Chapter 2). This applies to sectors traditionally less exposed to advanced technologies (but which, as noted earlier, are now inevitably becoming more, if not entirely, high-tech due to the advent of AI), but it is perhaps even more evident in domains that have historically been closer to science and naturally more knowledge-intensive.

In fact, the case of Google DeepMind's AlphaFold – specifically the transition from its fully open version 2 to version 3, which was initially fully closed (and later reopened, with, however, persistent commercial limitations for third parties) – is emblematic in this regard. While the earlier versions were widely hailed as a democratization of the biotech sector (Hume, 2025), offering significant development opportunities to biotech start-ups and levelling the playing field between innovative start-ups and large pharmaceutical companies, the strategic shift seen in version 3 – driven by Google's (or more precisely Alphabet's) own biotech ambitions through the founding of Isomorphic Labs in 2021 – marked a turning point with potential backlashes for innovative entrepreneurship. It is evident, in fact, that not many start-ups can compete with Google if the latter alone possesses the most advanced, cutting-edge AI models for predicting protein interactions and exploits them directly on the market (or indirectly through multi-billion-dollar deals with large pharmaceutical companies). In essence, Alphabet's dual role as both AI tool provider and market actor could result in a kind of 'vertically integrated AI monopoly' in many different domains, where the

best models are available only to Google and its affiliated partners. Indeed, biotech is only one among the almost infinite number of sectors potentially touched by the AI revolution. The more Google's AI ecosystem – and that of MAMAA or other AI giants – expands its boundaries, becoming increasingly horizontal as well as vertical, the more concrete the risk becomes of multiple AI monopolies or tight oligopolies emerging across a wide variety of sectors, dominated by the very same firms, which become stronger and stronger due to indirect network externalities and cross-sectoral AI data-driven learning effects.

AI-Based Entrepreneurship: Consolidated Patterns, Three Stylized Facts, and One Provocation

The above-mentioned risk is all the more concrete given that (some of) these Big Tech firms are also undisputed leaders in the domain of AI science itself (see Box 2.1), effectively positioning themselves as a sort of integrated whole along an extended supply chain that stretches from basic scientific research all the way to the market. This, however, is an aspect we shall examine in greater detail later on in this chapter. For now, and proceeding step by step, we begin by outlining basic (albeit largely unsurprising) geographical and temporal patterns associated with AI-based entrepreneurship, for then focusing on three stylized facts which, though considerably less predictable, appear to encapsulate some of the most distinctive characteristics of the phenomenon. The final point regarding the third stylized fact will lead us to put forward a possible provocation.

Patterns

Despite the variety of sources that can be used – such as PitchBook, Crunchbase, CB Insights, the Stanford AI Index or OECD innovation scoreboards – the different methods and metrics they employ, and, more fundamentally, the varying degrees of stringency in defining what constitutes a start-up whose core business is based on AI (and which therefore may reflect to a greater or lesser extent the surrounding hype), the available evidence consistently shows that AI-based entrepreneurship has experienced a sharp surge since the years following 2005. Particular growth has occurred in specific

segments of the ecosystem, such as data analytics, computer vision, and AI for enterprise software, whereas, as we know, the segment of start-ups focused on generative AI emerged at a later stage, with rapid acceleration beginning around 2022. For the reasons outlined above, estimates remain difficult and probably pleonastic. Nonetheless, in order to provide a sense of scale, suggesting that, at the time of writing (i.e. 2025), the number of currently active AI-based start-ups of some significance worldwide lies between 50,000 and 100,000 does not seem to be an overstatement. If anything, it may well be a rather conservative approximation. From a geographical perspective, the United States clearly leads the way (as we have already started to observe in Box 5.3). Even when looking at a more granular level, there are few surprises: the areas of San Francisco, and New York (and Greater Boston to a lesser extent) remain the most developed. The fact that these are also the three regions in which the US venture capital industry has historically been concentrated (Grilli, 2025) can hardly be regarded as a mere coincidence. After the United States, the other major player in AI-based entrepreneurship is China, followed (by most measures) by India, Japan, and South Korea; while the European context appears relatively lagging behind, with probably the only exception of the United Kingdom.

Moreover, if one moves beyond the mere phenomenon of start-up creation to also consider their subsequent success, the location of AI-based start-up unicorns broadly mirrors the geographical patterns observed for start-up births.

Indeed, based on the most generous estimate of the exercise we performed in Chapter 5, through which OpenAI's Deep Research identifies, from the CB Insights list of unicorns, those falling within the scope of AI-based ventures (essentially based on PitchBook and Fortune), there are approximately 370 AI-based unicorns globally. The majority are located in the United States and China, followed by the United Kingdom and India. Unsurprisingly, the pattern suggests a strong association between origin and success in AI-based entrepreneurship.

Beyond these broad temporal and geographical patterns, then, a number of deeper regularities begin to emerge – regularities that help elucidate the distinctive mechanisms underpinning AI-based entrepreneurial dynamics, whether operating in the 'predictive analytics' or 'creativity-centric' quadrants, or in hybrid forms that may also display agentic features.

1° Stylized Fact – AI-Based Innovative Entrepreneurship Is Strongly Rooted in AI Science

AI science is not only central to how innovation in AI unfolds, but is also shaping the dynamics of innovative entrepreneurship. In other words, there is a fairly strong correspondence between the scientific domain of AI and the ability to 'generate' AI-based ventures – something that, in principle, would be consistent with a SM I dynamic (although, as we shall see later, potentially stripped of its creative-destructive force). Following the KSTE, we have already indirectly highlighted the role of the firm as an 'incubator' of 'unexploited knowledge' made available for AI-based innovative entrepreneurship in Chapter 5 (i.e. the role of specific work experience in deep learning start-ups, see Box 5.1). Here, we emphasize – still in line with the KSTE – the role of AI science in this respect. As mentioned, it appears to be a significant one, further supported by the fact that its relevance is often corroborated through the use of diverse proxies.

We do not aim to be exhaustive in this account, but one of the first particularly compelling pieces of proof, albeit somewhat indirect, is offered by Gofman and Jin (2024). Their study investigates the effects of the so-called AI brain drain, namely the migration of leading AI professors from academia to industry, on the entrepreneurial outcomes of their former students. Analyzing a novel dataset of over 500 AI entrepreneurs and more than 400 AI start-ups in North America, the authors find that when tenured faculty, particularly those affiliated with top-ranked institutions and specialized in deep learning, leave academia, students from the affected universities are significantly less likely to found AI-based start-ups, and when they do, they tend to raise substantially less early stage funding. The authors attribute this effect to a loss of advanced, domain-specific knowledge, which diminishes students' ability to acquire the technical expertise and research mentorship necessary to compete effectively in the AI landscape. Notably, the negative impact is strongest when faculty departures occur before students enrol, indicating that even mere exposure to high-level AI academic expertise plays a crucial role in shaping entrepreneurial capability in that very same domain. The (this time positive) flipside of the coin is that founding teams including at least one academic have been found to be more 'generalist', that is, their offerings tend to span a broader market scope compared to non-academic start-ups. More specifically, Chattopadhyay et al. (2025) show that academic founders tend to formulate and solve

more general problems, creating technologies grounded in broadly applicable scientific knowledge. As a result, their ventures are associated with greater potential for market applicability, and possibly, in this sense, with more disruptive potential. Overall, viewed from different angles, both studies underscore the profound interdependence between scientific research and AI-based entrepreneurship.

More in line with the KSTE and the typically localized nature of knowledge spillovers (Jaffe et al., 2003), or, if preferred, the common 'home bias' of entrepreneurs (Michelacci and Silva, 2007), the relationship is also documented at the territorial level – this time in a European context, as opposed to the US setting discussed earlier – whether AI knowledge is proxied by AI patents (D'Alessandro et al., 2026) or by the presence of universities specialized in AI-related subjects (Colombelli et al., 2023). The former shows that regional stocks of AI patents significantly foster innovative start-up formation, but only within AI-related industries. The latter finds that the presence of universities specialized in computer science and engineering is a key driver of local AI-based start-up emergence in Italy. The role of AI science in fostering the emergence of innovative AI-based start-ups (at the local level) is also confirmed by our own independent analysis, presented in Box 6.1, which focuses on yet another proxy, namely scientific publications in the domain of AI.

Box 6.1 From AI scientific papers to AI-based start-ups

In 2025, as part of a joint research effort, I conducted a study examining the 'geographical' relationship between the production of scientific papers in the field of AI and the subsequent creation of AI-based start-ups over time. Specifically, we retrieved publications related to AI from Scopus for the period 2005–2022, using the term 'Artificial Intelligence' in their title, abstract, or keywords. In parallel, we identified AI-based start-ups in Crunchbase by searching for the same term within company descriptors. With regard to the former dataset, we collected 305,482 scientific documents classified across eighteen disciplines, approximately two-thirds of which originated from the core computational fields of computer science, mathematics, and engineering. From Crunchbase, we gathered information (including precise location data) on 55,087 AI-based start-ups founded between 2005 and 2023. Both datasets were mapped in a common geographic, temporal, and thematic space using geoparsing (via CorText Manager; see Breucker et al., 2016) and

text-mining techniques. The geoparsing process identified 14,088 city-level urban locations (including major urban areas; hereafter 'cities') producing AI publications and 2,004 cities with AI-based entrepreneurial activity. The text mining searched for keywords derived from the taxonomy proposed by Samoili et al. (2020) within publication metadata and start-up descriptions, and used them for thematic classification. AI was thus analyzed as a whole, but also disaggregated into its three main categories: hardware, software, and the specific case of perception technologies (image and video processing).

A series of panel data estimations show that the number of scientific publications is highly positively and significantly correlated with the creation of AI-based start-ups across all four cases: AI as a whole, hardware, software, and perception technologies. In other words, the number of AI-based start-ups founded in a given city c, year t, and technological domain d is found to be strongly correlated with the number of scientific publications produced (or co-authored) in the same city c and year t, within the corresponding AI domain d. In particular, looking at AI as a whole, in the random-effects model, the coefficient for number of publications (0.056, $p < 0.001$) indicates that each additional scientific publication is associated, on average, with an increase of about 0.06 start-ups in the same city–year. The fixed-effects estimate, although smaller in magnitude (0.028), remains highly significant and confirms that the relationship also holds within cities over time, once persistent local characteristics are controlled for. Considering both models, this roughly implies that between 14 and 21 AI-related scientific publications are associated, on average, with the creation of the first new AI-based start-up. The domain in which the estimated effects are strongest is software. According to the random-effects estimates, each additional publication in this domain is associated, on average, with an increase of about 0.10 start-ups in the same city–year. Thus, in this latter case, it 'takes' only around 8 AI software-related scientific publications to 'make' the first AI-based software start-up (or about 13, when considering the fixed-effects estimates).

These results remain robust to several more sophisticated model specifications, as well as to the introduction of time lags in the independent variable (number of publications). While (obviously) no causal claim is made, this remarkably consistent co-evolution – observed at such a fine-grained geographic level and with a technological detail rarely achieved in the KSTE literature – provides clear evidence that the KSTE framework, particularly in its scientific dimension, applies strongly to AI-based entrepreneurship.

2° Stylized Fact – The '(Golden) Talent Trap' Affecting AI Engineers, Together with the Increasing Automation of Software Programming, Jointly Risks Generating Adverse Selection in AI-Based Innovative Entrepreneurship

As already discussed in Chapter 1, and in line with an SM I view of the entrepreneur, the decision to become an entrepreneur may not be entirely shaped by purely economic motivations, but it is difficult to assert that such motivations do not represent a fundamental component in orienting this choice. Following this perspective, a potential entrepreneur (an innovative one, in our case) will decide to take this path if the opportunity cost, that is, the expected return as an employee in the labour market (wage), is lower than the expected financial return from entrepreneurial activity. In formulaic terms, drawing, among others, on Eesley (2016) and with some simplification, we can say that (innovative) entrepreneurship will be pursued by an individual only if:

$$NPV_{Employment} < (1 - \delta) NPV_{Start-up} - C_0$$

The net present value of expected earnings (wages, in the first instance) in the labour market must be lower than the expected net present value of the venture's future returns, scaled by the factor $(1 - \delta)$ to reflect the entrepreneur's retained equity share, net of entry costs.[2]

The second stylized fact, increasingly evident in the more advanced segment of AI-based entrepreneurship, particularly in *deep learning*, is that there are growing indications that the inequality above is becoming less favourable to entrepreneurship for highly skilled AI engineers (or, more generally, AI talent in charge of the human architecture of the digital 'statistical machines', see again Figure 2.1), while turning more favourable for individuals with only intermediate or basic computational and statistical skills. This dynamic raises the risk of a progressive 'adverse selection' process in the long run, whereby the most qualified technical profiles are less likely to engage

[2] While we remain within a strictly economic logic, it is clear that a broader perspective – one that also accounts for more hedonic or intrinsic motivations (see Chapter 1) – would require this basic equation to be accordingly extended. It is also clear that the inequality as it is written holds for a risk-neutral individual. We could of course complicate the framework to embrace all different possibilities, but such extensions, however, would not entail any substantial disruption to the core reasoning developed here.

in AI-based entrepreneurial activity. The risk stems primarily from what may appear to be a contradiction, but in fact reveals a growing segmentation within the software engineering labour market. In fact, while incumbent firms are increasingly laying off software engineers in the wake of productivity gains brought about by gen-AI tools (e.g. GitHub Copilot, automated testing, code generation platforms), some of these same firms – and here we are particularly referring to AI giants and other well-funded incumbents – are aggressively recruiting elite AI engineers at unprecedented wage levels (see, among others, Rikap, 2024b). This reflects a bifurcation in the labour market, where some programming work is being automated (or, in any case, largely 'commoditized', see Box 7.2 in the next chapter), while deep AI expertise remains in short supply and is becoming increasingly central to the corporate strategy of these incumbents.

Thus, looking at the inequality above, it is indeed true that (as we know) a high level of human capital should, on average, correspond to a relatively higher net present value of the start-up. However, at the same time, the opportunity cost of founding an innovative venture in the deep learning field has certainly increased, and we believe disproportionately so. Moreover, launching such a venture may involve increasingly high entry costs and, particularly in the case of high-potential start-ups, require substantial upfront investment. This, in turn, could lead to greater dilution of ownership for the prospective entrepreneur.[3] Conversely, the opportunity cost for individuals with relatively lower levels of software-based skills has certainly decreased, for the reasons discussed above. The logic, which follows the opposing forces acting on the terms of the inequality, is summarized in Figure 6.2.

[3] Such dilution may occur not only through equity allocated to external investors, but also through 'work for equity' contracts, which are often used to attract specialized talent, particularly AI engineers and data scientists, in financially constrained start-ups that cannot offer competitive salaries, especially in light of the now astronomical wage levels offered by AI giants, and the MAMAA group in particular. As a well-known and telling example of this broader and well-established trend, in 2025 Meta reportedly offered a rich four-year contract to Matt Deitke, a PhD student at the University of Washington and co-founder of the AI-based start-up Vercept. According to several media outlets, including The New York Times, the total package was valued at 250 million US dollars, with 100 million allocated in the first year alone (Isaac et al., 2025).

	Opportunity costs post AI ('Talent trap')	(Economic) benefits and returns post AI	Comparison between opportunity costs and (economic) benefits post AI	Increased risk of 'adverse selection' post AI
	$NPV_{Employment}$	$(1-\delta)NPV_{Start\text{-}up} - C_0$	All combined, it is more likely that the following condition will be verified for more individuals:	Differentials in terms of number of entrepreneurs:
AI-engineering skills				
High	↑↑		$NPV_{Employment} >> (1-\delta)NPV_{Start\text{-}up} - C_0$	
		δ ↑↑ and C_0 ↑↑ (probably more so, the higher the $NPV_{Start\text{-}up}$, certainly not diminishing with it)		$\Delta^{H}_{Entrepreneurs} < \Delta^{L}_{Entrepreneurs}$
Low	↓↓		$NPV_{Employment} << (1-\delta)NPV_{Start\text{-}up} - C_0$	

Figure 6.2 '(Golden) talent trap' and 'adverse selection' in advanced AI-based entrepreneurship. Conceptual diagram illustrating how, in the post-AI context, rising opportunity and entry costs for highly skilled AI engineers (compared to lower-skilled ones) increase the risk of adverse selection in advanced AI-based entrepreneurship.

Thus, as an ultimate consequence, the entrepreneurial ecosystem faces, at least in the short term (i.e. until the labour market adjusts) a growing risk of brain drain: as we will further elaborate in the following sections, the best minds are increasingly drawn into the gravitational pull of AI giants, depriving the AI start-up space of some of its most promising founders, while simultaneously encouraging the entry of medium- to low-skilled software programmers into the entrepreneurial arena.[4]

3° Stylized Fact – AI Giants Shape AI-Based Innovative Entrepreneurship at Multiple Levels

The influence of AI giants, particularly the MAMAA companies, over the phenomenon of AI-based entrepreneurship is undoubtedly significant, precisely because it operates across multiple and intersecting levels, which may easily reinforce themselves. In the previous section, we have referred to the risk of AI-vertical foreclosure in relation to Google and the degree of openness of its AlphaFold model, but the domain of AI science is only one among several possible areas of influence, strategic direction, and potentially even chokepoints for the foreclosure of innovative start-ups.

The global and ecosystemic nature of these AI giants, from basic science all the way down to worldwide consumer applications, is well-emphasized in the literature. Various terms and constructs have been used to describe and analyze this level of dominance in AI, ranging from 'ecosystem orchestrators' to 'global firm-based ecosystems' (e.g. Jacobides et al., 2021), and passing through 'corporate innovation systems' (Lundvall and Rikap, 2022; Rikap, 2024b), among others. The common thread running through these and other definitions is that (some) Big Tech firms, by potentially covering the entire AI vertical technology stack as outlined in Chapter 2 – and exerting particular dominance over its infrastructural and production layers – have succeeded in creating ecosystems (also thanks to the presence of various indirect network externalities operating at multiple levels) where others may participate and even profit, but only within limits. The lion's share

[4] After all, adverse selection phenomena related to the entrepreneurial choice, affecting specific sectors and driven by particular institutional factors or contingent market conditions, have already been documented (e.g. Colombo and Grilli, 2017).

of value in these increasingly pervasive AI ecosystems remains firmly in the hands of the platform leader, whose reach continues to expand as it orchestrates and absorbs the contributions of the various actors. The dominance of these tech giants is, of course, well-known beyond the context of the 'AI revolution' and clearly precedes it by some margin (e.g. Ezrachi and Stucke, 2022). Yet AI does not, at present, appear to represent a point of rupture in this trajectory. On the contrary, it seems to mark a moment of further entrenchment and amplification of their position. See Jacobides et al. (2021, emphasis mine):

> A related observation is that, although Big Tech firms' business models benefit directly from AI and create value for the entire sector (e.g. social media or digital marketing), they also generate business for their complementors. As such, the growth of digital ecosystems (Adner, 2017; Jacobides et al., 2018) goes hand in hand with the growth in AI. Firms that collaborate with Big Tech find ways to benefit from Big Tech's data. This […] suggests that Big Tech operate as 'kingpins' (Jacobides and Tae, 2015) – firms that create benefits for themselves and their segment (or, here, their complementors) by advancing technology while *skewing the distribution of profits*. This leads to *a feedback loop between technological edge, resource accumulation, and market dominance*. (Jacobides et al., 2021, p. 421)

More specifically, we can identify *five* main areas of action through which the AI giants may exert their 'ecosystemic power' and, in doing so, both directly and indirectly, influence the dynamics of AI-based entrepreneurship (and, to some extent, those of the AI-served type as well).[5] The first, which we have already touched upon, is *science*. In fact, AI science is not only produced through close collaboration between academia and industry, as we have noted repeatedly (a relatively rare

[5] It is fairly common in the literature to treat these AI giants, in their ecosystemic capacity, as a single bloc (see Rikap 2024b for a similar view, as well as for one of the few exceptions), and we shall likewise not refrain from doing so – even though there are (currently) significant differences in the strategies pursued by AI giants within (and across) the five areas identified (see Srnicek, 2025). Rivalry between them within each of these areas – and beyond – can be intense, yet it does not preclude technologically deep cooperation, as illustrated by the multi-year agreement announced in 2026 under which Apple's foundation models will rely on Google's Gemini models and cloud infrastructure (Google-Apple joint statement, 2026). This coexistence of competition and collaboration arguably strengthens the case for treating AI giants as a homogeneous bloc for the purposes of our analysis of AI-based entrepreneurship. A systematic examination of their competitive dynamics lies beyond the scope of our investigation.

occurrence in the broader scientific landscape), but, even more peculiarly, the 'industry' in question is, to a substantial extent, composed of a small number of firms, namely the AI giants themselves or companies closely tied to them. The phenomenon is well-documented in the literature, using various proxies – from publication affiliations and their quality (see various *AI Index Reports* produced by Stanford University), to the composition of the boards of leading international conferences in the field (Rikap, 2024b) – and spans multiple fields and subfields of AI, although it is arguably most evident in the domain of natural language processing (Bouschery et al., 2023; Liang et al., 2024). In a nutshell, AI giants or former start-ups backed by them, often in strict collaboration with the academic world, are responsible for a large share of training runs, algorithmic innovation, and impactful research in AI. We can therefore reasonably assert that the potential steering activity of AI giants over the other participants in the ecosystem – and, for our purposes, over innovative start-ups – begins from the very foundations. This clearly also applies to the second area of influence: *inferential infrastructure* (the left-most position in Figure 2.1). By this term we refer to the combined set of HW (e.g. processing units), SW services (e.g. cloud deployment platforms), and data and model assets (e.g. datasets, multi-layer neural network weights and parameters, and real time inputs) that enable AI models to perform predictive or generative tasks at scale (see Chapter 2). AI giants dominate the field across all segments, starting with NVIDIA with its GPUs (alongside other large-scale AI chip producers, e.g. Google with TPUs) in processing units, and Amazon, Microsoft, and Google in cloud services, to cite just a few emblematic examples, whose positions are not expected to change in the very short term – a trend further consolidated by mega-projects such as Stargate, emblematic of the near-complete alignment between Big Tech firms and large-scale AI infrastructure development. As regards data and model assets, most leading AI firms are key actors in the development of foundation models (e.g. LLMs; see Chapter 2). Moreover, MAMAA companies, as dominant platform-based businesses, are especially well-positioned to exploit the two types of (indirect) network externalities discussed in Chapter 2, which Hagiu and Wright (2023) and (2025) identify as (i) across-user learning and (ii) within-user learning data feedback loops (Jacobides et al., 2021; Rikap, 2024b). It is thus fairly legitimate to assert that, taking this second area as a whole, innovative AI-based

start-ups can easily find themselves in a condition of both strong dependence (consider, for instance, cloud services) and clear disadvantage, particularly as regards data, vis-à-vis AI giants. Progressing along the AI value chain, a third key channel through which AI giants influence AI-based innovative entrepreneurship lies in their strong pull on AI *talent*. As discussed in the second stylized fact, this dynamic can give rise to 'adverse selection' effects. Without necessarily reaching such extreme outcomes, it is clear that the talent war waged by Big Tech companies coupled with the generous recruitment policies of many incumbents anxious not to miss the 'AI train' – whether hype-driven or not – pushes wages upward, thereby potentially constraining AI-based start-ups in their access to qualified expertise (Fossen et al., 2024). While one might be inclined to view this situation as transitory, with the labour market expected to adjust over a longer time horizon, the effects of this phase of limited AI talent availability for AI-based entrepreneurship may have much stickier, longer-lasting, and potentially irreversible consequences in terms of both the pursuit and the actual realization of innovative entrepreneurial ideas (primarily in the domain of AI-based start-ups and then, by extension and indirectly, in that of AI-served start-ups).[6] Alongside shaping the ways in which AI-based start-ups access human capital, AI giants also exert considerable influence over the financial capital available to them, both directly and indirectly. The fourth area of influence, then, manifests itself most notably within the *VC* market. Indeed, AI giants, and Big Tech companies more broadly, are (very) active (corporate) VCs (e.g. Prado and Bauer, 2022; Gao et al., 2024). Yet their influence over entrepreneurial finance is not limited to direct investments alone. In many cases, other VC investors may deliberately refrain from backing certain start-ups whose business models are perceived as too close to those of the MAMAA firms because direct competition with or early acquisition by such powerful incumbents may significantly reduce the expected value of the start-up, effectively creating what has been described in the literature as a sort of 'kill zone' (Kamepalli et al., 2020), that is, a strategic 'no-go' area for other investors. On the other hand, the opposite effect

[6] See Bessen et al. (2023) for extensive empirical evidence in entrepreneurship studies documenting how competition for talent (not strictly related to AI) between large firms and start-ups, by pushing salaries upward, may have profoundly negative effects on the growth performance and ultimate success of the latter.

is also possible, at least in principle. Proximity to the core business of a dominant tech incumbent may increase the likelihood of being acquired (early or not) by that same incumbent in the future. In this case, VC might flow precisely towards those start-ups operating within or around incumbent spheres, betting on potential 'exit through acquisition' scenarios (Prado and Bauer, 2022). Which of the two effects prevails – in general, and in the AI domain in particular – remains an open question, and probably depends on several contingencies (e.g. Prado and Bauer, 2022, Gao et al., 2024). What is certain, however, is that in either direction, besides their direct influence as investors, the presence of AI giants also exerts a strong – albeit indirect – influence on the geography and dynamics of entrepreneurial finance. Finally, the fifth and final area of influence, which is certainly not the least important, concerns the *acquisition* of start-ups. Historically, Big Tech firms have been major acquirers of innovative start-ups (e.g. Jacobides et al., 2021; Ezrachi and Stucke, 2022; Prado and Bauer, 2022; Gautier and Maitry, 2024; Ioramashvili et al., 2024).[7] Unsurprisingly, in recent years, this 'corporate shopping' attitude has increasingly, and most intensively, extended to the world of AI (e.g. Gautier and Maitry, 2024; Rikap, 2024b), with all the associated implications for social welfare in the field of innovative entrepreneurship.[8] Box 6.2 presents evidence

[7] A single figure may serve as more illustrative than many others in this regard. From its founding in early September 1998 to the end of August 2025 – a span of 27 years – Google (Alphabet) has acquired no fewer than 261 companies. Wikipedia maintains an up-to-date list of these acquisitions, with the caveat that any inaccuracy is most likely an underestimation. This corresponds, on average, to roughly one acquisition every 38 days. Other Big Tech companies have shown a lower level of acquisition activity when measured in terms of the absolute number of deals, yet their engagement is certainly far from negligible (see Ioramashvili et al., 2024, their table 1).

[8] In broad terms, two antithetical perspectives can be identified regarding the impact of Big Tech acquisitions of start-ups on social welfare. A benevolent view may interpret them as a mere mechanism of technology transfer (e.g. Cabral, 2021), while a more critical one sees them as a distortion of normal competitive dynamics (e.g. Rikap, 2024a). We do not reject the former position outright, and in several cases, acquisitions, especially when there are strong complementarities between acquirer and target, may indeed produce (social) benefits. However, our position, which will be further elaborated in Chapter 7, warns against considering acquisition as the natural evolution of entrepreneurial dynamics (and in this aspect it aligns much more with the latter), not only for reasons of economic policy, but above all for more 'philosophical' reasons related to the very nature of capitalism.

Box 6.2 MAMAA as VCs and acquirers of AI-based innovative start-ups

Starting from the 'alpha', using the already cited PitchBook database and its information on 1,476 'deep learning' start-ups founded since 2010 (as of January 2025), the share of companies funded by the MAMAA group is certainly not numerically significant (68 is the number of investees that can, on a first assessment, be directly attributed to them, i.e. 4.6 per cent). Yet their involvement becomes far more visible when considering the amount invested. Start-ups co-invested by MAMAA for at least one round (i.e. deal) raised, on average, approximately three times the capital raised by the remaining VC-backed firms (€79.2 million vs. an average of €24.9 million across a total of 1,122 start-ups). This figure is likely to represent a lower-bound estimate, given the large number of missing values – particularly on the MAMAA side, as regards the deal size amounts. In fact, looking only at the 37 deals involving MAMAA for which deal size data are available, the average amounts to about €60 million, while, as already mentioned (see Box 5.3), the average deal size in the entire database (in this case 2,853 deals) is €10.8 million.

Interestingly, estimates from a Negative Binomial Regression and a Log-linear Model reveal that the presence of MAMAA in a previous round increases the number of investors in the subsequent funding round while shortening the time needed to reach it. Specifically, *ceteris paribus*, being invested in by MAMAA in a given round increases the expected number of investors in the next round by more than 50%, while the time needed to receive a new round is slightly less than 25% lower than if such investors had not previously invested. Similar findings are obtained when the analysis is extended to include other AI giants (e.g. the three Chinese BAT, ByteDance, Huawei, NVIDIA, IBM, and Oracle).

These results both point to the strong signalling effect that an investment by these AI giants is capable of exerting in the market – an effect that may well extend beyond the amount invested – and represent further evidence of how MAMAA (and AI giants in general) can significantly influence the dynamics of entrepreneurial finance.

Turning to the 'omega', drawing on the same Crunchbase dataset of 55,087 AI-based start-ups used in the analysis presented in Box 6.1, it is possible to identify a total of 2,835 acquisitions of AI-based start-ups over the period 2005–2023, of which 83 (around 3 per cent) were carried out by a MAMAA acquirer. Again, this is certainly not a large percentage, although it should be noted that the count does not

include acquisitions made by affiliated or satellite companies (and, after all, this still represents the mere numerical acquisition activity of only five firms in a specific sector, however large and powerful they may be, so higher figures were probably never to be expected). In any case, and similarly to what has been observed above, it is interesting to note that the 'innovative weight' of these acquisitions is probably considerably higher. Based on the few available data, a start-up acquired by a MAMAA firm shows an average patent activity of almost eighty-eight patents, compared to an average of around thirty-eight across the full sample, and fifty-six when considering only the subset of acquired firms.

on the alpha and the omega – from the perspective of the financial cycle – of the MAMAA firms' engagement with AI-based start-ups, namely the role that AI giants play in shaping both the entry and the exit of nascent entrepreneurial ventures in AI.

The Provocation: From Schumpeter's 'Private King' to the 'Take the Money and Run' Entrepreneur

The relevant acquisition activity described earlier could clearly not take place if demand did not, in some way, meet supply – that is, if innovative start-ups were not themselves willing to monetize up front their entrepreneurial effort once it proves promising. Put differently, it seems to us that we are quite far nowadays from the Schumpeterian portrait of the entrepreneur as the romantic hero outlined in Chapter 1, that is, the entrepreneur as someone willing to go against the tide, driven by the 'will to conquer' new (market) territories and establish themselves as the undisputed ruler of a 'private kingdom'. In other terms, in our view, we have witnessed over time an important cultural shift, whereby a new narrative is increasingly taking hold: the successful entrepreneur is no longer the one who, by resisting the temptations of the incumbents, eventually manages to displace them, but rather the one who, by attracting their interest, eventually manages to be bought by them. To be sure, the scientific literature still lacks direct and conclusive evidence to support this cultural shift beyond doubt, and it is precisely for this reason that we present it here as a (plausible) provocation.

Box 6.3, nevertheless, offers some preliminary indications in this direction.

Box 6.3 What has 'success' in entrepreneurship become? Evidence from press and science

Using Nexis Uni, a well-known database of international business and financial news, in October 2025 I conducted a textual analysis aimed at tracing the evolution of media narratives surrounding start-up success.

Specifically, the research examined the number of articles over time in the News Section of the database that contained, within the same body of text, the expressions 'successful exit' and 'start-up(s)'. The purpose was to assess whether, over time, the dominant perception of what constitutes a 'successful' start-up has shifted – from an enterprise achieving independent growth to one whose success is defined by being acquired.

Clearly, the frequency of co-occurrence of these terms within the same article does not constitute conclusive evidence, yet it represents a strong indication of a changing public perception when observed over a long time horizon and compared on a year by year basis. Indeed, there is no apparent reason to assume that such a frequency would vary due to factors other than a shift in collective attitude towards start-up success.

Over the period from 1997 to 2025, the database returned 3,364 articles. Their distribution over time is presented in Figure 6.3a. From a single article in 1997, the number rises to 439 in 2023, 321 in 2024, and 368 in 2025 (as of October 6). Despite some fluctuations, the trend shows a decisive and sustained upward trajectory after 2010. Interestingly (and perhaps unsurprisingly, given the analyses discussed in the preceding pages and those that follow in Chapter 7), the companies most frequently mentioned across these 3,364 articles were Google (163 occurrences), Microsoft (140) and Meta (113).

The analysis was conducted using English-language terms, but replicating the test in other languages (e.g. German, French, Spanish) and corresponding media sources revealed a remarkably similar pattern, suggesting that the evolution of this narrative is not confined to the English-speaking world.

Finally, as a further robustness check of this shift in perception, in Figure 6.3b I replicated exactly the same exercise over the same time span, but this time within the scientific domain, using the Scopus database (search within 'all fields'). Out of a total of 1,955 documents, the number rose from a single publication in 1997 to 183 in 2023, 229 in 2024, and 185 in 2025 (as of October 7), thus displaying a trend that is virtually identical to the one observed in the news media.

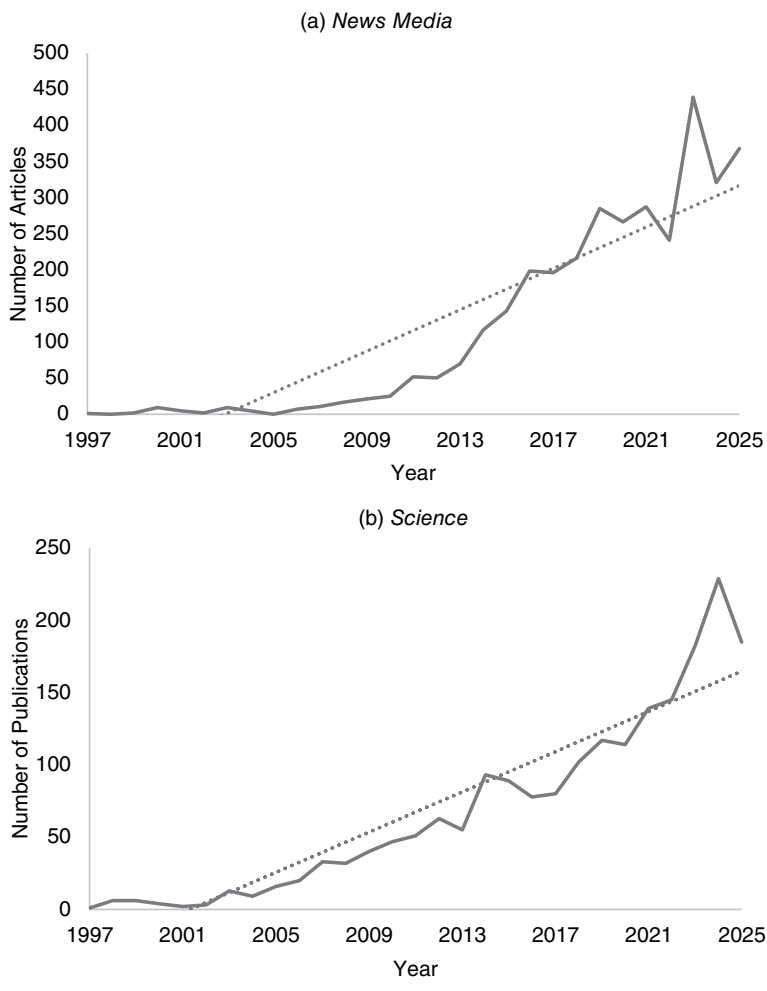


Figure 6.3 Evolution of media and science narratives on start-up success (1997–2025). Evolution of the number of (a) news articles (b) and scientific documents containing the expressions 'successful exit' and 'start-up(s)' within the same body of text. Linear trends as dashed lines. Source for (a) Nexis Uni (data retrieved on 6 October 2025; the 2025 count is therefore partial). Source for (b) Scopus (data retrieved on 7 October 2025; the 2025 count is therefore partial).

Beyond this specific evidence, what we do observe more broadly are also indirect signals – such as the changing equilibrium over time between IPOs and mergers & acquisitions (M&As), with the latter increasingly dominating as the preferred mode of entrepreneurial exit (e.g. Signori and Vismara, 2018). Regardless of whether or not the shift is as deep or widespread as it appears, one key point remains, which can be illustrated by referring to the early history of Microsoft – a story masterfully reconstructed, among others, by Rohlfs (2003). Had Bill Gates, then leading a tiny start-up, accepted IBM's early offer to license its operating system on an exclusive basis in exchange for a substantial short-term gain, IBM would in all likelihood, still be a major player in the personal computer market today, or at the very least would have endured far longer than it did. And Microsoft, quite possibly, would be little more than a footnote.[9] Whether social welfare ultimately increased as a result of Microsoft's decision – as opposed to an unknowable counterfactual – is difficult to say. But the bottom line, in our view, is the following: forms of 'intellectual monopoly' (Rikap, 2024a), or 'corporate innovation system' (Lundvall and Rikap, 2022; Rikap, 2024b), in which innovation is either born inside or brought back inside the perimeter of a corporate 'blob' – one governed by a single global orchestrator (Jacobides et al., 2021) – are, by their very nature, very unlikely to foster the emergence of radical innovations that might threaten the stability of the ecosystem itself.

The cultural shift on the supply side evoked above, if ultimately confirmed, appears clearly endogenous – that is, largely shaped by demand, and in particular by the strategic behaviour of the AI giants, which leaves entrepreneurs and VCs with few realistic alternatives

[9] Interestingly, and probably already signalling the change in attitude discussed in Box 6.3, another MAMAA firm – specifically Alphabet (then Google) – almost failed to exist in its current form. At the outset, the founders, Larry Page and Sergey Brin, wishing to continue their studies at Stanford, sought to sell their newly created start-up for one million dollars to the then market leader, AltaVista, which turned down the offer and thereby missed the opportunity of a lifetime (Cabral, 2021). Our point, as will be emphasized later, is that the substantial advantage in *inferential infrastructure* now available to AI giants renders such misjudgements, while not impossible, considerably less likely than they were at the end of the last century.

other than (early stage) acquisition. In this respect, we emphasize that, rather than stemming from purely siloed dominance in the five areas of influence mentioned earlier, entrepreneurial lock-in by AI giants may also be reinforced through what we define as a 'Systemic Lock-in Model', whose dynamics may lead to a possible entrapment of the AI entrepreneur, the theme of Chapter 7.

PART III

Policy Challenges

7 *The Perils for the AI Entrepreneur*

So What?

The analytical framework constructed up to this point may be synthesized as follows. Although AI – conceived as a mere digital 'statistical machine' – undeniably has the potential to transform the entrepreneurial process, and consequently the very identities of those who enact it, from the conception of the entrepreneurial idea to its execution, it should not fundamentally alter the role of the entrepreneur within a capitalist society. Put differently, technology in itself does not seem to possess characteristics capable of undermining entrepreneurship as a constitutive actor of capitalism. Indeed, from a purely technological perspective, quite the opposite may be argued: the modelling capacities of AI in both *predictive* and, above all, *generative* data processing can serve to enhance entrepreneurial creativity. AI does not eliminate ontological uncertainty – which only the entrepreneur can ultimately confront – but a range of 'AI entrepreneurial buddies' may act as valuable complements in the generation and implementation of ideas and thus in entrepreneurial action, particularly when supported by *agentic* AI. Moreover, as we have seen, the three fundamental 'capital' pillars of success in innovative entrepreneurship are likely to undergo changes in their internal composition and in their relative weight vis-à-vis one another. Yet, taken as a whole, there is little reason to assume that they would cease to constitute the very 'pillars' of entrepreneurship, even in the face of the AI revolution. And also in this case, AI could bring about highly positive effects (consider, for example, the potential greater ease of access to VC even for entrepreneurs without the 'right' connections). More or less the same type of argument also holds for psychological components, whose weight, although moderated by new factors (i.e. trust in AI tacitness), we believe remains important.

Thus, recalling Bowles's three ingredients for capitalism to 'work well' (i.e. 'good skill, hard work, and good luck'; Bowles, 2013, p. 24), there is no reason to think that AI – much like other recent technological revolutions (e.g. the Internet) – could, by itself, if we pause for a moment to consider only its technological features, overturn the recipe for 'good' capitalism (Baumol et al., 2007). Accordingly, highly successful innovative start-ups – whether AI-based or AI-served – have been created in the recent past, continue to be created, and will almost certainly be created in the future (obviously with varying intensity across the four quadrants of our taxonomy in Figure 6.1).

Thus, what we have delineated so far is, overall, a picture of change, yet one that displays far more light than shadow in the relations between AI, entrepreneurship, and capitalism. Some shadows, of course, we have encountered. First and foremost, what we have defined as the *incumbent advantage* in the availability of data, primarily affecting more Kirznerian forms of entrepreneurship. But also possible phenomena of adverse selection in AI-based innovative entrepreneurship, and the omnipresence of the AI giants. This chapter is therefore devoted to delving more deeply into these (and other) potential perils for the AI entrepreneur. Fully aware that, as is often the case with technological developments, should issues emerge, they are less likely to stem from the technology itself than from the economic processes it triggers and the institutional arrangements within which it becomes embedded (see in particular the First, Fourth, and Sixth 'Laws' of the already cited Kranzberg, 1986).

We will therefore begin this investigation by building on the five areas of influence of the AI giants identified in the previous chapter (*science*, *inferential infrastructure*, *talent*, *VC*, and *acquisitions*) to better delineate a second threat, which we define as the *systemic lock-in* imposed by the AI giants – a risk that also bears upon Schumpeterian entrepreneurship. We will then continue with what appear to us as other significant risks, pertaining to the macroeconomic and sociological spheres, and which may in some way weaken the 'innovative reach' of AI entrepreneurs. This analysis is not necessarily meant to contradict the positive aura we believe AI may exert on the dynamics of innovative entrepreneurship. Rather, its aim is instrumental in

identifying areas of possible intervention by policy makers, so as to make prospects even brighter in terms of future social welfare, which will be the topic of our last chapter.

Towards a 'Systemic Lock-in Model' of AI-Based Entrepreneurship

The 'Systemic Lock-in Model' of AI giants, by exploiting complementarities across the five domains and generating super-additive effects, may significantly limit the competitive threat posed by innovative AI-based start-ups (and also that posed by AI-served start-ups, in light of the platform pervasiveness of a corporate 'blob' operating across all four quadrants of our proposed taxonomy in Figure 6.1 and, as a result, across an expanding range of sectors and markets). A graphical representation of the model, together with selected examples of these cross-domain synergies, is provided in Figure 7.1.

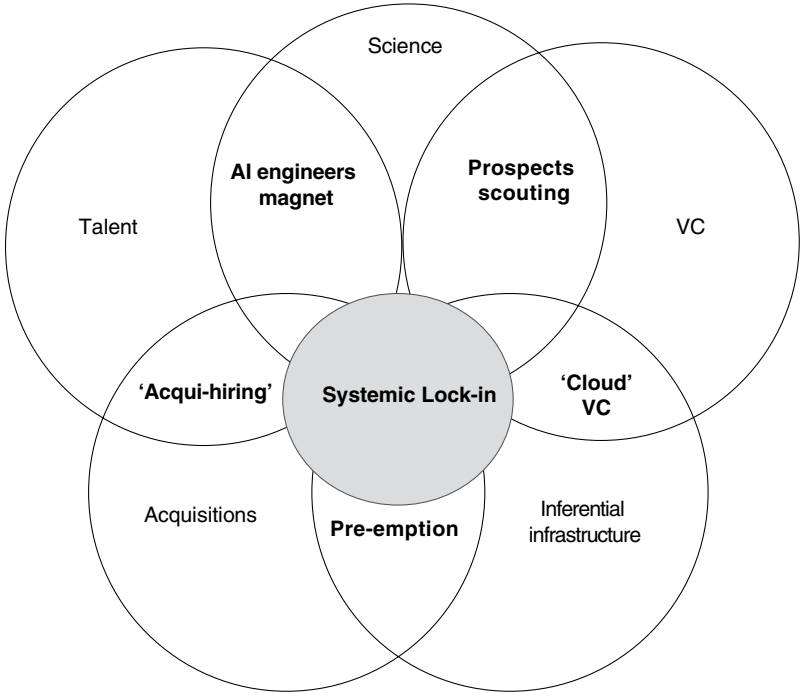


Figure 7.1 The 'Systemic Lock-in Model' of AI-based entrepreneurship

We illustrate here only some actual and potential complementarities across the five areas in which AI giants operate, but by no means do we intend these examples as exhaustive of the full range of synergies that may arise.

Acquisitions and Talent: 'Acqui-hiring'

The first important super-additive effects manifest themselves in the simultaneous presence of AI giants in both acquisitions and talent sourcing. The intersection of these two domains gives rise to what has become widely known as *acqui-hiring*: acquisitions conducted less for the intrinsic value of a product or a technology of an innovative start-up than for the purpose of internalizing highly skilled human capital. By combining their financial power with their already strong position as global magnets for AI talent (see next point), AI giants are able to reinforce their dominance along two trajectories at once. On the one hand, they weaken the pool of human capital potentially available to independent AI-based start-ups, depriving them of some of the very skills needed to scale and compete. On the other hand, they further enlarge their own in-house research and engineering capacity, thereby strengthening the complementarities with other areas of their ecosystemic power.

Talent and Science: AI Engineers Magnet

Another area where complementarities generate super-additive effects within the 'Systemic Lock-in Model' is certainly the intersection between talent and science. We have already noted how, also thanks to extremely high salaries, the AI giants are able to attract leading scientists; but this is possible above all because these AI giants allow scientists and AI engineers to continue being precisely that – and at the highest levels. As we have already emphasized, the fact that the AI giants' laboratories and research units publish in the most prestigious venues, frequently in collaboration with leading universities, offers engineers and researchers the opportunity to continue building a recognized scientific profile while working in industry. Moreover, beyond traditional academic publishing, several AI giants also cultivate legitimacy in the open-science and open-source communities, a theme often particularly cherished by the brightest

minds in the field.[1] Initiatives of this kind not only showcase scientific excellence but also typically create reputational spillovers, reinforcing the idea that work conducted within these companies continues to circulate in the broader research commons (e.g. Jacobides et al., 2021). This dual opportunity, that is, career advancement within a corporate setting combined with academic and open-science visibility, represents a powerful *magnet for AI engineers*. And again it reinforces the concentration of talent within the AI giants and further depletes the pool available to independent start-ups. At the same time, the continuous inflow of highly skilled researchers strengthens the scientific output of the giants themselves, thus perpetuating the cycle: talent fuels science, science attracts further talent, and the lock-in deepens.

Science and VC: Prospect Scouting

Being positioned at the absolute frontier of AI science, indeed, actively contributing to pushing that frontier forward, gives the AI giants a clear advantage in identifying which entrepreneurial prospects and ideas are most deserving of attention. Their visibility into emerging research domains allows them to act as early identifiers of potential breakthroughs, and their presence as early stage investors, in turn, feeds back positively into shaping which scientific avenues should receive the greatest reinforcement. Thus, this dual positioning may easily create a privileged *scouting* mechanism, placing them in a stronger position than other investors to recognize promising entrepreneurial prospects and, ultimately, to channel capital accordingly. Importantly, such an arrangement may also generate 'kill zones' or, if the opposite effect prevails, risk narrowing the diversity of the entrepreneurial landscape, where start-ups whose research trajectories align with the priorities of the giants are more likely to secure monetization, either through initial backing or later acquisition.

[1] As we know from Chapter 2, Google (Alphabet), for instance, has been a leading contributor to TensorFlow, Meta has released PyTorch and the LLaMA family of LLMs. Microsoft not only supports open-source frameworks but also acquired GitHub in 2018, the world's largest platform for collaborative software development. Over time, GitHub has become the default repository for open-source AI tools, libraries, and even model checkpoints, thereby allowing Microsoft to sit at the very heart of the global open-source ecosystem. Yet these are only a few among the many possible examples.

VC and Inferential Infrastructure: 'Cloud' VC

The simultaneous presence of AI giants in both the VC segment and inferential infrastructure can also generate substantial super-additive advantages. As Rikap (2024b) has underlined, the cloud credits offered by AI giants (initially Amazon, later emulated by Microsoft and Google) to innovative start-ups may operate as a useful form of (relatively) low-cost seed VC (which we simplify here as *cloud VC*). This practice has a double benefit (see also Srnicek, 2025): on the one hand, it can be seamlessly scaled up through more substantial investments once the start-up proves to be promising; on the other hand, almost by definition, it inevitably ties the innovative firm to the infrastructure of the AI giant providing the credits, producing what interviewees in Rikap (2024b, p. 1750) describe as 'high stickiness, in other words, a lock-in effect'.

Inferential Infrastructure and Acquisitions: Pre-emption

Indeed, possession of the inferential infrastructure confers a great structural observability advantage: cloud telemetry, API logs, marketplace analytics, and billing signals reveal which start-ups are scaling fastest, which models or datasets are resonating, and where compute intensity is spiking. This signal privilege reduces search and evaluation costs, sharpens timing, and may improve deal selection relative to outside investors. Armed with this foresight, AI giants can also clearly pursue *pre-emptive acquisitions* (e.g. Cabral, 2021) – neutralizing nascent threats before they mature – while simultaneously folding distinctive datasets, model artefacts, and pipelines into their own stack. Whether they take the form of *killer acquisitions* (e.g. Cunningham et al., 2021) in which the technology or product of the target firm is discontinued, or of *reverse killer acquisitions* (e.g. Padilla, 2024) in which the opposite occurs – namely, the technology or product of the target replaces that of the acquirer, which nevertheless appropriates the benefits within its own ecosystem – both dynamics illustrate the strategic use of acquisitions made by AI giants in shaping innovation trajectories. While, in theoretical terms, some scholars remain sceptical about the occurrence of such pre-emption phenomena in digital markets, given their high degree of uncertainty (e.g. Cabral, 2021), others are more inclined to consider them possible despite such uncertainty, as noted by Shapiro (2018):

> Acquisitions like these [those made by MAMAA, author's note] can lessen future competition, even if they have no such immediate impact. To illustrate, suppose that the target firm has no explicit or immediate plans to challenge the incumbent firm on its home turf, but is one of several firms that is best placed to do so in the next several years by developing innovative new products or by improving or modifying its existing products. Not even the target firm knows for sure how its product offerings will evolve. Does it seem so far-fetched that the dominant incumbent firm, whose market capitalization will fall sharply if successful entry occurs, would pay a premium to acquire the target firm in order to avoid the risk of facing this pesky rival in a few years' time? Not to me. Nor does it seem far-fetched that a dominant incumbent firm can reliably identify the firms that are genuine future threats […]. (Shapiro, 2018, p. 740)

Looking at his latter point, the capacity to recognize such menaces is, of course, reinforced by the superior position of AI giants vis-à-vis all other actors, precisely in terms of the *inferential infrastructure* we are discussing here – and one that Shapiro most likely did not yet have in mind. In this regard, the extant empirical evidence on *killer acquisitions* by Big Tech firms in digital markets is actually mixed. While some contributions, such as Gautier and Maitry (2024), lend support to the idea of a 'killing instinct' among MAMAA firms, others arrive at a far more cautious, if not opposite stance (e.g. Ederer et al., 2025). Our descriptive analysis in Box 7.1, however, seems to suggest that in the specific domain of AI, Shapiro's concerns may be particularly warranted.

Thus, once again, the ultimate result of this last component of the model might well be a reinforcing loop: infrastructure → superior signals → earlier, better-aimed acquisitions → deeper infrastructure and data advantages → tighter lock-in for the remaining innovative start-ups. The aggravating factor is that the objective advantage created for the giant AI by such a reinforcing loop (as well as by the others highlighted) increases its incentive to pursue this path over time: the greater its market power becomes, the more it has to lose from any potential 'systemic' threat, and the more it is therefore interested in eliminating such threats by one means or another.

Outcome

The overall effect is that the application of the 'Systemic Lock-in Model' by AI giants further consolidates and entrenches their dominant market

Box 7.1 AI 'killer' acquisitions

Classifying an acquisition as a 'killer acquisition' is never a straightforward task, since the discontinuation of a product or entrepreneurial project may result from a wide range of contingencies, many of which are difficult to disentangle and not necessarily driven by malign intent (Gautier and Maitry, 2024). Consequently, any proxy selected to identify such a phenomenon inevitably has its strengths and weaknesses, and is therefore open to criticism and should in any case be applied with caution. Nevertheless, it seems to us that a particularly strong signal of a potential killer acquisition may be represented by the closure of the target firm following its acquisition, as this is arguably the most 'neutral' among the possible proxies.

Drawing on the same Crunchbase dataset of 55,087 AI-based start-ups used in the analysis presented in Box 6.1, we have seen (in Box 6.2) that acquisitions carried out by MAMAA amount to 83 out of a total of 2,835 acquisitions of AI-based start-ups that took place during the period 2005–2023 (not including satellite companies). Defining killer acquisitions as those cases where the acquired firm subsequently closed, we can attribute 816 AI-based start-ups to this category. Within this subset, MAMAA firms account for 48 transactions (around 6 per cent). These figures (which must, of course, be interpreted with all due caution) suggest that 48 of the 83 acquisitions undertaken by MAMAA between 2005 and 2023 – approximately 58 per cent – were very likely killer acquisitions, whereas the incidence of killer acquisitions in the total sample of acquisitions is significantly lower (816 out of 2,835, at 29 per cent). More than one third of these suspected killer acquisitions by MAMAA (18 out of 48, approximately 38 per cent) were announced during the final five years of the observation period (2019–2023).

Extending the scope to AI giants (beyond MAMAA, this set may be taken to include IBM, NVIDIA, and Oracle together with the major Chinese players – BAT, that is, Baidu, Alibaba, and Tencent – plus ByteDance and Huawei, while excluding activities of satellite companies), the proportion highlighted above remains similar (55 per cent), with this broader group accounting for 61 killer acquisitions (as defined here) out of a total of 111 acquisitions. With regard to timing, the share of suspected killer acquisitions announced in the most recent period is again approximately 38 per cent (23 out of 61).

Thus, all in all, and even allowing for all the necessary caveats set out in the incipit of this descriptive investigation – and therefore also conceding the possibility that these metrics may overestimate the true extent of the phenomenon – it seems to us that these figures, in any case, remain a cause for serious concern.

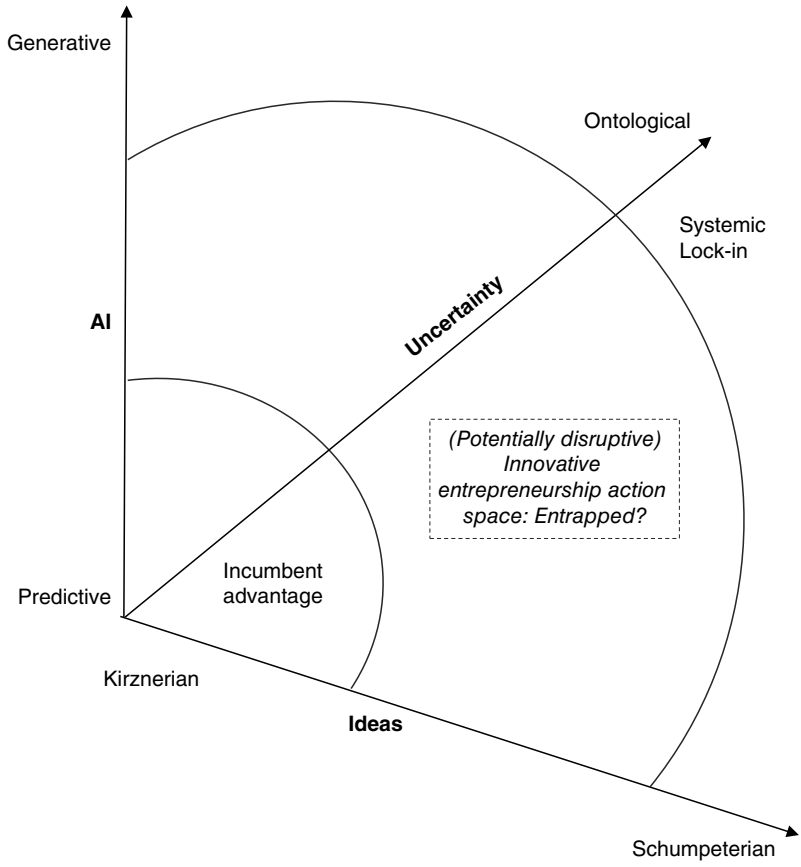


Figure 7.2 The AI entrepreneur and the risk of entrapment

positions, to an extent that surpasses the reinforcing effects typically associated with network externalities. While network markets typically evolve towards *winner-takes-all* (or *most*) outcomes, dominant positions should nevertheless remain subject to potential disruption through technological breakthroughs, in line with a Schumpeterian form of 'competition for the market', which enhances social welfare via the channel of dynamic efficiency. The operation of the model described here, however, clearly *undermines from within* the likelihood of such processes taking place – particularly by constraining the agency of innovative start-ups engaging in *creative destruction* in an SM I fashion.

Figure 7.2 therefore completes our framework by highlighting the risk that a non-negligible portion of the '(potentially disruptive) innovative entrepreneurship action space' may become trapped, if not by

the former barrier we have identified (the *incumbent advantage*), then at least by the latter (the *systemic lock-in*).

Yet, as said, these are not the only risks that can be discerned on the horizon for the AI entrepreneur.

Macroeconomic and Political Economy Constraints

States' 'AI Wars' and the Possible (Further) Tendency to Consolidate National Champions for AI Sovereignty

The first additional risk to innovative entrepreneurship that we highlight is – if the expression may be allowed – an 'old' one, corresponding to a classic dilemma in industrial policy, yet re-emerging in the domain of AI in markedly amplified form. First, in contemporary capitalism (or at least in many of its variants), the boundaries (and indeed the objective functions) between the State and the large firms appear increasingly blurred (Lundvall and Rikap, 2022; Srnicek, 2025). Second, one must keep in view the nature of AI as a general-purpose technology within an ever more datafied digital economy, and hence its pervasive reach across production processes, not only today but, even more, in prospect. Third, there is the characteristic structure of digital markets, marked by economies of scale, network externalities, and the ensuing *winner-takes-all* (or *most*) dynamics, of which AI is no exception.

The combined effect of these three elements may easily lead the global superpowers (first and foremost the United States and China, and secondarily the European Union, albeit from a comparatively disadvantaged position) not merely to refrain from weakening their own major incumbents so as to bolster somehow domestic markets, but to do everything possible to strengthen them – including substantial public investment and targeted resource allocation – in order to prevail in the global 'AI war'.[2] In short, a policy of fostering *AI (supra-)national champions*, with the purpose, if not of achieving technological sovereignty, at least of securing some bargaining advantage in relations

[2] Accordingly, the emphasis, both in the scientific community and in the popular press, is often not so much on firms as on States, with firms, in this view, being merely their expression – or, to continue the military metaphor, their armed wing. Consider the emblematic titles: *Why the United States Must Win the Artificial Intelligence (AI) Race* (Hull et al., 2022) and *Who is Winning in AI – China or America?* (The Economist, 2025b).

with other States. Indeed, Srnicek's (2025) historical-political account underscores that the relationship between States and AI giants has evolved through different phases, oscillating between visions of global competition, periods of intensified antitrust scrutiny or enforcement,[3] and ambitions of national self-sufficiency along the AI vertical technological stacks. The reality is therefore complex, the next developments are difficult to anticipate, and the outcome of these dynamics is far from certain. Yet, *de minimis*, it is hard to imagine a near-term future in which the AI giants will not remain – and likely become even more – entrenched as pivotal actors in the global political economy.

Seniority-Biased Technological Change and the Entrepreneurial Pipeline

A second additional risk – still insufficiently evidenced on a large scale to draw general conclusions, yet in principle plausible and already documented in its early stages – concerns the possible emergence of what Hosseini and Lichtinger (2025) describe as a *seniority-biased technological change* and what Garicano (2025) terms the 'AI-Becker problem' (see also evidence by Klein Teeselink, 2025). That is, the risk that AI-driven solutions adopted and implemented by incumbents might substitute not so much for senior workers, but rather for junior employees and mid-skilled profiles. If confirmed, this dynamic would pose a problem not only for the aforementioned companies in the long run, but would also clearly represent a serious issue within our own domain. As already noted, in line with the KSTE, the firm often plays an 'incubator' role for entrepreneurial talent. In other words, the innovative entrepreneur frequently stems from a process of (specific) human-capital accumulation that takes place within the incumbent firm (see, once again, the examples illustrated in Chapter 1). If this

[3] It suffices here to recall the accommodative stance towards the tech giants in the period preceding China's regulatory crackdown of 2021; or the more recent evolution of US antitrust, which, from positions that were less interventionist, witnessed the rise of the New Brandeis Movement. By contrast, the European approach has been easier to sustain – and thus more consistent – since, lacking (at least in relative terms compared with the other two superpowers) major 'champions', it has long allowed itself a certain continuity in constraining (non-European) Big Tech's power, culminating most recently in the Digital Markets Act.

process of growth is missing for younger cohorts, the long-term pipeline of potential successful entrepreneurs could be significantly weakened. This concern is reinforced by the general evidence that a younger age (or at least not too advanced in years, e.g. Azoulay et al., 2020), especially if combined with some 'business acumen', is conducive to innovative entrepreneurship (Grilli, 2022).

The plausibility of this risk also becomes more apparent when we focus on the functioning of digital labour markets. If AI adoption leads to the erosion of entry-level and mid-skilled positions inside firms, the same dynamics of displacement and downgrading can be observed in external labour markets, where certain AI-related tasks are increasingly standardized and, for this reason, are demanded as 'gig work' by incumbents. Box 7.2, which examines AI freelancers on Upwork, documents this phenomenon, showing a global supply (and hence market) where certain AI technical skills have been commoditized. The increasing relevance of 'gig work' in modern advanced economies is what Bowles and Carlin (2021) identify as a defining trait of a new form of capitalism, which they term 'shrinking capitalism'.[4]

Box 7.2 AI freelancers in a crowdwork platform: the case of Upwork

The characteristic of an increasingly global collaborative nexus, where the demand for a particular job expressed in one geographical context can be met by labour supply from a completely different part of the world, is perhaps best exemplified by crowdwork platforms. The COVID-19 pandemic has certainly stimulated the shift to remote work and freelancing, and one of the most well-known platforms in this field at an international level is certainly Upwork.

Launched in 2015 from the merger of Elance and oDesk, Upwork is a global freelancing platform that facilitates collaboration between

[4] We are, however, not fully in agreement with this expression. In our view, it is not capitalism, in the strict sense, that shrinks – indeed, on closer inspection, market-governed economic relations expand enormously – but rather the boundaries of the modern firm itself that shrink. This (in some respects slight) disagreement stems from the fact that the authors consider the defining trait of capitalism to have its basis in *firms* (with large numbers of employees governed by owners or managers), whereas the definition we adopt places greater emphasis on *market relations* as the essential feature of capitalism.

businesses and independent professionals across different geographic regions. With the latest data indicating millions of freelancers and clients (i.e. firms) across more than 180 countries, the platform enables businesses to access a diverse pool of skills while providing professionals with opportunities to engage in international projects across various fields, including AI services.

Our independent analysis (conducted in November 2023 on the platform data) confirms that, even in the AI services sector, the AI labour market appears rather global. AI freelancers – who can to some extent be considered self-employed entrepreneurs – come from all six continents. Specifically, most of AI freelancers are based in Asia (61.4%) and in particular, in India (31.4%). The old continent counts for 17.1%, while the US for 12.7% of the AI freelance global labour supply. Using Natural Language Processing (NLP) techniques through CorText Manager (Breucker et al., 2016), a semantic analysis on their job titles' posts reveals that the most common keyword used by freelancers to propose themselves is 'Full Stack Developer'. Generally, this label means a professional skilled in developing both front-end and back-end systems, proficient in multiple programming languages, and capable of integrating different software components. The term is often paired with ML/AI-related terms, though it does not necessarily imply medium- or high-level expertise or the capacity to design AI architectures from scratch. Keywords from more specific technical domains (e.g. Python language programming) or 'soft' AI-driven content creation tasks (e.g. prompt engineering and/or content writing) appear with much less frequency. Looking at the demand (i.e. client) side of the 'AI gig work' market, a very basic hedonic wage regression model run on 2,173 posts, suggests that an AI job involving 'Full Stack Developer' skills does not significantly increase (neither in economic nor in statistical terms) the minimum wage offered by companies compared to other typologies of jobs and tasks (i.e. design, marketing and management).

All in all, in addition to the 'global' and 'remote' nature of the market, what this last result seems to suggest is a large-scale commoditization – or at least a growing standardization – of certain technical AI-related skills, that, until recently, were not exactly trivial. These insights also align well with the findings of an OECD report using Lightcast online vacancy data (Borgonovi et al., 2023), which analyses AI skill demand across multiple countries and industries. The report reveals how some technical AI competencies are increasingly concentrated in a narrow set of roles and are fairly standardized across different employers, which may suggest a growing routinization of certain types of AI-related work.

In a 'gig economy', there may well be a tendency to reduce (even high-profile, not necessarily AI-based) competencies into modular tasks that are increasingly fragmented, standardized, and routinized, with downwards pressure on remuneration and recognition of skills. Of course, the performance of these (potentially repetitive) tasks, always carried out outside the (ever-changing) commissioning firm, may well limit opportunities for younger professionals to accumulate the differentiated human and social capital within organizations that so often seeds innovative entrepreneurship.

Sociological Mediation of Entrepreneurial Ideas and Identities

Algorithmic Polarization and Homogenization in Digital Service Economies

As noted in Chapter 2, in increasingly service-based digital economies, where social platforms occupy a significant role, algorithms can more readily than in the past produce phenomena of *polarization* and *homogenization* (Nechushtai et al., 2024), dynamics that, in certain respects, resonate with populism (Guriev and Papaioannou, 2022). This can clearly exert an impact on entrepreneurial dynamics in several nuanced ways. On the one hand, establishing an initial successful idea may become easier and quicker within a specific niche, with less need for adaptation to different cultures and habits inside that niche. On the other hand, entrepreneurs may find it harder to introduce novel products or services that set global dominant trends – that is, that do not belong to predetermined communities of interest or that manage to transcend their boundaries. The result, at the level of society as a whole, may well be the consolidation of *informal institutions* that, more readily than in the past, incentivize a form of 'small-bore' and 'short-haul' entrepreneurship – less ambitious in scope – and, ultimately, dampen the 'think big' attitude typical of successful innovative entrepreneurs (Elia and Grilli, 2026), those who become successful precisely because they start with the intention of 'changing the world'. As Dyer et al. (2008) put it (emphasis mine):

> Meg Whitman, former CEO of eBay, has worked directly with a number of innovative entrepreneurs in our sample, including Pierre Omidyar (eBay founder), Niklas Zennstrom and Janus Friis (founders of Skype and Kazaa), and Peter Thiel and Elon Musk (founders of PayPal). When asked how these

folks differ from typical executives Whitman said 'My experience is that they get a kick out of screwing up the status quo. They can't bear it. So they spend a tremendous amount of time thinking about *how to change the world*'. (Dyer et al., 2008, p. 323)

Thus, social-platform and marketplace algorithms may indeed incentivize the targeting of specific groups of consumers/citizens, further exacerbating these silos over time and increasingly characterizing them in a feedback loop (which we consider negative) that is strongly path-dependent, and therefore increasingly difficult to reverse as time goes by.[5] Such homogenizing pressures on already polarized ideas and audiences can set the stage for a brand new risk for entrepreneurship in times of AI, namely that of a revised Veblenian entrepreneurial stance. The innovation embodied by this entrepreneur would therefore risk being not only more likely 'small-scale' in its very premises, but also increasingly *thin* in its pace.

Veblenian Entrepreneurship and (Socially) Thin Innovation

This leads us to another significant threat. Up to this point, we have spoken of Kirznerian and Schumpeterian entrepreneurship. The former is probably incremental, rooted in the recognition of entrepreneurial ideas that emerge from market imperfections, inefficiencies, or simple disequilibria. The latter is more disruptive and radical, distinguished by its capacity to create entirely new markets or to unseat incumbents in existing ones. Both, of course, form the backbone, if the expression may be allowed, of innovative entrepreneurship. When Veblen wrote in 1899 the famous *The Theory of the Leisure Class*, however, he discerned in the society of his time a third type of entrepreneurship (which we may, since then, define as Veblenian): that of the entrepreneur as an 'end in itself'. The aim is not to 'change the

[5] Empirical evidence (especially in the US, which, as we know, is also one of the most polarized contexts in the recent past) testifies to a certain (relatively) new difficulty on the part of young firms in 'scaling', that is, in reaching high growth and productivity (see Hartmann et al., 2022, for a discussion). Although this phenomenon certainly does not originate with the AI revolution but rather predates it, having deeper roots and probably more than one contributing cause (Decker et al., 2016; Sterk et al., 2021), it is clear that, for the reasons evoked here, it is difficult to expect a reversal of the trend. If anything, the opposite seems more plausible to us.

world' through innovation, nor, more prosaically, to become the sovereign of a private domain, but rather to appear to do so, or to pretend to be(come) so. In this conception, entrepreneurship ceases to be *a means* and becomes *an end*.[6] The entrepreneur becomes a client of the so-called entrepreneurship industry, whose flagship product is precisely the accreditation of individuals into the ranks of entrepreneurs in the eyes of society. Citing Hallonsten (2023):

> [E]ntrepreneurship has therefore become an arena of itself for Veblenian 'conspicuous consumption' (Hartmann et al., 2022). By consuming the products of the entrepreneurship industry, individuals can simply consume their way to the entrepreneurial identity without making any substantial contribution of the kind typically identified with entrepreneurial acts. This kind of 'Veblenian entrepreneurship' is an activity driven most of all by a want of identity, which means that the main activities of a 'Veblenian entrepreneur' is to build and maintain a lifestyle with certain specific attributes. (Hallonsten, 2023, p. 52)

Many recent contributions (e.g. Hartmann et al., 2022; Hallonsten, 2023; Frost et al., 2025), explicitly or implicitly drawing on Veblen, have pointed to the growing weight of this type of 'status' entrepreneurship in contemporary advanced economies, which – paraphrasing Baumol (1990) – we may, at best, characterize as a form of 'unproductive entrepreneurship'. The innovation embodied by this type of 'innovative' entrepreneur (and here quotation marks have never been more appropriate) has been described in various ways; the one we probably prefer is *empty innovation* (Hallonsten, 2023) which refers to the pretentious appropriation of the 'innovation label', aimed at legitimizing firms' activities that, upon closer inspection, are not innovative at all. Building on this term, we propose a new one, in order to give it a slightly different connotation. We define (socially) *thin innovation* as the generation of new products, services, or processes that, while formally novel, are 'thin in substance'. Such innovations, though some may prove commercially successful, are by their very nature substantively devoid of offering meaningful contributions to human well-being or social progress. In contrast to 'empty

[6] Veblen's thought is of course more complex, and it is simplified here (hopefully without undue trivialization) and, to some extent, bent to our purposes – yet in line with a recent literature that, in fact, offers an interpretation similar to the one adopted here (Hartmann et al., 2022; Hallonsten, 2023). For more detailed and thorough accounts of Veblen's work, see, for example, Camic (2020).

innovation' (Hallonsten, 2023), which primarily designates the rhetorical inflation of the innovation label to rebrand conventional activities as novel, thin innovation may denote actual instances of innovation, but whose novelty is directed towards 'small problems' rather than 'grand challenges', to non-essential and often induced rather than pre-existing needs.[7] The risk is that when 'innovationism' (the obsession with innovation) pervades a society – as is often the case in modern advanced economies (Hallonsten, 2023) – *thin innovations* may absorb a significant share of the natural entrepreneurial spirit available within it (and here we are close to Baumol, 1990), with the effect that a non-negligible portion of innovative entrepreneurship is, from the outset, incapable of aiming at durable improvements in welfare, productivity, or collective problem-solving.[8] Considering the opportunity costs of resource use and relevant counterfactuals, these entrepreneurial forms may plausibly evolve into 'destructive entrepreneurship' (Baumol, 1990), as they ultimately risk eroding social welfare in the long run.

[7] We are perfectly aware that the concept here introduced of *thin innovation* is, by construction, not endowed with strong ontological objectivity, that is, its contours are not intersubjectively fixed but remain, to a degree, perspectival. What one observer may deem thin another may not. Yet this does not preclude convergence at the macro level on a diagnosis of a broader *thinness of innovation*. In this respect, a useful parallel lies with positional goods: their valuation, too, contains an inherently subjective component – status and distinction depend on socially mediated perceptions – yet this did not prevent Hirsch (1976) (without even providing a precise definition; see also Tyssedal, 2026 on this point) from indicating their prevalence as a defining trait of advanced economies of the time. In the same spirit, while assessments of *thin innovation* are contestable case by case, in the aggregate such innovations may amount to a systemic misallocation of creative and productive capacities.

[8] *Thin innovation* is clearly related to other expressions that can sometimes be used in similar contexts, but do not convey the same meaning we intend to express. For example, 'ephemeral innovation' is often adopted to address transient, post-material needs or short-lived consumer desires, yet can still generate a non-negligible social value – albeit a temporary one. Indeed, our concept of *thin innovative entrepreneurship*, as detailed in these pages, should not be conflated with that of *thin innovation systems*, albeit there are points of tangency between the two. The latter refers to regional or sectoral innovation environments characterized by a lack of density in actors, networks, and institutions, which makes them structurally weak and less capable of sustaining (real) innovation. Similarly, with thin innovative entrepreneurship, we refer not to the absence of an ecosystem, but rather to the orientation of entrepreneurial efforts towards narrow, low-impact, or induced needs, which results in innovations that are formally novel yet substantively thin in their contributions to welfare and social progress.

More on (Socially) Thin 'Innovative' Entrepreneurship: Why Might AI Heighten Its Risk?

The answer to this question depends on a series of factors at multiple levels, all linked to the common thread of AI, some of which we have already to some extent unravelled. Taken together, they represent – hence the risk – that perfect alignment which may give rise to a kind of 'perfect storm' favouring *thin innovative entrepreneurship*. Several converging forces may contribute to this scenario, spanning from macro-level economic shifts down to the behaviour of individual entrepreneurs and the dominance of AI giants. These will now be considered one by one.

Macro-Level Pressures: AI and Technological Unemployment

A central factor is certainly the impact of AI on labour markets. We have already argued that predicting which jobs will emerge in the future is a daunting, if not impossible, task. Yet this is different from asserting that AI will indeed trigger major disruption in the world of work. On this point there is unanimous consensus (e.g. Felten et al., 2021, 2023; Anelli et al., 2024). In fact, automation (driven by AI) will in the short run displace large numbers of workers, while new jobs will probably be created at a slower pace (e.g. Acemoglu and Restrepo, 2019). Accordingly, the grey literature (e.g. the World Economic Forum and its *Future of Jobs Report*) projects at least in the short run a net negative outlook: millions of jobs eliminated, with fewer created quickly enough to compensate. The history of technology, for its part, shows a long-term tendency towards labour-saving in favour of capital input, and consequently a reduction rather than an increase in working hours. This is not a mechanical outcome, but the historical record nonetheless suggests that the negative balance, though perhaps reduced in absolute terms, is hardly destined to vanish. In short, AI is poised both to destroy and to create jobs, but the transition will be turbulent and is likely to involve substantial *technological unemployment* in the interim – a concept that goes back to John Maynard Keynes (1930), when he famously (though in retrospect overstated in magnitude) expected that advances in productivity would make it possible to work only a handful of hours each week.

This wave of displacement matters for entrepreneurship because workers who lose jobs because of AI may be pushed into starting their own small ventures out of necessity. And if there is one well-established finding in the literature, it is that such 'necessity entrepreneurs' typically start small-scale, survival-focused ventures rather than innovative, high-growth firms (see O'Donnell et al., 2024). AI's rapid deployment thus risks producing an influx of new entrepreneurs with scarce chances of success where the innovative activity, if present at all, is likely to be *thin*. Even for those still employed, AI will reshape existing occupations, fostering instability and insecurity, and many may be tempted to hedge by starting side businesses. Yet again, many of these ventures risk being marginal, driven more by income replacement or supplementation than by genuine innovation.

The Shift to a Digital, Post-Essential Economy

A second macro-level element to consider stems from the typical tendency of advanced economies to move towards what we might call 'discretionary production and consumption systems': a growing focus on goods and services not strictly necessary for basic living needs (such as food staples, basic clothing, housing, or utilities). In other words, as societies evolve into high-income, their consumption patterns naturally become more heterogeneous (Saviotti, 2002; Chai et al., 2023), often concentrated in niche markets (Neiman and Vavra, 2023), and, consistent with Engel's (1856) law, less oriented towards necessities and more centred on discretionary categories. Where these niche markets increasingly resemble sealed compartments, for the AI-related reasons already noted, and consumers display ever more varied preferences towards discretionary goods, there is the risk that the army of prospective necessity entrepreneurs mentioned earlier will find easy to target them. Their ventures may be innovative in a narrow sense, but it is certainly innovation of a modest scale, often lacking real substance or lasting value – in short, arguably *thin*. Exactly like our example of the 'SoleScribe Snack Shoe' in its first original version (Boxes 3.2 and 4.1), designed to illustrate how easy entrepreneurship can become with the help of AI. Thanks to the enhanced generative (and agentic) capabilities of AI, and its combinatory and convergent thinking power, the entrepreneur, increasingly risks, over time, becoming *postmodern*. In this

sense, the product offered, exactly as in our example of the environmentally friendly snack-shoe,[9] becomes the outcome of a (micro-) narrative – or rather, it is itself 'narrative', that is, 'a kind of discourse', to borrow one of Lyotard's (1984, p. 3) expressions. The shoe – especially when the Syntilay contamination described in Box 5.2 is absent – is not really meant to improve walking (as is the case with Syntilay in Chapter 6, whose core offering lies in its individually tailored fit), but to 'experience nature creatively and respectfully': it becomes far *more sign than substance*.[10] To quote another of the seminal fathers of postmodernism, Baudrillard (1998, note, incidentally, his anticipatory reference to separate 'groups' – a tendency that, as we have discussed, may be further amplified in the AI era):

> [Y]ou never consume the object in itself (in its use-value); you are always manipulating objects (in the broadest sense) as signs which distinguish you either by affiliating you to your own group taken as an ideal reference or by marking you off from your group by reference to a group of higher status. (Baudrillard, 1998, p. 61)

The problems of accumulating the specific human capital required for knowledge-intensive entrepreneurial projects (see Table 5.1) – arising from the aforementioned seniority-biased technological change and the 'AI-Becker problem' – may only further exacerbate this 'postmodern' risk. The same holds for the increasingly digital and service-oriented nature of the post-essential era. In contemporaneous advanced economies, where agriculture and manufacturing no longer dominate, services account for the bulk of national wealth – with digital services, in particular, representing a substantial share of discretionary consumption (e.g. media, communication technologies,

[9] Readers more familiar with the literature on 'postmodernity' and 'postmodernism' will not fail to notice how 'the shoe' – or rather its representation in art (particularly in Van Gogh's *A Pair of Shoes*, Warhol's *Diamond Dust Shoes*, and Magritte's *Le Modèle Rouge*) – is precisely one of the first examples chosen by Jameson (1991, pp. 6–9) in his landmark book within this transdisciplinary current, to illustrate how different representations of the shoe evoke entirely different worlds (including inner ones), and thus become entirely different 'symbols'.

[10] In this respect, it is clear that the application of 'design-driven innovation' paradigms (in particular the concept of 'innovation by meaning'; e.g. see Verganti, 2017), when directed towards *thin innovations*, may easily degenerate into a search for new symbols rather than the solution of (socially relevant) problems.

entertainment). Within this context, even if consumer demand does not revolve exclusively around the next trending app or gadget, such products today occupy a disproportionately larger niche compared with previous periods (e.g. Fang, 2019; Cochoy et al., 2020; Stocchi et al., 2022). The (potential) phenomenon of the younger generations' 'digital dementia' (Spitzer, 2012), mentioned in Chapter 2 – which, in due time, will itself age to make room for an even more digital generation – can only reinforce the consumer appetite that sustains this very form of *digital thin innovation*. Thus, entrepreneurs may become ever more inclined to develop minor digital applications rather than pursue any genuine technological breakthrough, in what risks becoming a societal vicious cycle.[11] In this respect, the already noted low barriers to entry in low-tech and less knowledge-intensive sectors (including the limited technical skills required; see again Chapter 5), combined with the growing assistance provided by 'AI entrepreneurial buddies' – which increasingly enable forms of solopreneurship not only for the technologically unemployed but also for the expanding ranks of freelancers in the 'gig economy' – constitute additional elements that render a *thin innovation* scenario progressively less unrealistic.[12]

AI Giants Do Not Necessarily Dislike Thin 'Innovative' Entrepreneurship

The last, though certainly not least, element we draw attention to is that AI giants have little, if any, incentive to curb this type of innovative entrepreneurship on their platforms – in fact, quite the opposite.

[11] The mobile app market is already characterized by an enormous number of applications, the majority of them marked by little or no success and, in any case, by a very short lifespan, with overall average profitability remaining low (see the incipit of Goulart et al., 2026; needless to say, 'low profitability' refers to the sellers, though not for the AI giants that market them). For the very reasons just set out, AI is bound to accelerate this trend which, from the perspective of global social welfare, inevitably evokes a scenario of misallocation of resources and productive capacities.

[12] Although we are here making a general point about this postmodern risk – which, in principle, may affect all four types of AI start-ups identified in our taxonomy in the previous chapter – it is probably legitimate to expect that the threat is most relevant to our fourth category of 'creativity-augmented' start-ups, where AI-enhanced creative capabilities are directed towards thin (as we might consider the case of our 'SoleScribe Snack Shoe' example) rather than thick innovation.

To the extent that thin innovative entrepreneurship, while not particularly ambitious, does not jeopardize the quality of the platforms themselves and continues to rely on services provided by the AI giants (first and foremost their *inferential infrastructure*), there is no intrinsic reason for them to halt it. *Thin innovations* are both safe and profitable for the ecosystem's firms, and are thus, in a sense, actively encouraged by these actors, which have no inherent interest in placing limits on the potential frivolity of this 'disorderly expansion of capital' (Srnicek, 2025).[13] Part of this innovation, to be sure, may not only be frivolous and therefore, in some respects, innocuous. While not producing any major impact, some *thin innovation* may also take the form of what Ezrachi and Stucke (2022) describe as *toxic innovation* – that is, innovation that destroys (social) value rather than creating it. Emblematic is their following passage:

> The innovations within the Tech Barons' ecosystem are invariably skewed to attracting and sustaining our attention – whether watching Apple TV, YouTube or Prime videos, posting on social media sites, texting, playing online games, or browsing the web. That means less time for other activities – like sleep, studying, and doing things with other people that can actually improve our well-being. Millennials, reportedly, spend an average of 5.7 hours using their smartphones each day, of which 48 minutes are dedicated, on average, to texting. (Ezrachi and Stucke, 2022, p. 118)

We do not necessarily agree with all the examples provided, nor do we regard all the digital activities listed above as inevitably toxic in the sense given by the authors. Nonetheless, we believe their core message holds: namely, that much innovation channelled through platforms is not necessarily oriented towards socially desirable outcomes. Indeed, as the authors themselves highlight elsewhere, the *predictive* (and increasingly *generative*) capacities of AI may further accelerate this 'race to the bottom' in exploiting human instincts.

Finally, even when innovative entrepreneurship turns out to be neither *toxic* nor, for that matter, *thin*, AI giants can still rely on the

[13] The expression of 'disorderly expansion of capital' was first popularized by Xi Jinping and the Chinese Communist Party following the 2021 crackdown on Chinese Big Tech. As Srnicek (2025), drawing on Zhang (2024, p. 98), observes, this regulatory intervention also sought, alongside the officially stated reasons, to curb the proliferation of frivolous digital products on AI giants' platforms – products that reinforced China's celebrity fan culture and risked undermining social stability and welfare.

'Systemic Lock-in Model' discussed earlier. They can readily intercept promising ventures – whether to block them, absorb them, or institutionalize them as allies – and then position themselves as promoters of thicker forms of innovation, which ultimately serve only to reinforce their ecosystem. The dynamic is general, but it is clearly illustrated in the mobile app segment: countless start-ups emerged, but those that became too important (Instagram, WhatsApp, YouTube, etc.) were swiftly acquired by MAMAA, while the rest were left to fill niche roles within Apple's or Google's platforms. A similar logic may easily characterize the world of LLMs, with heavyweight, equity-based alliances such as OpenAI–Microsoft (and more recently, NVIDIA) and Anthropic–Google.[14]

A Postmodern AI Entrepreneur Playing into a 'Fenced Capitalism'?

All in all, we can synthesize these risks in the broader peril that, at the societal level, too much of innovative entrepreneurial activity unfolds within the boundaries of the AI systems imposed by the tech giants. This dynamic potentially extends their *longa manus* across different sectors and domains, (implicitly or explicitly) encouraging the pursuit of thin innovations in low-tech areas by a sort of postmodern entrepreneur – certainly relying on their (generic) human and social capital, but often acting in solitude with the sole help of 'AI entrepreneurial buddies'. By contrast, thicker forms of innovation, pursued by entrepreneurs endowed with (specific) human capital in high-tech and knowledge-intensive sectors, risk to an even greater extent than in the

[14] The cases cited here are merely the most emblematic examples of the AI giants' activism towards LLMs, and it seems reasonable to expect that further partnerships (and perhaps consolidations) will follow in due course. Relatedly, it is also worthwhile to recall here another of the most prominent recent entrepreneurial initiatives in the AI world: Perplexity.ai, the successful AI-powered search engine that relies on different LLMs. This start-up has already been mentioned in Chapter 6, and in the taxonomy developed there it serves as an example (among others) of an increasingly multifunctional AI-based start-up combining retrieval and generative functions and exhibiting agentic features. What we emphasize here is that, even in this case, the influence of the AI giants cannot be considered negligible, given the significant investments by Jeff Bezos (Amazon) and NVIDIA, among others. Moreover, much speculation has identified it as a potential acquisition target for Apple and Meta (Verdict, 2025).

past ending up in the meshes of the control exercised by the 'Systemic Lock-in Model' of the AI giants, which may easily take the form of 'control beyond ownership' (Cirillo et al., 2025) or 'control without ownership' (Srnicek, 2025). Needless to say, important innovations may and do occur in sectors outside their influence (though with the risk of becoming fewer and fewer) and also within their 'territory', but more likely steered to their benefit, and not necessarily (albeit, of course, we lack precise counterfactuals) to the benefit of social welfare. As once again noted by Ezrachi and Stucke (2022):

> Tech Pirates [i.e. highly innovative start-ups, author's note] will become rarer. While many start-ups might proclaim themselves to Tech Piratehood, their spirit will be institutionalized. They will align with the prevailing Tech Baron's value chain. To put it differently, in all but uniform, they will have joined the navy. (Ezrachi and Stucke, 2022, p. 100)

The result is what we label as the risk of a *fenced capitalism*. It is capitalism, since it still contains what we regard as its distinctive traits already recalled (i.e. private ownership, profit orientation, use of markets). It is capitalism which, though with some reservations, can still function 'well', to borrow Bowles's (2013) phrase, insofar as success still arguably depends on a combination of the three ingredients (i.e. ability, effort, and fortune) – although the weight of luck is likely increasing particularly in the sphere of thin innovative entrepreneurship, possibly to levels that may even become socially dangerous. But all of this remains fenced in – not so much and not only by *formal institutions* (i.e. laws, norms, regulation), as has always been the case in (all varieties) of modern capitalism, but increasingly by the 'rules of the game' set by the corporate world itself. This, clearly, cannot but reverberate on the *informal institutions* that shape society. And, in turn, this entire dynamic cannot but concern the *entrepreneurship policy* of the future.

8 *AI Entrepreneurship Policy and Institutional Reforms*

On the 'Market Failure' Argument

Public intervention in support of innovative entrepreneurship has historically found its main rationale in the classical 'market failure' argument (Parker, 2009; Grilli, 2014). As recalled in Chapter 1, innovative start-ups – because of their very intrinsic characteristics – are generally regarded as facing greater difficulties than other endeavours in accessing labour markets (e.g. Rocha and Grilli, 2023) and, above all, in securing the financial resources required to transform an entrepreneurial idea into a tangible venture (e.g. Storey and Tether, 1998), with the possibility, but never the certainty, of success. Given their potential to generate positive spillovers for the overall system, and the consequent risk of global underinvestment with respect to social welfare (e.g. Hall, 2002), they deserve policy attention – something that, in fact, has consistently been the case across advanced economies over time.

Our analysis must inevitably begin here. Based on the arguments developed so far, what can we say about the extent to which the picture outlined earlier changes as a result of AI?

The first consideration is that certain failures and hurdles that start-ups encounter in the labour and capital markets may, to some extent, be circumvented or ameliorated in specific sectors or for particular types of activities. For example, this may occur due to a reduced need for capital and labour inputs for ventures (particularly, though not exclusively, in less technologically intensive sectors), or due to a lesser need for a pre-existing 'social capital' endowment when AI-augmented 'solopreneurs' or small-size teams seek access to stakeholders' doors. Or again, even when the needs for inputs and complementary assets remain intact, AI may reduce certain forms of informational frictions – such as screening and matching between investors or job seekers on the one hand, and prospective entrepreneurs on the other – and thus

diminish the asymmetric information between the parties that initially gave rise to these market failures.

The second consideration, however, is that this cannot be regarded as a general rule. While AI may attenuate traditional market imperfections in some contexts, in others it risks exacerbating them. For example, as discussed in Chapter 5, it is reasonable to expect that AI will likely lead to an increase in both the *intensive* margin (i.e. more firms financed per investor) and the *extensive* margin (i.e. more investors active in the market) of the VC industry, yet venture capital remains almost by definition an elitist form of finance (Wood, 2020). At the same time, initial financial investments may remain or even become (very) high for prospective ventures (particularly, though not exclusively, in more technologically intensive sectors), and the same applies in the labour market, where the search for specific high-level human capital beyond founders can be impervious (e.g. the 'golden talent trap' for highly skilled AI engineers and similar profiles) and is still characterized by important information asymmetries, given the inherent Knightian uncertainty of truly innovative projects. Moreover, screening and matching AI algorithms may still suffer from human-induced biases, and if not, they are in any case always subject, at least to some unavoidable degree, to 'brittleness' due to *concept* and/or *data* drifts. Thus, a first general observation we can draw is that the 'market failure' rationales for innovative entrepreneurship policy have not necessarily disappeared, and therefore any claim that short-run policy interventions in this domain are becoming less needed would be largely misguided. In other words, the notion that AI might reduce the need for subsidies, grants, or targeted support for innovative start-ups is, in our view, inaccurate. Indeed, it is precisely in the sectors where the societal returns to entrepreneurship are potentially the highest – such as knowledge-intensive and AI-based industries – that public intervention is likely to be most needed.

Bring AI into Entrepreneurship Policy-Making

In fact, an ideal well-functioning innovative entrepreneurship policy, one that is maximally additional to the collective benefit, should be able to identify only those entrepreneurial projects left aside by private investors (in their various forms, such as VCs, banks, or crowdfunding, etc.), but which, if publicly funded, may turn out to enhance

social welfare (Colombo et al., 2011; see Parker, 2009 for a very similar view). An outcome, naturally, that does not exclude, but rather often coincides with, the commercial success of the entrepreneurial endeavour. Obviously, the scientific literature has long questioned the real possibilities for policy makers to even come close to this ideal state (Parker, 2009), with positions ranging from more optimistic (e.g. Lerner, 2002) to much more sceptical (e.g. Shane, 2009; see Lerner, 2020 for a discussion). However, it is not the purpose of our analysis here to delve into this vast literature. But more simply, our second observation is that it is obvious that, if there is to be an active policy, it would be shortsighted for policy makers not to use the same type of AI tools that private investors, particularly VCs, employ in the phases of selection and, where relevant, in coaching or otherwise supporting the investment made. Or rather, not necessarily the identical tools, but the same category of instruments, since policy makers, for the reasons already highlighted, pursue a different objective function. Yet nothing prevents them – and indeed our point is to encourage this – from doing so: just as VCs collect their own data and build predictive models for the screening and due diligences phases, policy makers can do the same. There are already interesting initiatives in this direction, albeit still at a developmental or experimental stage, on both sides of the Atlantic – for instance, the European Innovation Council (EIC) AI Platform tool and the 'AI-enabled diligence evaluation' system (AIDE) proposed by the US Department of Defense. Likewise, policy makers could build and/or fruitfully employ new or existing AI tools of various kinds to support the ventures they finance (what we might term 'AI policy buddies' for entrepreneurial support). Pilot experiences of this kind already exist in different parts of the world, though for now mostly in the form of AI-driven chatbots providing personalized advice to entrepreneurs. For instance, in the UK, the Department for Science, Innovation and Technology has launched a chatbot through which entrepreneurs can ask questions about regulations, taxes, and available support programmes, receiving tailored answers that collate information from across government websites. Although in a different context, the analysis by Otis et al. (2024), already discussed in Chapter 5, shows that under certain conditions the contribution of such instruments can be genuinely additional. Nothing stops the number, scope, and depth of the support offered by these 'AI policy buddies' from expanding

further, potentially serving not only entrepreneurs but also policy makers themselves in carrying out coaching functions envisaged by more interventionist entrepreneurship policies (e.g. governmental VCs, public incubators). Of course, assistance in both the selection and coaching functions comes with all the limits already discussed for VCs in Chapter 5, but clearly also with the opportunity for heterogeneity and superadditive effects arising from hybrid intelligent human–machine interactions, as we highlighted there.

Why 'More Entrepreneurs' Remains Bad Public Policy

The abovementioned AI capabilities applied in policy-making may also help solve a problem that AI itself has, admittedly unintentionally, generated. But, as a matter of fact, a great deal of hype and 'AI-washing' surrounding the technology has undoubtedly emerged, and clearly the decision as to which innovative projects deserve public support and which do not is becoming increasingly complex and costly in certain contexts. Moreover, as we have seen in Chapter 7, the possibility of a 'composition problem' cannot be ruled out a priori: the drastic reduction of entry barriers in some sectors may easily favour *thin* rather than *thick innovations*, pursued more by *postmodern or Veblenian entrepreneurs* than by *Schumpeterian ones*. In this sense, in our view, Shane's (2009) famous dictum – already powerfully conveyed in the title *Why Encouraging More People to Become Entrepreneurs Is Bad Public Policy* – can only be regarded as further reinforced. After all, the empirical literature has often shown that mere subsidies to entry may generate distortive effects on market dynamics (e.g. Santarelli and Vivarelli, 2002) and that, conversely, more articulated and comprehensive policy instruments – those that do not simply target entry but rather support the growth activities of start-ups across multiple domains – may exert more beneficial effects (Grilli et al., 2023; see Audretsch et al., 2020 and Matricano, 2024, for a review of some of these most recent policies).[1] These beneficial effects may also be indirect (but not less

[1] Admittedly, although there are indeed multiple geographical contexts and historical periods in which more systemic policy approaches to innovative entrepreneurship have been undertaken, much of the recent evidence – which, with different facets and nuances (see Matricano, 2024) overall offers a

important), for instance, by encouraging individuals endowed with high (specific) human capital to pursue an entrepreneurial career in innovative sectors which, in the absence of such policies, they would likely not have considered (e.g. Grilli et al., 2023). It is clear that in the presence of *adverse selection* phenomena in AI-based innovative entrepreneurship, such as those highlighted in Chapter 6, if there is to be an AI-based entrepreneurship policy, the idea of adopting a more systemic approach is, for this reason as well, further reinforced.[2]

(Potential) Perils for the AI Entrepreneur Equate to Areas of Concern for (Entrepreneurship) Policy

As we noted in Chapter 1, the goal in innovative entrepreneurship is 'quality' rather than the mere 'quantity' of entrepreneurs (Shane, 2009), and an 'entrepreneurial society' (see Economidou et al., 2018) should be designed to enable those with the capabilities to experiment and (potentially) succeed as innovative entrepreneurs to realize their opportunities, while dissuading others who are less suited to this path. Crucially, such a society should also be able to channel entrepreneurial talent towards socially relevant activities, thereby enhancing both private and social returns (Colombo et al., 2011).

In Chapter 7, we identified several potential perils for the AI entrepreneur across different levels of analysis. From an industrial organization perspective, we underscored the *incumbent advantage* and the *systemic lock-in* enforced by AI giants. From a macroeconomic standpoint, we highlighted the risks of what is, in some respects, an inevitable *turmoil in the labour market* brought about by a general-purpose technology such as AI – accompanied by a range of non-neutral effects on different categories of workers (e.g. junior employees and mid-skilled profiles), as well as the potential emergence of a 'human capital gap' in innovative entrepreneurship (e.g. the *seniority-biased*

reassuring picture of results – is based on the case of the so-called Italian 'Start-up Act' (e.g. Giraudo et al., 2019; Manaresi et al., 2021; Biancalani et al., 2022).

[2] The wedge in entrepreneurial opportunity costs between individuals with high and low AI-engineering skills, as reported in Figure 6.2, will certainly not be eliminated but at least alleviated, thereby improving overall social welfare.

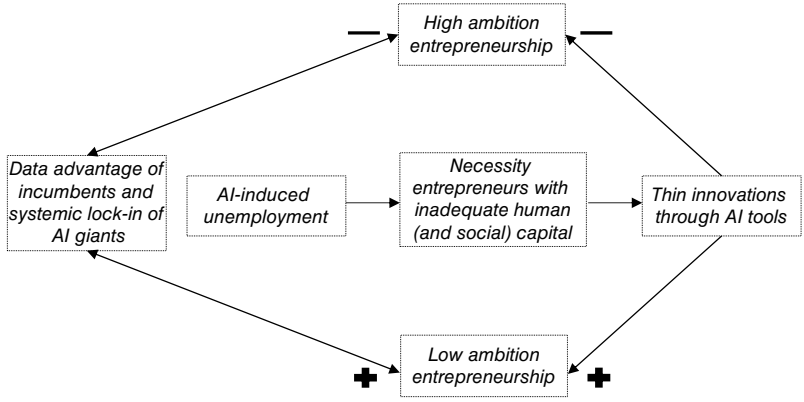


Figure 8.1 The vicious-cycle risk of postmodern entrepreneurship

technological change). Finally, at a sociological level, we noted the heightened likelihood – greater than in the past, precisely because of AI's capabilities – of *postmodern entrepreneurship* phenomena proposing *thin innovations* potentially targeting closed and polarized communities of interest.

And moreover, as only briefly mentioned earlier, there is also a broader meta-risk. Many of these factors – macro-level labour disruption, post-essential consumption, low-barrier entrepreneurship, and the dominance of AI giants – may converge to form a sort of perfect storm that might curtail the otherwise positive impact that AI could, in principle, exert on innovative entrepreneurship and the dynamism of capitalism. In fact, each of these forces may easily reinforce the others. The data incumbent advantage and, more critically, the systemic lock-in by AI giants may render ambitious and genuinely innovative entrepreneurial efforts increasingly impervious. AI-induced unemployment may expand the ranks of necessity entrepreneurs, in some measure deprived of the possibility to accumulate relevant human (and social) capital, and thus more often engaged in thin innovations on non-essential products, where AI tools make it cheaper and faster to launch low-value ventures. Finally, the ecosystem structure of AI giants, potentially sustained by State action, tends to privilege incremental and complementary forms of innovation. Figure 8.1 offers a visual representation of this potential *vicious-cycle risk in postmodern entrepreneurship*.

Assigning a probability distribution to the magnitude of these possible perils in modern advanced economies, and to the vicious circle that might arise from them, is obviously beyond the scope of this chapter, and we shall not attempt such an exercise, not least because it would be highly contingent upon a wide range of contextual factors, including, most prominently, the geopolitical dimension.[3] What we do strongly emphasize, however, is that these perils should translate into areas of concern for social welfare that should guide policy makers' attention. Regardless of whether one perceives them as imminent threats today or not (provided, of course, that one is not persuaded that – now or in the future – all of them represent entirely null risks, something I hope by this point to have convincingly dispelled), the old adage still applies: *prevention is better than cure*. This traditional proverb holds true in the personal sphere, of course, but the principle, as we know, also applies in the fields of public health and environmental policy. We are not necessarily supportive of applying the 'precautionary principle' to policy in all cases and under all circumstances, especially when it comes to fully developing a technology.[4] Yet policy can certainly act upon the technology's envisaged applications and effects (see Horowitz, 2025, for a similar view), and ultimately upon the institutions that will – or will eventually – shape its usage. There is, therefore, no reason to believe that the proverb should not equally apply to (entrepreneurship) policy aimed at those institutions in the age of AI. It follows, then, that (innovative entrepreneurship) policy-making can act proactively in this domain. But how?

[3] It is worth noting that the probability of the meta-risk of a 'perfect storm' cannot be regarded as negligible, even though it is conditioned on the simultaneous occurrence of all underlying events. This probability is not merely the product of independent probabilities; rather, it results from events that, for the reasons outlined in the text, may reinforce one another and are therefore likely to exhibit positive dependence. In the simplified case of three underlying events, each with probabilities x, y and z of occurring, the probability of the meta-risk is therefore not given by the mere product xyz (as under independence), but by their joint occurrence under positive dependence. In such circumstances, the overall probability may rise well above xyz, potentially approaching the upper bound of $min\ (x, y, z)$.

[4] We are even more reluctant when precautionary regulation on technology (AI in our case) is unilateral rather than global, thereby affecting only certain institutional contexts (as exemplified by the EU AI Act of 2024). We will return to this point in the final section of this chapter.

The use of both sets of parentheses above is, of course, intentional, as it conveys an important point regarding the answer to this question: it is by no means guaranteed that an ad hoc policy will exert greater effects on its immediate target – in our case, innovative entrepreneurship – than broader institutional reforms. This may depend on general reasons related to policy design, but in our context it also reflects motives inherent to the domain itself, such as the high degree of (Knightian) uncertainty (Bradley et al., 2021) or definitional issues concerning what qualifies as an innovative start-up (Audretsch et al., 2020). In any case, the point has two important consequential implications. The first, quite naturally, is that 'entrepreneurial' problems do not necessarily need to be addressed through policies specifically designed for that purpose.[5] The second, equally important, is that given the centrality of entrepreneurship within the capitalist process, even wide-ranging reforms must safeguard its orientation towards 'productive' forms and, in any case, take into account the effects they generate on the kind of entrepreneurship that sustains 'good' capitalism (Baumol et al., 2007). This, in truth, does not seem to occur very often. Indeed, broad-spectrum institutional reforms pursued with objectives other than entrepreneurship have been found to produce 'unintended consequences' (mostly negative) for innovative entrepreneurship in many different areas (see table 1 in Bradley et al., 2021 for some examples).[6]

Bring Entrepreneurship into AI Regulation and Policy-Making

The case of AI regulation and policy-making appears to be no exception to the limited attention typically devoted to innovative

[5] This affirmation is, of course, fully consistent with the institutional view of entrepreneurship (Baumol, 1990). Yet expressed in such a direct and straightforward form, I must, in truth, credit it to the late David J. Storey – one of the leading figures in the study of entrepreneurship policy – and to a conversation we had in 2014 at the University of Sussex.

[6] This issue is part of a broader discourse within the policy sciences, where the coordination and integration of mutually influencing measures, though difficult to achieve, is nevertheless regarded as an objective of primary importance (e.g. Peters, 2018).

entrepreneurship within broader policy frameworks. So far, entrepreneurship has been virtually absent from discussions on AI regulation, a feature that also seems to characterize the diverse regulatory approaches adopted in different institutional contexts – from the ethically interventionist stance of the European Union, to the more laissez-faire orientation of the United States, to the State-surveillance model of China (Krarup and Horst, 2023; Wang et al., 2025). The regulatory agenda (where present) has understandably focused on issues such as privacy, security, liability, and ethical concerns. The point here is not to downplay the centrality of these classical issues of AI regulation, but rather to emphasize that the entrepreneurial dimension should be regarded as an integral part of the discussion. This does not seem to have occurred, judging from the legislative texts and regulatory measures adopted so far. Indeed, while the textual occurrence of words and expressions associated with innovative entrepreneurship can by no means be taken as evidence of substantive engagement, it nevertheless represents an important signal in this regard. Yet the evidence available suggests that, if not entirely absent (Wang et al., 2025), this dimension has so far been addressed only in a highly marginal way (Saheb and Saheb, 2023). While the broader concept of innovation does, of course, find some resonance, terms such as 'start-ups', 'entrepreneurship', 'innovative entrepreneurship', and their hybrid combinations are, by contrast, almost never mentioned. Our own investigation, presented in Box 8.1, confirms this result – on a different corpus of texts than those cited in the literature above – if further confirmation were needed.

In our view, this calls for a change of course. As we have noted, the empirical literature (admittedly not abundant) has documented cases in which broader regulatory and policy interventions have produced effects on entrepreneurship that were not always positive. Within the field of our interest, perhaps the most widely investigated general policy in recent years with regard to its effects on innovative entrepreneurship has, in all likelihood, been the General Data Protection Regulation (GDPR). Officially introduced by the European Union in 2018, the GDPR aimed at strengthening individuals' control over their personal data, harmonizing data protection rules across Member States, and seeking to promote greater accountability in the digital economy.

Box 8.1 Where is entrepreneurship in AI regulation(s)?

There exist several international sources that provide a structured overview of AI regulation across countries, such as the OECD.AI Policy Observatory, the AI Laws Around the World Tracker (AIPRM), the Global AI Legislation Tracker by the International Association of Privacy Professionals (IAPP), and the already cited *AI Index Reports* of the Stanford University.

To conduct our qualitative analysis on the relevance of innovative entrepreneurship within this context, we selected the *Artificial Intelligence and Democratic Values 2025 Report* by the Center for AI and Digital Policy (CAIDP). This latter is a non-partisan institute engaged in promoting democratic values and human rights in AI governance – actively contributing to initiatives such as the UNESCO Recommendation on the Ethics of AI, the Global Partnership on AI, and the Council of Europe AI Treaty.

The report – approximately 1,500 pages long – aims to comprehensively catalogue all existing AI regulatory frameworks and policies worldwide, condensing them into detailed country-by-country profiles. This approach, by its very nature, lends itself particularly well to cross-geographical comparison. A systematic textual and lexical analysis was therefore conducted using NVivo.

The first, rather unsurprising result, is the virtual absence of terms such as 'start-up(s)' or 'entrepreneurship' from the profiles of the three most relevant economic areas: China, Europe (in particular United Kingdom, Germany, France and Italy), and the US. When examining the top 100 most frequently cited words in each of these profiles, these two terms are absent across all cases. This finding, however, is not particularly surprising – especially in the European context – given that across the 144 pages of the 2024 EU AI Act, comprising approximately 90,000 words, the term 'start-up(s)' appears only once.

Interestingly, aside from the obvious prominence of the two words 'artificial' and 'intelligence', the most frequently cited word (including its variants) in the CAIDP report is 'technology' in both the American and Chinese contexts. In contrast, the most recurrent term in the European context is 'government', while 'technology' only appears in tenth position. Although, as noted, words are not facts, such a lexical divergence seems to reflect a deeper difference in institutional and policy orientation: while in the US and China the regulatory discourse appears more closely tied to technological and industrial considerations (of course with many substantial differences between the two), in Europe it remains primarily framed within the language of public governance and administrative oversight.

The Case of the EU Regulation

The results emerging from GDPR policy evaluation exercises on the EU society are certainly not easy to interpret; yet, with regard to innovative entrepreneurship, the evidence appears to lean strongly towards the negative, considering a variety of measures (e.g. Jia et al., 2021).

Our purpose again here is not to provide an exhaustive review of this literature – comprehensive surveys of the broader effects of the regulation already exist (e.g. Johnson, 2022; Li et al., 2025), notwithstanding its relatively recent introduction – but rather to highlight the mechanism which may originate this negative 'unintended consequence', since it may easily be extended to other types of (general) regulatory intervention. Compliance with the GDPR inevitably entails costs for firms. As Bessen et al. (2020) show, many AI start-ups had to reallocate scarce resources or even create new positions to deal with GDPR-related issues, despite their small size and limited means. Since these compliance costs are, by their very nature, largely fixed, they weigh disproportionately on start-ups that cannot spread them across a large scale of operations. In relative terms, this asymmetry constitutes an indirect advantage for (large) incumbents, while representing an additional burden for entrepreneurial ventures.[7] Similarly, negative consequences for innovative start-ups – though the mechanisms at play in this case probably remain unclear and call for further research – also seem to arise from more recent EU legislative and regulatory interventions, such as the Digital Markets Act of 2022 (e.g. Yang, 2026) and the EU AI Act, announced in April 2021 and effectively entered into force in 2024. In Box 8.2, we document that even the mere announcement of the latter may have generated a significant negative impact on the birth rate of AI-based innovative start-ups in Europe.

All this (growing) evidence, in turn, suggests two things. First, that there are no easy recipes in the policy design context, since the policy maker often has to deal with trade-offs. In the present case, the evident

[7] It is indeed possible – if not plausible – that compliance costs will decline over time as firms acquire the relevant competencies and a genuine market emerges, with specialized providers offering compliance-related services and others demanding them. However, the transitional phase may be prolonged, and the price to be paid will, by definition, remain relatively more burdensome for start-ups.

Box 8.2 The 'unintended consequences' of the EU AI Act on European AI-based start-ups

The European Union Artificial Intelligence Act represents the first attempt to introduce a comprehensive regulatory framework for AI within the EU single market. Initially proposed in April 2021 and formally approved in 2024, it entered into force on 1 August 2024 (with provisions entering into operation gradually). Its structure relies on a tiered, risk-based classification distinguishing four categories: unacceptable, high, limited and minimal risk. Each category is associated with specific regulatory obligations, where high-risk applications (i.e. those in healthcare, education, critical infrastructure and employment), have to comply with the strictest requirements, including conformity assessments, enhanced transparency and human oversight. Concerns about the potentially adverse effects of the Act on Europe's capacity to innovate in AI have been voiced across multiple fronts and by various sources, not least the Draghi report (European Commission, 2024), and have likely contributed to recent reform proposals by the European Commission (e.g. the 'Digital Omnibus'). As a matter of fact, the analysis presented here – while preliminary and, as we argue later in the chapter, in need of much deeper investigation – certainly does not appear to dispel such worries.

Drawing once again on the Crunchbase database, which provides information on 55,087 AI-based start-ups founded between 2005 and 2023, we examine the potential effects of the announcement of the EU AI Act in April 2021 on the creation of AI-based firms, focusing specifically on a comparison between Europe and the United States. Over the period of interest, Crunchbase reports the formation of 12,130 AI-based start-ups in Europe (including the UK) and 19,067 in the US, while 23,890 were founded in the rest of the world.

Figure 8.2 presents a comparative analysis of start-up formation in the two regions. Figure 8.2a reports results for the full sample of firms, for which we have precise information on the year of foundation. Figure 8.2b focuses instead on the subsample for which reliable information is available at the month–year level (a total of 11,025 start-ups). In both cases, a preliminary visual inspection suggests the emergence of a widening gap between the two areas from the time of the announcement onwards.

To substantiate this descriptive evidence and assess whether the observed divergence is statistically robust, we also conduct a difference-in-differences (DiD) econometric analysis over a time window spanning

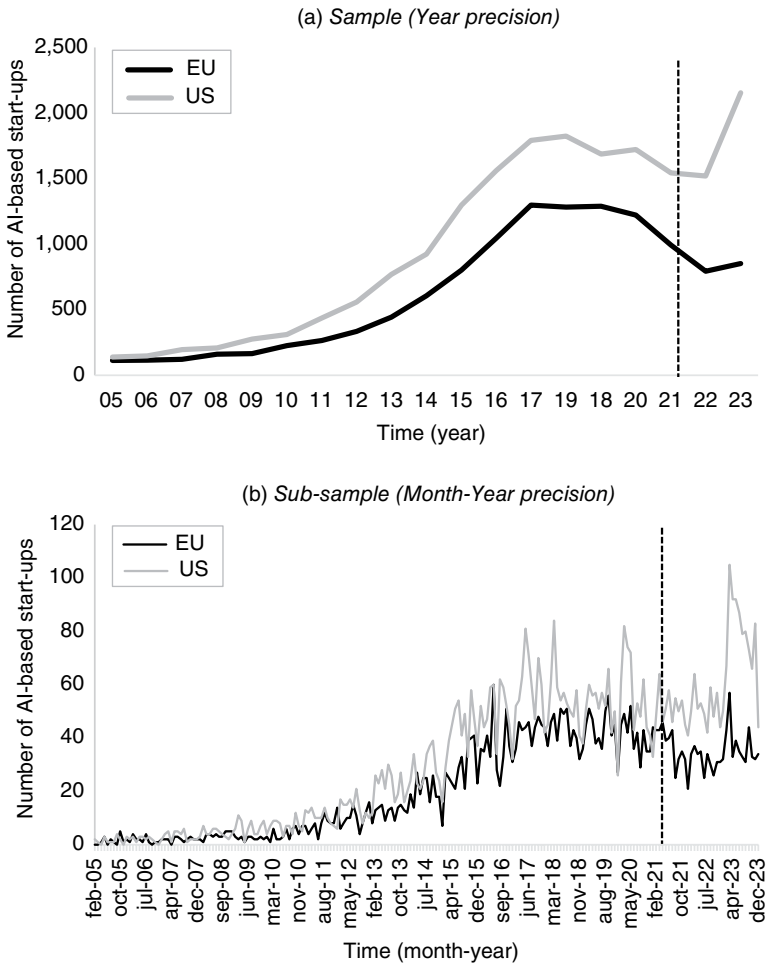


Figure 8.2 Comparison EU vs. US on the number of AI-based start-ups created (2005–2023). Evolution of the number of AI-based start-ups in the EU (including the UK) and the US; (a) includes all start-ups (a total of 31,197), while (b) considers only those for which reliable month-year information is available (a total of 11,025 start-ups). In the latter case, beyond missing data on the month of foundation, we also exclude January for methodological consistency, as it systematically exhibits anomalous peaks across both geographical areas. As discussed in the text, however, the econometric findings remain robust even when these observations are incorporated through Monte Carlo simulations. The dashed line indicates the announcement of the EU AI Act (April 2021, approximated in (a)). Source: our elaboration based on Crunchbase data.

2018–2023, ensuring temporal symmetry. Formally, we estimate the following two-way fixed effects specification:

$$Y_{it} = \alpha + \beta\left(Treatment_i \times Post_t\right) + \mu_i + \lambda_t + \varepsilon_{it}$$

The dependent variable captures the number of start-ups created in i (Europe – US) at time t (month – year); *Treatment* is a dummy variable that identifies Europe-based start-ups (and 0 for US ones); *Post* is a step time dummy variable equal to 1 from April 2021 onwards (i.e., following the official proposal of the EU AI Act presented to the European Parliament), and 0 otherwise; μ_i are unit fixed effects, controlling for time-invariant characteristics specific to each unit; λ_t are time fixed effects, capturing shocks or common trends affecting all units at a given point in time; ε_{it} is the idiosyncratic error term.

The estimation results suggest that the AI Act (more precisely, its initial proposal) may have imposed an additional reduction of up to approximately one third in the monthly volume of AI-based start-up formation in Europe, over and above the pre-existing gap between the two regions.

The implementation of several robustness checks – including, notably, tests for the parallel tend assumption and a series of Monte Carlo simulations in which a random month of foundation was assigned to firms excluded from the main analysis due to missing information on such dimension – corroborates the results reported above, both in terms of magnitude and statistical significance.

trade-off is between privacy, consumer protection, and security on the one hand, and greater entrepreneurial dynamism on the other, with all that this may entail (e.g., an increased scope for 'creative destruction' dynamics, SM I). Second, it is of paramount importance for policy makers to elucidate the most relevant trade-offs, since policy design can then be adjusted accordingly to achieve higher levels of social welfare. Put simply, once the trade-off at stake is understood, nothing prevents policy makers – given the inherently malleable nature of regulation – from introducing exemptions for specific categories of actors, namely, small firms below a certain size threshold.[8] To privilege one

[8] The recent 'Digital Omnibus' reform proposal put forward by the European Commission – which, in its current form (as of November 19, 2025), would introduce targeted amendments to the EU AI Act and the GDPR that exempt,

side or another in a policy trade-off – that is, to favour one category of actors over another – will always remain, in the end, a political decision. Yet, it is evident that a precise understanding of all the possible ramifications can only facilitate the identification of Pareto-dominant solutions, or at least the softening of the trade-off itself.

Getting Institutional Reforms Right (for Innovative Entrepreneurship)

It is in this spirit that we now turn to the examination of institutional reforms. The objective here is not necessarily to advance fully-fledged policy proposals. Rather, in several instances, our aim is to outline and critically reflect on the plausible pros and cons of selected measures in relation to innovative entrepreneurship, and to stimulate evidence-based research capable of guiding policy makers towards more informed decisions. In doing so, and applying a precautionary principle to policy design – this time, rather than to technology – we suggest that, in order to minimize potential negative and unintended consequences for entrepreneurship, priority should be given to policy interventions that expand the opportunities and endowments available to start-ups, instead of those that focus predominantly on reducing or constraining the existing endowments of incumbents. In order to structure this discussion, Table 8.1 outlines a policy framework that maps each of the four 'areas of concern' we have identified – *incumbent advantage*, *systemic lock-in*, *labour market turmoil*, and *postmodern entrepreneurship* – onto a corresponding 'area of intervention': *data accessibility (and model openness)*, *antitrust enforcement (and gatekeepers ex-ante regulation)*, the *management of technological unemployment*, and *educational systems*. Taken together, these domains provide a structured set of instruments through which, in principle, policy makers can address the underlying challenges. Clearly, these areas of intervention and their associated instruments extend far beyond the specific context of our inquiry, as they lie at the very centre of current debates in academia, policy circles, and public opinion more broadly. But the point is precisely this: to add to this debate also the perspective of innovative entrepreneurship.

simplify, or otherwise ease regulatory obligations for small and medium-sized enterprises and small mid-cap companies – appears to move precisely in this (in our view, desirable) direction.

Table 8.1 ***Institutional reforms and AI entrepreneurship: a policy framework***

Areas of concern	Areas of intervention	Instruments (and challenges)
Incumbent advantage	Data accessibility (and model openness)	Promote and govern vertical 'data commons' and explore and establish more nuanced approaches to IPRs regimes.
Systemic lock-in	Antitrust enforcement and/or ex-ante regulation	Strengthen the 'market investigation' antitrust tool, focusing on the most problematic segments of AI giant dominance (e.g. cloud computing services), and adopt ad hoc, evidence-based remedies to minimize the risk of negative unintended consequences.
Labour market turmoil	Management of technological unemployment	Encourage entrepreneurship-oriented research on universal basic income (UBI) and related safety-net programmes in order to broaden the evidence base for policy-making on social protection in the labour market.
Postmodern entrepreneurship	Educational systems	Integrate AI into (entrepreneurship) education by embedding it across curricula and orienting learners towards grand societal challenges. Cultivate AI-enhanced critical thinking, dynamic orchestration capabilities, and a strong entrepreneurial mindset, while countering both algorithmic aversion and a-critical acceptance.

In doing so, we adhere to a (neo-)institutional view. It is possible, even in the short run, to redirect entrepreneurial trajectories by altering the rewards associated with the different goals that individuals may pursue within a society (Baumol, 1990). Yet any policy intervention (and thus any change in *formal institutions*) that reshapes rewards will, over time, inevitably exert an influence on the cultural aspects

(and, more generally, on the *informal institutions*) that characterize a society, and thereby also on the goals that individuals come to pursue within a given 'institutional matrix' (North, 1990; Williamson, 2000).

1° Policy Proposition: Balance Data Openness and Intellectual Property to Sustain Entrepreneurial Dynamism

Generally, the literature identifies two approaches to limiting firms' informational power (e.g. Prainsack, 2019). The first is an *individual-centric* approach, which seeks to reassert personal control over data and strengthen individual rights over its circulation and use. The second is a *collective* approach, in which data are treated as a shared resource and access is organized through collectively governed arrangements rather than individual permissions. In the context we examine here, that is, the reduction of the informational wedge that separates incumbents from start-ups, the collective route, under certain conditions, appears more promising, as it offers a more realistic pathway to rebalance informational asymmetries.

However, it must be stated from the outset that, in our view, the *incumbent advantage* remains, in many respects, structurally persistent – and, in any case, resistant to full elimination. We have already seen how incumbency enables firms to leverage past (and often exclusive) data to optimize products, personalize customer experiences, and better anticipate market needs (Norbäck and Persson, 2023), with learning effects that clearly compound over time. In this regard, we believe that even the most radical visions of *personal data sovereignty* – such as Solid, the initiative launched by Tim Berners-Lee, the inventor of the World Wide Web, to restore a 'person-centric' architecture of information flows on the Internet – cannot fully dissolve this dynamic. Even under such extreme forms, these instruments cannot enable a form of 'digital forgetfulness' on the part of the seller, whereby the latter truly forgets (or is forced to forget) the transaction and the data associated with it.

Conversely, we believe that emerging experiments with *vertical data commons* – sector-specific pools of information governed through cooperative rules – may not entirely solve the problem but at least have a higher likelihood of alleviating it (especially if data are already 'processed', i.e. structured and immediately 'usable'; see Mihet et al., 2026). In this sense, the recent (and still evolving) initiative of the European

Union's Data Spaces, with the establishment of (still relatively) 'open' information reservoirs in several strategic vertical sectors (e.g. agriculture, tourism, media, among others), represents the most advanced institutional expression of this logic: federated infrastructures designed to rebalance access to data, also for the benefit of start-ups seeking to develop truly data-driven innovations. As we know, 'commons' may easily turn into tragedies (Hardin, 1968) if not properly governed (Ostrom, 1990); and needless to say, the success of this type of initiative, like any policy intervention, will depend on a myriad of contextual and operational factors (e.g. Duflo, 2017). Among them, not least, the quality of the data put in common and the involvement of incumbents, who should in some way be 'rewarded' in the short term for their data contribution but potentially bypassed in the long run, precisely as a result of the Schumpeterian dynamics they helped to trigger. Clearly, in the event that these open data pools succeed (which, for the reasons discussed earlier, is far from obvious), within an ecosystem where information becomes a shared resource, provided that no binding bottlenecks persist in computational capacity – both hardware and software – along the *inferential infrastructure* (Mihet et al., 2026), the locus of competitive advantage in the innovation game will inevitably shift from data to something else. This 'something else' may well include certain elements of an AI model's architecture – how algorithms select, weight, and interpret signals that are, in principle, available to everyone – as well as, of course, the ability to act accordingly by effectively implementing the idea (as we know since Chapter 3, from the story of 'move 78' after 'move 37' in the game of Go, even predictive AI can still foster creativity when combined with field experience and specific human capital). In other words, once the data asymmetry begins to narrow, the true source of competitive advantage increasingly depends on how data are made to speak. The natural consequence is that, while data should arguably remain open to start-ups to ensure fairness and dynamism, the models that translate data into actionable knowledge may well not be entirely open – and this precisely for the sake of social welfare. Total model openness would otherwise risk triggering the 'algorithmic deforestation' described by Hunt and Kurdoglu (2025): a cognitive monoculture in which diversity of perspectives collapses and the competitive landscape becomes barren. Preserving a degree of model opacity could therefore not be a concession to secrecy but a condition for maintaining the epistemic biodiversity upon which

meaningful (Kirznerian) innovation depends. Clearly, AI models are only one (albeit an important) dimension on which start-ups could be able to secure an adequate margin of appropriability, and other components along the value chain may play a similar role. Yet the principle that certain complementary innovative elements should remain protectable – or at least not easily imitable – is a classical result dating back to Teece (1986) and has even been further advanced by the open innovation literature (see, for instance, the concept of 'selective revealing' in Alexy et al., 2013).[9]

The delicate – indeed, extremely delicate – balance between what should be rendered more open (and therefore less protected) and what should, conversely, be more tightly enclosed within IPRs represents, in our view, one of the defining policy challenges of the AI era, with clear repercussions for innovative entrepreneurship. While the scientific and public debate on the use of IPRs has long crystallized around two opposing poles – on the one hand, the Schumpeterian view that regards them as indispensable incentives for innovation by providing the necessary rewards (Schumpeter, 1911/1934); on the other, more open approaches that either downplay their centrality (Boldrin and Levine, 2002) or consider them merely one instrument among others (Cohen et al., 2000) – modern advanced economies will increasingly need to ground social welfare maximization (and, in particular, dynamic efficiency through entrepreneurial action) in a more nuanced equilibrium. What will matter most, from a policy perspective, is not the strength of IPRs per se, but their relevance across different domains of innovation. Thus, the policy challenge lies less in calibrating the strength of IPRs than in defining their scope, determining in which fields protection fosters genuine innovation and where,

[9] Paradoxically, the idea of the 'commons' – once associated with socialist imaginaries of collective ownership – re-emerges here as a mechanism for making capitalism more efficient, fair, and dynamic. By opening access to the raw material of prediction, data commons can curb the monopolistic concentration of knowledge while preserving entrepreneurial vitality. In this sense, they may end up serving capitalism rather than undermining it. Put differently, for capitalism to remain innovative and socially dynamic, it may need to internalize an institutional logic originally conceived to counter it: collectively governed reservoirs of information, analogous to the public infrastructures that sustained nineteenth-century industrial growth. Data commons thus function not as a negation of capitalism, but as its epistemic corrective, preventing excessive concentration of predictive power that could inhibit (Kirznerian) forms of entrepreneurship.

instead, openness better serves the public interest. For instance, data scraping still occupies a legal grey zone (at best) in many jurisdictions (OECD, 2025) – a condition that, for the reasons discussed in Chapter 3, certainly does not favour the prospective entrepreneur. Yet, despite its inherent limitations, scraping can still be an important tool for acquiring the informational material necessary for experimentation, provided it is lawful and not rendered excessively impervious from a technical standpoint. Rights of exploratory access should therefore be safeguarded, particularly for innovative entrepreneurs.[10] On the other hand, once novel insights, products, or models emerge from this process, the resulting intellectual artefacts must be securely owned by their creators. Without such assurance of ownership, the incentive to invest in the entire process, from data scraping and data cleaning to model design and further developments, may be severely undermined. The challenge for policy makers, then, would be to design regimes in which scraping is enabled and protected for exploratory and innovative uses, while appropriation without transformation remains prohibited. Public policies governing the degree to which innovation (and the underlying data) can be formally appropriated should strive for the subtle art of maintaining a porous frontier between *openness* and *ownership*, ensuring that the same institutional mechanism – for example, the 'data commons' – can sustain both collective learning and individual entrepreneurial action.

2° Policy Proposition: Empower Antitrust with the 'Market Investigation' Tool in Key AI Markets

As discussed in Chapter 7, the 'Systemic Lock-in Model' of AI giants – within which the most promising forms of innovative entrepreneurship risk becoming entrapped – stems from the super-additive effects and hardly matchable competitive advantages that arise from their simultaneous presence across all domains relevant to AI. Antitrust is, of course, the traditional area of intervention aimed at preventing this deadly grip from suffocating market competition and deteriorating social welfare. Being, by nature, an ex-post form of regulation – apart

[10] A stricter enforcement of property rights on this front will inevitably benefit large incumbents (including, most prominently, the AI giants), who are far better positioned than any start-up to negotiate licensing agreements or to withstand the legal and financial consequences of illegitimate data use.

from the preventive control exercised over M&A operations – designed to sanction anti-competitive conduct or abuses of dominant positions, it has been increasingly questioned over time as to its actual ability to deal with the specific features of digital competition and, consequently, of AI markets.

On the one hand, the two traditional views of antitrust – the classical Harvard structure–conduct–performance paradigm and the Chicago School approach advocating for its substantial downsizing – have given rise to a far more interventionist and, in some respects, ideological perspective, namely the New Brandeis Movement, which advocates rather structural and intrusive remedies. On the other hand, these criticisms have also prompted a shift of emphasis towards ex-ante forms of regulation, regarded by many as more suitable than ex-post ones for governing access to AI giants' platform-based ecosystems conceived as essential facilities and, therefore, as sort of natural monopolies. The EU's Digital Markets Act is perhaps the emblem of this shift, but several other examples can be found – such as the UK's Digital Markets, Competition and Consumers Act, Japan's Act on the Promotion of Competition for Specified Smartphone Software, or Brazil's Draft Bill Regulating Digital Platforms, to name just a few.

Providing a full account of both debates – namely, the evolving views on antitrust paradigms and the differing perspectives on ex-ante rather than ex-post forms of digital regulation – would take us too far afield and well beyond the scope of our inquiry. Our aim here is simply to highlight a few key observations – arising from the considerations developed so far – concerning the relationship between antitrust (or regulatory) enforcement and innovative entrepreneurship in the age of AI, which may guide policy makers in this matter.

As widely acknowledged in recent analyses (e.g. Padilla, 2024), there is a general consensus that competition law (with nuances varying across jurisdictions) has been subject to a certain degree of under-enforcement, particularly in digital and high-tech markets. The starting point, therefore, is that if something has to be done – and we are fairly convinced of this, given the 'Systemic Lock-in Model' outlined in these pages – *where* and *how* to intervene is far less obvious, precisely because of the potential 'unintended consequences' on the entrepreneurial side. This brings us to a first consideration: these potential unintended consequences need to be better explored, both empirically and conceptually.

First of all, most of these unintended consequences are documented in merely quantitative terms, with little attention paid to possible qualitative effects on the entrepreneurial phenomenon. In other words, it is one thing to say that the number of firms created decreases as an (unintended) consequence of a policy; but such a result tells us nothing about what kind of start-ups fail to emerge, and which, on the contrary, are encouraged by the policy. Put differently, an unintended consequence that reduces the number of new start-ups, but results in those that do emerge being more robust, perhaps more innovative, or otherwise more likely to succeed, would radically change the evaluation of that unintended consequence: from negative to positive (or at least not as negative).[11]

Second, the reasons behind an unintended consequence often remain unexplored and are, I would say, far more frequently inferred in a speculative rather than a direct manner. For example, hypothesizing that the Digital Markets Act in the EU may reduce the number of start-ups created because 'stricter regulations on M&A involving major platform firms could reduce exit opportunities for start-ups' (Yang, 2026, p. 360) presupposes that an entrepreneur with a potentially valuable idea – one that could later attract the acquisition appetite of AI giants – consciously decides to abandon it altogether, fearing that regulatory tightening on a handful of firms would limit the possibility of monetizing it sufficiently. As we discussed in Chapter 6, the presence of AI giants can indeed influence VC decision-making and, consequently, the dynamics of innovative start-ups, including their acquisition prospects. However, at the level of the entrepreneur(s), frankly, such reasoning requires a whole chain of backward inferences that, when multiplied one by another and taken together, seems, at first sight, a somewhat implausible line of explanation – unless supported by direct evidence.

This, in turn, leads us to a third, and certainly no less important, consideration. Even when acknowledging the existence of an unintended consequence and identifying the specific reasons for its

[11] For instance, Bertoni et al. (2025) document how a specific antitrust measure introduced to make collusion more difficult and thus markets more competitive (i.e. a leniency programme) reduces – at least in the short term – the number of new firms created, while at the same time the smaller number of entrants exhibit better survival, productivity, and growth performance compared to start-ups founded before the policy.

occurrence, what may appear as a negative effect in the short term might not necessarily remain so in a longer-term perspective. Take again the example mentioned earlier: granted – though not necessarily conceded – that the aforementioned is indeed the underlying reason, the increased difficulty of being acquired could, over time, encourage innovative start-ups to develop more self-reliant business models and pursue more independent growth paths. In other words, a negative unintended consequence in the short run may not turn out to be such (or at least not as negative) in the long run, if it helps to curb the already noted – and somewhat provocatively labelled – tendency of the 'take the money and run' entrepreneur, and instead nudges emerging innovative founders back towards the Schumpeterian ambition of building their own 'private kingdom'.

For all these reasons, it becomes essential to delve deeper into the nature, mechanisms, and long-term implications of such unintended consequences. At the same time, the inherent difficulty of preventing them – especially under legislative approaches more grounded in the *rule of law* than in the *rule of reason* – makes it preferable to design more ad hoc regulatory instruments: probably narrower in scope, yet capable of reaching greater depth.

In this respect, a particularly promising example comes from the recent evolution of some antitrust authorities – towards more ex-ante forms of regulation – notably the UK's, through their 'market investigation' tool (see Motta et al., 2022 for a comprehensive analysis and geographical coverage of the instrument or similar ones). Unlike traditional antitrust enforcement, it does not require identifying a culprit: its purpose is not to punish past misconduct but to determine whether specific features of a market – rather than individual behaviour – are restricting or distorting competition. Under such tool, once reasonable grounds for concern are established, the antitrust opens a formal investigation, conducted by an independent panel of experts in law, economics, accounting, and business analysis (adequately supported by specialized staffs). This panel gathers evidence from firms, consumers, and other stakeholders, and is tasked with assessing whether specific characteristics of a selected market generate an adverse effect on competition and, in that case, with proposing the most appropriate remedies. Owing to its interdisciplinary composition and analytical mandate, the panel should in principle be particularly well-equipped to weigh the potential unintended consequences of the measures it

recommends. Focused on well-defined market segments – as illustrated, for instance, by the UK's use of this tool in the cloud services market (CMA, 2025) – the process is designed to enable the formulation of precise and proportionate remedies, whether structural (e.g. divestitures) or behavioural (e.g. interoperability). Of course, the devil is in the details – especially when it comes to translating policies that look effective on paper into ones that truly are in practice (Duflo, 2017). Nevertheless, it seems to us that, at least in principle, this tool displays several promising features. Its intrinsic combination of autonomy, depth and granularity of analysis, and contextual sensitivity – stemming precisely from the heterogeneity of expertise represented within the panel – should make it a valuable model for beginning to erode the 'Systemic Lock-in Model' of AI giants, while minimizing the risk of potential unintended consequences – including, though not limited to, those affecting the realm of innovative start-ups.

3° Policy Proposition: Encourage Entrepreneurship-Oriented Research on UBI and Related Safety-Net Programmes

AI's impact on labour markets is expected to be profound and disruptive, as it will both destroy and create jobs, but the transition is likely to involve significant *technological unemployment*, at least in the short term and potentially lasting longer (Acemoglu and Restrepo, 2019). This has clearly fuelled the debate on the need for labour market safety-net programmes capable of accompanying these transformations and mitigating their social impact. It has also rekindled the discussion around more radical policy schemes, epitomized by the idea of a universal basic income (UBI), succinctly described by Furman and Seamans (2019, p. 180) as follows:

> A universal basic income has three characteristics. First, it is available nearly universally, limited only by criteria such as citizenship or possibly age, but generally not by income or by other factors like disability or employment status. Second, it provides cash, not in- kind benefits like many existing programs that target food, heating oil, housing, and the like. Third, it is unconditional, so unlike many existing programs does not require the applicant to be working, looking for work, attending school, or other forms of conditioning. Most UBI proposals are also intended to be substantial enough to raise people to something like the poverty line. (Furman and Seamans, 2019, p. 180)

The literature highlights both the pros and cons of this and other, less radical (or rather, less unconditional) schemes. It is beyond the scope of this work to review them all, to discuss the feasibility of one model over another, or to engage with their moral or ethical implications. Instead, our aim is to emphasize that, in this area, from a purely economic perspective, policy makers have not, in our view, been sufficiently supported by researchers working on entrepreneurship-related topics. To clarify our point, when looking specifically at UBI, our domain of interest – innovative entrepreneurship – is often associated with both one of the key advantages and one of the key disadvantages of the measure. On the one hand, it is often argued that having a basic income could unleash creativity and entrepreneurial activity, acting on both the intensive and extensive margins: encouraging greater risk-taking among prospective entrepreneurs and convincing more individuals to pursue innovative ventures, reassured by the knowledge that, in case of failure, a basic income would still provide them with a safety net. This view is often countered by a first, well-grounded critique rooted in the literature on entrepreneurial finance, which warns that this very increase in risk-taking may prompt lenders and investors to become more cautious, potentially reducing their willingness to provide funds (Furman and Seamans, 2019). The more substantial criticism (to some extent connected with the previous one),[12] however, is that a UBI could risk turning citizens not into entrepreneurs but rather into mere (small-scale) consumers, thereby 'crystallizing' society and the capitalist system itself, and ultimately undermining the inherent dynamism that entrepreneurship brings as a social elevator (e.g. Kearney and Mogstad, 2019). The expanding ranks of freelancers in an economic system increasingly shaped by the 'gig economy' clearly amplify the risk of this latter drawback. The point we wish to stress, however, and the reason we call for more substantial scholarly engagement, is that research has so far failed to address this topic adequately, offering limited analytical and policy guidance. Admittedly, part of this gap can be explained by a classic chicken-and-egg paradox: it is difficult to investigate a phenomenon that has not yet been fully implemented. Yet, conversely, introducing such a measure would require a solid understanding of its potential implications across different contexts.

[12] In a closed economic system, greater consumption levels translate into lower savings, which reduce the supply of capital and raise interest rates (e.g. Hawal and Hashmi, 2025).

There are, of course, studies on a few localized 'experiments' involving UBI-like measures (e.g. Feinberg and Kuehn, 2020, on the Alaska Permanent Fund Dividend programme), as well as related but not fully central strands of literature that may offer some insights – for instance, research on lottery winnings and entrepreneurship (e.g. Golosov et al., 2024). However, an analysis of Scopus data reveals a striking imbalance: as of late 2025, while just under 1,400 documents include the term 'universal basic income' in their title, abstract, or keywords, the term 'entrepreneur' (including all derivations and related expressions such as 'start-up' or 'venture') appears in fewer than one hundred, and fewer than half of these fall within the fields of economics or business, with contributions from leading journals in entrepreneurship or innovation studies being virtually absent. Given that the 'game' in this and other related measures is largely played within the field of innovative entrepreneurship, a more focused and targeted research effort by the scientific community could substantially strengthen the evidence base available to policy makers, making it less vulnerable to ideological biases or preconceived narratives that too often cloud the debate on the subject. And, clearly, in the absence of a spontaneous initiative from the academic community, policy makers themselves have at their disposal a range of instruments, from research grants to targeted funding programmes, to stimulate greater engagement in this direction.

4° Policy Proposition: Embed AI and Grand Societal Challenges at the Core of (Entrepreneurship) Education

AI and education represent a long-standing field of application and inquiry, with roots reaching back to the 1970s and 1980s (as noted by one of its seminal figures, Schank, 2006). It is therefore by no means a by-product of the recent 'deep learning' wave that we have identified as the 'AI revolution'. While a full historical reconstruction lies beyond our scope, it is evident that this latest technological turn has reinvigorated both the field itself and the broader debate surrounding it. And if there were only one other domain (among the many, of course) where the transformative potential of AI most clearly triggers change, it would undoubtedly be education. In this regard, we share the view of those scholars and observers who argue that the *AI revolution* should rhyme with, and indeed be accompanied by, a *revolution in education*. Indeed, as early as 2017, Harari emphasized education,

and the reform of educational systems, as indispensable conditions for ensuring that AI technologies evolve within societies in ways that genuinely serve, and ultimately advance, the well-being of humanity. Building on these premises, we can now turn to our specific area of interest. We believe that the relationship between AI, education, and (innovative) entrepreneurship can be characterized by two main and mutually reinforcing effects – one primarily concerning the demand for innovative goods, and the other the supply side.

As for the former, it is clear that the educational system (alongside other factors, naturally) plays a crucial role in shaping the informal institutions that underpin every society. Its action should therefore be viewed as the first and most fundamental bulwark for the demystification of AI, capable of countering both *algorithmic aversion* and, even more importantly, the risk of *a-criticality* (Larson et al., 2024), the latter being a precursor to the excessive *polarization* of modern societies and to the resulting *homogeneity* within polarized communities. Only through a genuine integration of this new technology into the educational system – in our view, from the earliest school cycles (see, in this regard, the Chinese experience already mentioned in Chapter 5), in both teaching methods and content – will it be possible to mitigate, if not eliminate, these risks. Indeed, there is already, albeit necessarily preliminary, evidence that LLMs in particular may foster greater critical thinking among students (see, for example, the case study by Xiao and Hu, 2026), precisely because their effective use requires the ability to keep questioning, to ask again after each answer, and to resist the temptation to settle for the first response received (i.e. through *Socratic questioning*).[13]

Moreover, the educational system should also be regarded – as also emphasized by several institutions (e.g. UNESCO, 2022) and scholars

[13] In this respect, we believe that the teaching of philosophy – especially its emphasis on questioning, reasoning and epistemic reflection – is becoming increasingly important in contemporary education. Empirical evidence from 'philosophy for children' programmes suggests that philosophical inquiry can foster precisely the kind of dialogical and critical mindset we are referring to (Daniel and Auriac, 2011). In our view, such dispositions become even more crucial in an AI-mediated educational environment: integrating philosophical instruction with AI literacy from the earliest school stages may help children develop the critical mindset required to counter AI *sycophancy* (Naddaf, 2025), interact thoughtfully with AI systems and foster the emergence of important forms of co-intelligence, as well as hybrid or symbiotic intelligence (Zhou et al., 2021; Jarrahi et al., 2022; Mollick, 2024).

(e.g. Rieckmann, 2012) – as a key lever for ensuring that individuals' social and community-oriented attitudes, as well as their cultural values, help to avert, on the demand side, the risks of the 'thin innovation' dynamics already discussed. If, indeed, we remain firmly convinced that within the capitalist system one cannot (or rather, should not) entirely depart from Friedman's famous dictum that 'the business of business is business', it is equally clear that the pursuit of profit is by no means necessarily at odds with the broader objective of fostering socially 'thicker' forms of innovation. On the contrary, profit can serve as a powerful engine for achieving societal goals – as exemplified by the race to develop vaccines during the recent pandemic – in a typical heterogony of ends between the individual and society, a notion that can ultimately be traced back to Adam Smith (1776) and his 'invisible hand'. This, however, presupposes that individuals genuinely aspire to, and are willing to pay for, *thick*, rather than *thin*, forms of innovation.

And clearly, the opportunity to pursue one form of innovation over the other must also be recognized, seized, and (why not) fuelled by the supply side, and in our specific context, by the AI entrepreneur. To avert the risks of *postmodern entrepreneurship* and move instead towards a successful form of *thick* innovative entrepreneurship, education plays a pivotal role, especially through its integration with AI. First, ongoing research on entrepreneurship education suggests that AI tends to amplify, rather than weaken, the key characteristics that continue to distinguish (and, in our view, will continue to do so) successful entrepreneurial behaviour, as discussed in Chapter 4 (see in particular Figure 4.3). The use of AI technologies appears to magnify (or at the very least complement) the entrepreneurial mindset (i.e. in particular *self-efficacy*, see for example, among others, Xie and Wang, 2025), and can help *entrepreneurial savoir-faire* (e.g., in pitch training and emotional control by rehearsing with the aid of AI-generated self-avatars, see Lauer and Horvath, 2025). Second, although we believe that the development of the *dynamic AI-enabled orchestration capability* – already highlighted in Chapter 5 as a key element of a renewed competence-based view of the start-up – naturally benefits from practical experience, we also consider that an important part of this process (at least in terms of creating the appropriate 'fertile ground') comes from the integration of AI and entrepreneurship within specific study programmes. In our view, this does not mean

simply bringing the two domains together as separate silos under a single label, but rather aiming for a genuine integration between technical components and economic–managerial elements wherever possible.[14] Third, and certainly not least, as already noted in Chapter 4 – drawing on Graham's (2012) own words – entrepreneurial ideas should be grounded in 'problems', preferably ones that are personally experienced. (Entrepreneurship) Education therefore has the crucial task of ensuring that tomorrow's innovative entrepreneurs cultivate the boldness (and the *need for achievement*) required to confront significant and ambitious problems for themselves (Elia and Grilli, 2026), but also of ensuring that, at the same time, these problems at least point in the same direction as the major societal challenges affecting our world (Dyer et al., 2008).

Safeguarding Innovative Entrepreneurship in an Era of Industrial Policy Weaponization

This final point leads us to a broader consideration regarding public policy – one that extends well beyond the boundaries of entrepreneurship yet inevitably shapes, sometimes profoundly, the very dynamics that concern us. Insofar as the States' 'AI Wars' discussed in Chapter 7 are waged also through instruments of industrial policy (the so called risk of 'weaponization' of industrial policy, see Vannuccini, 2025) – and, by extension, entrepreneurship policy – it is evident that the introduction of regulation in one jurisdiction rather than another, or the deployment of a mission-oriented policy here rather than elsewhere, generates a competitive dynamic of winners and losers that resembles a zero-sum game. In this respect, there is widespread agreement that a regulatory approach to AI should strive, as much as possible, to adopt a geographically global perspective (e.g. Frazier and Krishna, 2025; Vannuccini, 2025), mirroring the logic underpinning cooperative efforts in other domains such as climate change. From our vantage point as scholars of entrepreneurship, we join this view.

[14] The examples in this sense can be numerous and naturally depend on the specific design of each course. One clear instance, however, could be the tight coupling between the development of concrete AI models and equally concrete entrepreneurial ideas, while explicitly taking into account issues of innovation appropriability and regulatory compliance across different geographical contexts.

If AI – and the suite of ethical, societal, and economic issues that accompany it – comes to represent a major societal challenge in its own right (rather than merely a tool, as implicitly assumed thus far, to mitigate pre-existing challenges), and therefore may require regulation, such regulation should be at least coordinated across countries. Only through a degree of international alignment can we preserve a level playing field for innovative entrepreneurial activity and avoid distortions that would otherwise arise from fragmented, nationally driven approaches.

Conclusion

Previous industrial revolutions have largely been characterized by SM I dynamics, that is, the fundamental process of 'stepping into the entrepreneurial arena' undertaken by individuals driven by their own knowledge. In this regard, it is impossible not to notice the striking assonance between the words of Josiah Wedgwood, one of the pioneers of the first industrial revolution, and those of Bill Gates, a leading figure of the third, despite the vast differences in their business contexts – eighteenth-century English ceramics for the former and software in the United States in the early 1970s for the latter:

> I saw the [industrial, author's note] field was spacious, and the soil so good, as to promise an ample recompense to any who should labour diligently in its cultivation. (Time, 1948)

> And I had a lot of uncertainty, even though the software thing, there were many experiences that made that just such a unique thing that, wow, why do I see this when other people don't see it. (Wired, 2025)

The primary purpose of our inquiry has been to explore whether, decades from now, looking back at the 'AI revolution', it will be equally possible to find excerpts such as those above. And, in the Weberian spirit evoked in our Introduction, the ambition of these pages has certainly not been to pronounce the last word on the matter, but to offer a policy-relevant map of the terrain – one that is useful to have today and can help orient present choices towards the societal good. In this respect, beyond more targeted policy recommendations, we humbly believe to have offered two (provisional) answers of a more general nature.

The first is that there is no reason to assume that this technology – however disruptive and general purpose it may be – has the capacity to undermine the essential function of the innovative entrepreneur within the capitalist system (of whatever variety). As we have seen, it alters – and will continue to alter – entrepreneurship, in some respects profoundly and even positively, yet without impoverishing its role.

Knowledge sources remain central to innovative entrepreneurship; the key 'pillars' of success may change in their composition and relative weight, yet they endure. So too does the importance of typical entrepreneurial psychological traits and creativity, which may even see their role magnified. Therefore, AI does not erode the very foundations of the social system of which it represents the Hamlet. In principle, to the extent that the recognition of any ('spacious') field is associated with a significant degree of (ontological) uncertainty, the earlier excerpts may still hold relevance in the AI era.

The second is that the magnitude of this 'space' – and the extent to which it is directed towards socially relevant grand challenges – depends less on the technology itself than on how institutions will be able to preserve it. It is in this sense that we have sought to elucidate why this is the case and to shed light on the vulnerabilities and challenges faced by the AI entrepreneur, by clarifying the role that policy makers can play in ensuring that AI-enabled innovative entrepreneurship is leveraged in ways that benefit society. We do not necessarily believe that capitalism is 'the end of history' (Fukuyama, 1992). At the same time, we do not see its end as a necessary consequence of AI. Capitalism, as we know, exists in different forms (Baumol et al., 2007; Bowles, 2013) and varieties (Hall and Soskice, 2001) and is 'accustomed' to changing shape over time (Grilli et al., 2024). Whether capitalism in the AI era is a 'good' capitalism – entrepreneurial and dynamic – ultimately depends on human agency. Indeed, thanks to AI, scientific progress and the potential innovative capacity of economic systems are expected to advance dramatically. At the same time, AI technologies are likely to become increasingly accessible, as the history of innovation diffusion suggests, not least through the proliferation of 'AI(d) offerings' (including those marketed by start-ups). In this scenario, alongside – and indeed in spite of – the growing strength of incumbents and AI giants, entrepreneurship must continue to play its role by generating *genuine* (i.e. 'thick') innovation rather than becoming a 'simulacrum' of itself, to borrow another familiar term from postmodernism (Baudrillard, 1981). Only under this condition, we argue, is it possible to preserve the very nature of (good) capitalism and maintain, as far as possible, its function as a social elevator in a world where the upper floors appear increasingly out of reach and the elevator itself seems progressively less capable of carrying people upward.

References

Abonamah, A. A., Tariq, M. U., & Shilbayeh, S. (2021). On the commoditization of artificial intelligence. *Frontiers in Psychology*, 12, 696346.

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.

Acs, Z. J., Braunerhjelm, P., Audretsch, D. B., & Carlsson, B. (2009). The knowledge spillover theory of entrepreneurship. *Small Business Economics*, 32(1), 15–30.

Acs, Z. J., Desai, S., & Hessels, J. (2008). Entrepreneurship, economic development and institutions. *Small Business Economics*, 31(3), 219–234.

Acs, Z. J., Lafuente, E., & Szerb, L. (eds.) (2023). *The Entrepreneurial Ecosystem: A Global Perspective*. Palgrave Macmillan, Cham.

Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. *Journal of Management*, 43(1), 39–58.

Agrawal, A., Gans, J., & Goldfarb, A. (2017). How AI will change strategy: A thought experiment. *Harvard Business Review*, 3 October 2017: https://hbr.org/2017/10/how-ai-will-change-strategy-a-thought-experiment (accessed 14 January 2026).

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press, Cambridge, MA.

Agarwal, N., Moehring, A., Rajpurkar, P., & Salz, T. (2023). Combining human expertise with artificial intelligence: Experimental evidence from radiology. National Bureau of Economic Research (Working Paper No. 31422).

Ahir, H., Bloom, N., & Furceri, D. (2022). The world uncertainty index. National Bureau of Economic Research (Working Paper No. 29763).

Ahlstrom, D., & Bruton, G. D. (2006). Venture capital in emerging economies: Networks and institutional change. *Entrepreneurship Theory and Practice*, 30(2), 299–320.

Ajjan, H. (2025). GenAI in Action: Practical Applications of AI Tools for Teaching and Learning. Harvard Business Publishing Webinar, 24 April 2025.

Alaimo, C., & Kallinikos, J. (2024). *Data Rules: Reinventing the Market Economy*. MIT Press, Cambridge, MA.

Alemany, L., & Andreoli, J. J. (eds.) (2018). *Entrepreneurial Finance: The Art and Science of Growing Ventures*. Cambridge University Press, Cambridge, UK.

Alexy, O., George, G., & Salter, A. J. (2013). Cui bono? The selective revealing of knowledge and its implications for innovative activity. *Academy of Management Review*, 38(2), 270–291.

Altman, S., & Moskovitz, D. (2014). *Lecture 1 – How to Start a Startup*. YC Root Access YouTube Channel: www.youtube.com/watch?v=CBYhVcO4WgI (accessed 20 September 2024).

Alvarez, S. A., Barney, J. B., & Anderson, P. (2013). Forming and exploiting opportunities: The implications of discovery and creation processes for entrepreneurial and organizational research. *Organization Science*, 24(1), 301–317.

Amabile, T. M. (2020). Creativity, artificial intelligence, and a world of surprises. *Academy of Management Discoveries*, 6(3), 351–354.

Ambrosini, V., & Bowman, C. (2001). Tacit knowledge: Some suggestions for operationalization. *Journal of Management Studies*, 38(6), 811–829.

Anderson, B. R., Shah, J. H., & Kreminski, M. (2024). Homogenization effects of large language models on human creative ideation. arXiv preprint arXiv:2402.01536.

Andreani, M., Ellahie, A., & Shivakumar, L. (2025). Are CEOs rewarded for luck? Evidence from corporate tax windfalls. *Journal of Finance*, 80(4), 2255–2302.

Andreoli, J. J. (2025). Venture capital investment selection. In D. Cumming and B. Hammer (eds.), *The Palgrave Encyclopedia of Private Equity* (pp. 1234–1242). Palgrave Macmillan, Cham.

Anelli, M., Colantone, I., Gallego, A., & Stanig, P. (2024). What if you see it? Workers' perceptions of and reactions to generative AI. Paper presented at the *European Workshop on the Economics of AI and Automation*, 25 October 2024, University of Milano-Bicocca, Milan.

Ardichvili, A., Cardozo, R., & Ray, S. (2003). A theory of entrepreneurial opportunity identification and development. *Journal of Business Venturing*, 18(1), 105–123.

Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of one, many: using language models to simulate human samples. *Political Analysis*, 31(3), 337–351.

Asam, D., & Heller, D. (2025). Generative AI and firm-level productivity: Evidence from startup funding dynamics. Paper presented at the *European Central Bank Conference, The Transformative Power of AI*, 2 April 2025, Frankfurt.

Ashley, M. (2025). The future is solo: AI is creating billion-dollar one-person companies. *Forbes*, 17 February 2025.

Audretsch, D. (1995). *Innovation and Industry Evolution*. MIT Press, Cambridge, MA.

Audretsch, D., Belitski, M., & Fiedler, A. (2025a). The knowledge spillover theory of entrepreneurship and innovation: Taking stock and new directions. *Journal of Technology Transfer*, forthcoming. https://doi.org/10.1007/s10961-025-10215-9.

Audretsch, D., Belitski, M., Fiedler, A., & Khlystova, O. (2025b). Five research themes of knowledge spillover theory of entrepreneurship: An analysis of 30 years of research. *International Journal of Management Reviews*, 27(4), 544–561.

Audretsch, D., Colombelli, A., Grilli, L., Minola, T., & Rasmussen, E. (2020). Innovative start-ups and policy initiatives. *Research Policy*, 49(10), 104027.

Autor, D. (2024). AI could actually help rebuild the middle class. *Noema*, 12 February 2024: www.noemamag.com/how-ai-could-help-rebuild-the-middle-class (accessed 27 February 2025).

Awal, F., & Hashmi, A. R. (2025). Universal basic income in an entrepreneurial economy. Mimeo. University of Calgary.

Azoulay, P., Jones, B. F., Kim, J. D., & Miranda, J. (2020). Age and high-growth entrepreneurship. *American Economic Review: Insights*, 2(1), 65–82.

Backes-Gellner, U., Werner, A., & Mohnen, A. (2015). Effort provision in entrepreneurial teams: Effects of team size, free-riding and peer pressure. *Journal of Business Economics*, 85, 205–230.

Bajari, P., Chernozhukov, V., Hortaçsu, A., & Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. *American Economic Association (AEA) Papers and Proceedings*, 109, 33–37.

Baker, T., & Nelson, R. E. (2005). Creating something from nothing: Resource construction through entrepreneurial bricolage. *Administrative Science Quarterly*, 50(3), 329–366.

Bandura, A. (1997). *Self-efficacy: The Exercise of Self Control*. Freeman, New York, NY.

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732.

Baron, R. A., & Ensley, M. D. (2006). Opportunity recognition as the detection of meaningful patterns: Evidence from comparisons of novice and experienced entrepreneurs. *Management Science*, 52(9), 1331–1344.

Baron Capital (2022). *Ron Baron Interviews Elon Musk at the 29th Annual Baron Investment Conference*. Baron Capital YouTube Channel: www.youtube.com/watch?v=E-squeb0YJA (accessed 2 January 2025).

Baron-Cohen, S. (2023). *The Pattern Seekers*. Basic Books, New York, NY.

Barringer, B. R., Jones, F. F., & Neubaum, D. O. (2005). A quantitative content analysis of the characteristics of rapid-growth firms and their founders. *Journal of Business Venturing*, 20(5), 663–687.

Barry, Z. (2023). In the age of AI, do you really need a cofounder? *Fast Company*, 2 August 2023: www.fastcompany.com/90931283/in-the-age-of-ai-do-you-really-need-a-cofounder (accessed 27 February 2025).

Baudrillard, J. (1981). *Simulacres et Simulation*. Éditions Galilée, Paris.

Baudrillard, J. (1998). *The Consumer Society: Myths & Structures*. Sage Publications, London.

Bauer, K., Heigl, R., Hinz, O., & Kosfeld, M. (2024). Feedback loops in machine learning: A study on the interplay of continuous updating and human discrimination. *Journal of the Association for Information Systems*, 25(4), 804–866.

Baumol, W. J. (1968). Entrepreneurship in economic theory. *American Economic Review*, 58(2), 64–71.

Baumol, W. J. (1990). Entrepreneurship: Productive, unproductive, and destructive. *Journal of Political Economy*, 98(5), 893–921.

Baumol, W. J., Litan, R. E., & Schramm, C. J. (2007). *Good Capitalism, Bad Capitalism, and the Economics of Growth and Prosperity*. Yale University Press, New Haven, CT.

Baumol, W. J., & Strom, R. J. (2007). Entrepreneurship and economic growth. *Strategic Entrepreneurship Journal*, 1(3–4), 233–237.

BBC (2024). *'Godfather of AI' on AI "exceeding human intelligence" and it "trying to take over"*. BBC Newsnight YouTube Channel: https://youtu.be/MGJpR591oaM?si=xsr7MMi0RmCLn8kq (accessed 27 January 2025).

BBC (2026). *AI-trained Domestic Robots: From Science Fiction to Reality?* BBC World Service YouTube Channel: www.youtube.com/watch?v=ZOsm1qoEr7M (accessed 25 January 2026).

Becker, G. S. (1975). *Human Capital*. National Bureau of Economic Research, New York, NY (US).

Ben-David, I., Johnson, M. J., & Stulz, R. M. (2025). Models behaving badly: The limits of data-driven lending. *Review of Finance*, 29(3), 711–745.

Bennetot, A., Franchi, G., Del Ser, J., Chatila, R., & Diaz-Rodriguez, N. (2022). Greybox XAI: A Neural-Symbolic learning framework to produce interpretable predictions for image classification. *Knowledge-Based Systems*, 258, 109947.

Bertoni, F., Colombo, M. G., & Grilli, L. (2011). Venture capital financing and the growth of high-tech start-ups: Disentangling treatment from selection effects. *Research Policy*, 40(7), 1028–1043.

Bertoni, F., Grilli, L., & Marzano, R. (2025). Antitrust enforcement and new firm dynamics. SSRN working paper: http://doi.org/10.2139/ssrn.5309461.

Bessen, J., Impink, S. M., Reichensperger, L., & Seamans, R. (2020). GDPR and the importance of data to AI startups. NYU Stern School of Business Working Paper, SSRN: http://doi.org/10.2139/ssrn.3576714.

Bessen, J., Impink, S. M., Reichensperger, L., & Seamans, R. (2022). The role of data for AI startup growth. *Research Policy*, 51(5), 104513.

Bessen, J., Pöge, F., & Röttger, R. (2023). Competing for talent: Large firms and startup growth. Bocconi University Management Research Paper, SSRN: http://doi.org/10.2139/ssrn.4673494.

Beugelsdijk, S. (2007). Entrepreneurial culture, regional innovativeness and economic growth. *Journal of Evolutionary Economics*, 17(2), 187–210.

Biancalani, F., Czarnitzki, D., & Riccaboni, M. (2022). The Italian start up act: A microeconometric program evaluation. *Small Business Economics*, 58(3), 1699–1720.

Bick, A., Blandin, A., & Deming, D. J. (2024). The rapid adoption of generative AI. National Bureau of Economic Research (Working Paper No. 32966).

Bilgram, V., & Laarmann, F. (2023). Accelerating innovation with generative AI: AI-augmented digital prototyping and innovation methods. *IEEE Engineering Management Review*, 51(2), 18–25.

Bittner A., & Lau, B. (2021). Women-led startups received just 2.3% of VC funding in 2020. *Harvard Business Review*, 25 February 2021: https://hbr.org/2021/02/women-led-startups-received-just-2-3-of-vc-funding-in-2020 (accessed 14 January 2026).

Blank, S. (2023). Startups need both profound beliefs and customer input. *Entrepreneur & Innovation Exchange*, 21 August 2023: https://eiexchange.com/content/startups-need-both-profound-beliefs-and-customer-input (accessed 14 January 2026).

Blohm, I., Antretter, T., Sirén, C., Grichnik, D., & Wincent, J. (2022). It's a peoples game, isn't it?! A comparison between the investment returns of business angels and machine learning algorithms. *Entrepreneurship Theory and Practice*, 46(4), 1054–1091.

Boldrin, M., & Levine, D. (2002). The case against intellectual property. *American Economic Review*, 92(2), 209–212.

Bonelli, M. (2026). Data-driven investors. *Review of Financial Studies*, 39(7), 1909–1969.

Borgonovi, F., Calvino, F., Criscuolo, C., Nania, J., Nitschke, J., O'Kane, L., Samek, L., & Seitz, H. (2023). Emerging trends in AI skill demand across 14 OECD countries. *OECD Artificial Intelligence Papers*, no. 2.

Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, Oxford.

Botelho, T. L., Fehder, D., & Hochberg, Y. (2021). Innovation-driven entrepreneurship. National Bureau of Economic Research (Working Paper No. 28990).

Bottazzi, L., Da Rin, M., & Hellmann, T. (2008). Who are the active investors?: Evidence from venture capital. *Journal of Financial Economics*, 89(3), 488–512.

Bottazzi, L., Da Rin, M., & Hellmann, T. (2016). The importance of trust for investment: Evidence from venture capital. *Review of Financial Studies*, 29(9), 2283–2318.

Bourzac, K. (2024). Fixing AI's energy crisis. *Nature Outlook*, https://doi.org/10.1038/d41586-024-03408-z.

Bouschery, S. G., Blazevic, V., & Piller, F. T. (2023). Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management*, 40(2), 139–153.

Bowles, P. (2013). *Capitalism*. Routledge, New York, NY.

Bowles, S., & Carlin, W. (2021). Shrinking capitalism: Components of a new political economy paradigm. *Oxford Review of Economic Policy*, 37(4), 794–810.

Bradley, S. W., Kim, P. H., Klein, P. G., McMullen, J. S., & Wennberg, K. (2021). Policy for innovative entrepreneurship: Institutions, interventions, and societal challenges. *Strategic Entrepreneurship Journal*, 15(2), 167–184.

Brekelmans, B. (2025). Announcing agentic automation with bidirectional integration between Microsoft Copilot Studio & UiPath. *Microsoft Monthly Updates*, 14 April 2025: https://bit.ly/3NcaHRx (accessed 14 January 2026).

Breucker P., Cointet J., Hannud Abdo A., Orsal G., de Quatrebarbes C., Duong T., Martinez C., et al. (2016). CorTexT Manager. Version v2: https://docs.cortext.net.

Brill, S. (2024). *The Death of Truth*. Knopf, New York, NY.

Bronkhorst, J., Schaveling, J., & Janssen, M. (2019). Commoditization and IT product innovation strategies from an IT firm perspective. *Information Systems Management*, 36(2), 126–140.

Brusset, X., La Torre, D., & Broekaert, J. (2022). Algorithms, analytics, and artificial intelligence: Harnessing data to make supply chain decisions. In B. MacCarthy and D. Ivanov (eds.), *The Digital Supply Chain* (pp. 93–110). Elsevier, Amsterdam.

Brynjolfsson, E., & Hitt, L. M. (1998). Beyond the productivity paradox. *Communications of the ACM*, 41(8), 49–55.

Brynjolfsson, E., & Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23–48.

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., et al. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv preprint arXiv:2303.12712.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12.

Burton, D. M., Dahl, M. S., & Sorenson, O. (2018). Do start-ups pay less? *ILR Review*, 71(5), 1179–1200.

Burton, J. W., Stein, M. K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239.

Cabral, L. (2021). Merger policy in digital industries. *Information Economics and Policy*, 54, 100866.

Camic, C. (2020). *Veblen: The Making of an Economist Who Unmade Economics*. Harvard University Press, Cambridge, MA.

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, 66(2), 564–586.

Cardon, M. S., & Stevens, C. E. (2004). Managing human resources in small organizations: What do we know? *Human Resource Management Review*, 14(3), 295–323.

Carr, N. G. (2003). IT doesn't matter. *Harvard Business Review*, 81(5), 41–49.

Carr, N. G. (2004). *Does IT Matter?* Harvard Business School Publishing, Boston, MA.

Carroll, N., Holmström, J., & Matook, S. (2024). Special issue editorial: Transforming business with low-code and no-code. *MIS Quarterly Executive*, 23(3), v–xiii.

Castelo, N., & Ward, A. F. (2021). Conservatism predicts aversion to consequential Artificial Intelligence. *PLoS ONE*, 16(12), e0261467.

Castelvecchi, D. (2024). Huge randomized trial of AI boosts discovery – At least for good scientists. *Nature*, 636, 286–287.

Celi, L. A. (2025). Teaching machines to doubt. *Nature Medicine*, 31, 3964.

Center for AI and Digital Policy (CAIDP) (2025). Artificial intelligence and democratic values 2025. Report by CAIDP: www.caidp.org/reports/aidv-2025/ (accessed 8 July 2025).

Chai, A., Stepanova, E., & Moneta, A. (2023). Quantifying expenditure hierarchies and the expansion of global consumption diversity. *Journal of Economic Behavior & Organization*, 214, 860–886.

Chalmers, D., MacKenzie, N. G., & Carter, S. (2021). Artificial intelligence and entrepreneurship: Implications for venture creation in the fourth industrial revolution. *Entrepreneurship Theory and Practice*, 45(5), 1028–1053.

Chaney, A. J., Stewart, B. M., & Engelhardt, B. E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (pp. 224–232). Association for Computing Machinery, New York, NY.

Chattopadhyay, S., Honoré, F., & Won, S. (2025). Free range startups? Market scope, academic founders, and the role of general knowledge in AI. *Strategic Management Journal*, 46(4), 1027–1079.

Chen, V. T., & Goldstein, J. (2024). The theatre of entrepreneurship: Learning to perform the speculative self in university entrepreneurship programs. *Social Problems*, 71(1), 203–219.

Chen, X. J. (2021). How AI is transforming venture capital. *BRINK*, 14 June 2021: www.brinknews.com/how-ai-is-transforming-venture-capital/ (accessed 11 June 2025).

Choi, S., Kang, H., Kim, N., & Kim, J. (2022). How does artificial intelligence improve human decision-making? Evidence from the AI-powered Go program. Paper Presented at the *41st Annual Strategic Management Conference*, Toronto.

Chouard, T. (2016). The Go Files: AI computer clinches victory against Go champion. *Nature*, https://doi.org/10.1038/nature.2016.19553.

Cirillo, V., Durand, C., Guarascio, D., Rabinovich, J., & Rikap, C. (2025). Power, knowledge and technology in a finite world. *Review of Political Economy*, 37(4), 1467–1478.

Clark, A. (1996). *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.

Clark, A. (2025). Extending minds with generative AI. *Nature Communications*, 16(1), 1–4.

CMA (Competition & Markets Authority) (2025). Cloud infrastructure services. Final Decision Report (31 July 2025): https://assets.publishing.service.gov.uk/media/688b8891fdde2b8f73469544/final_decision_report.pdf (accessed 20 October 2025).

Coad, A., Frankish, J., Roberts, R. G., & Storey, D. J. (2013). Growth paths and survival chances: An application of Gambler's Ruin theory. *Journal of Business Venturing*, 28(5), 615–632.

Cochoy, F., Licoppe, C., McIntyre, M. P., & Sörum, N. (2020). Digitalizing consumer society: Equipment and devices of digital consumption. *Journal of Cultural Economy*, 13(1), 1–11.

Cockburn, I. M., Henderson, R., & Stern, S. (2019). The impact of artificial intelligence on innovation. In A. Agrawal, J. Gans and A. Goldfarb (eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 115–146). University of Chicago Press, Chicago, IL.

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.

Cohen, W. M., Nelson, R., & Walsh, J. P. (2000). Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not). National Bureau of Economic Research (Working Paper No. 7552).

Collins, C., Dennehy, D., Conboy, K., & Mikalef, P. (2021). Artificial intelligence in information systems research: A systematic literature review and research agenda. *International Journal of Information Management*, 60, 102383.

Collins, H. (2007). Bicycling on the moon: Collective tacit knowledge and somatic-limit tacit knowledge. *Organization Studies*, 28(2), 257–262.

Collins, H. (2010). *Tacit and Explicit Knowledge*. University of Chicago Press, Chicago, IL.

Colombelli, A., D'Amico, E., & Paolucci, E. (2023). When computer science is not enough: Universities knowledge specializations behind artificial intelligence startups in Italy. *Journal of Technology Transfer*, 48(5), 1599–1627.

Colombo, M. G., & Grilli, L. (2005). Founders' human capital and the growth of new technology-based firms: A competence-based view. *Research Policy*, 34(6), 795–816.

Colombo, M. G., & Grilli, L. (2010). On growth drivers of high-tech start-ups: Exploring the role of founders' human capital and venture capital. *Journal of Business Venturing*, 25(6), 610–626.

Colombo, M. G., Grilli, L., Murtinu, S., Piscitello, L., & Piva, E. (2009). Effects of international R&D alliances on performance of high-tech start-ups: A longitudinal analysis. *Strategic Entrepreneurship Journal*, 3(4), 346–368.

Colombo, M. G., Grilli, L., Piscitello, L., & Rossi-Lamastra, C. (2011). Introduction. In M. G. Colombo, L. Grilli, L. Piscitello and C. Rossi-Lamastra (eds.), *Science and Innovation Policy for the New Knowledge Economy* (pp. 1–15). Edward Elgar, Cheltenham, UK.

Colombo, M. G., & Shafi, K. (2021). Receiving external equity following successfully crowdfunded technological projects: An informational mechanism. *Small Business Economics*, 56(4), 1507–1529.

Colombo, S., & Grilli, L. (2017). Should I stay or should I go? Founder's decision to leave an entrepreneurial venture during an industrial crisis. *Industry and Innovation*, 24(2), 97–121.

Confino, P. (2024). Could AI create a one-person unicorn? Sam Altman thinks so – and Silicon Valley sees the technology 'waiting for us'.

Fortune, 4 February 2024: https://fortune.com/2024/02/04/sam-altman-one-person-unicorn-silicon-valley-founder-myth/ (accessed 24 July 2025).
Conti, A., & Messinese, D. (2025). The selective tailwind effect of artificial intelligence in entrepreneurship. SSRN working paper: http://doi.org/10.2139/ssrn.4958898.
Cooper, A. C., & Bruno, A. V. (1977). Success among high-technology firms. *Business Horizons*, 20(2), 16–22.
Cooper, A. C., Gimeno-Gascon, F. J., & Woo, C. Y. (1994). Initial human and financial capital as predictors of new venture performance. *Journal of Business Venturing*, 9(5), 371–395.
Cropley, A. (2006). In praise of convergent thinking. *Creativity Research Journal*, 18(3), 391–404.
Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. Harper Perennial, New York, NY.
Csikszentmihalyi, M. (1997). *Finding Flow: The Psychology of Engagement with Everyday Life*. Basic Books, New York, NY.
Cunningham, C., Ederer, F., & Ma, S. (2021). Killer acquisitions. *Journal of Political Economy*, 129(3), 649–702.
D'Alessandro, F., Vivarelli, M., & Santarelli, E. (2026). The KSTE + I approach and the advent of AI technologies: Evidence from European regions. *Journal of Technology Transfer*, 51(1), 414–451.
Daniel, M. F., & Auriac, E. (2011). Philosophy, critical thinking and philosophy for children. *Educational Philosophy and Theory*, 43(5), 415–435.
Danneels, E., Clarysse, B., & De Cock, R. (2024). A tale of two ventures: How founder experiences shape search in tech-anchored and market-anchored ventures. *Strategic Organization*, forthcoming. https://doi.org/10.1177/14761270241293096.
Darby, M. R. & Zucker, L. G. (2005). Grilichesian breakthroughs: Inventions of methods of inventing and firm entry in nanotechnology. *Annales d'Économie et de Statistique*, 79/80, 143–164.
Davidsson, P. (2023). Ditching discovery-creation for unified venture creation research. *Entrepreneurship Theory and Practice*, 47(2), 594–612.
Davidsson, P., & Honig, B. (2003). The role of social and human capital among nascent entrepreneurs. *Journal of Business Venturing*, 18(3), 301–331.
Davidsson, P., & Sufyan, M. (2023). What does AI think of AI as an external enabler (EE) of entrepreneurship? An assessment through and of the EE framework. *Journal of Business Venturing Insights*, 20, e00413.
de Jong, A., Song, M., & Song, L. Z. (2013). How lead founder personality affects new venture performance: The mediating role of team conflict. *Journal of Management*, 39(7), 1825–1854.
de Jong, J. P., & Marsili, O. (2015). The distribution of Schumpeterian and Kirznerian opportunities. *Small Business Economics*, 44, 19–35.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2016). Declining business dynamism: What we know and the way forward. *American Economic Review*, 106(5), 203–207.

Deeds, D. L., & Hill, C. W. (1996). Strategic alliances and the rate of new product development: An empirical study of entrepreneurial biotechnology firms. *Journal of Business Venturing*, 11(1), 41–55.

Dell, M. (with C. Fredman) (1999). *Direct from Dell: Strategies that Revolutionized an Industry*. HarperCollins, New York, NY.

Dell'Acqua, F., Ayoubi, C., Lifshitz-Assaf, H., Sadun, R., Mollick, E. R., Mollick, L., Han, Y., et al. (2025). The cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise. Harvard Business School Working Paper Series no. 25-043.

Diakanastasi, E., Karagiannaki, A., & Pramatari, K. (2018). Entrepreneurial team dynamics and new venture creation process: An exploratory study within a start-up incubator. *Sage Open*, 8(2), 1–17.

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290.

Duflo, E. (2017). The economist as plumber. *American Economic Review*, 107(5), 1–26.

Dushnitsky, G., & Stroube, B. K. (2021). Low-code entrepreneurship: Shopify and the alternative path to growth. *Journal of Business Venturing Insights*, 16, e00251.

Dyer, J. H., Gregersen, H. B., & Christensen, C. (2008). Entrepreneur behaviors, opportunity recognition, and the origins of innovative ventures. *Strategic Entrepreneurship Journal*, 2(4), 317–338.

Economidou, C., Grilli, L., Henrekson, M., & Sanders, M. (2018). Financial and institutional reforms for an entrepreneurial society. *Small Business Economics*, 51(3), 279–291.

Ederer, F., Seibel, R., & Simcoe, T. (2025). Digital (killer?) acquisitions. Mimeo. GitHub: https://florianederer.github.io/digital.pdf.

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., & Larson, J. (2025). From local to global: A graphRAG approach to query-focused summarization. arXiv preprint arXiv:2404.16130.

Eesley, C. (2016). Institutional barriers to growth: Entrepreneurship, human capital and institutional change. *Organization Science*, 27(5), 1290–1306.

Eisenstein, M. (2023). Peptone: Finding opportunity in disorder. *Nature Biotechnology*, https://doi.org/10.1038/d41587-023-00011-x.

Elia, S., & Grilli, L. (2026). Go and conquer the world? Entrepreneurs' international experience and the target of different types of foreign markets. *Industry and Innovation*, 33(2), 234–268.

Engel, E. (1856). Die physische beschaffenheit der militärpflichtigen bevölkerung im königreich sachsen. *Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministerium des Inneren*, 2(4–7), 6161.

Esen, T., Dahl, M. S., & Sorenson, O. (2023). Jockeys, horses or teams? The selection of startups by venture capitalists. *Journal of Business Venturing Insights*, 19, e00383.

European Commission (2024). *The Future of European Competitiveness*. Publication Office of the European Union, Luxembourg.

Ezrachi, A., & Stucke, M. E. (2022). *How Big Barons Smash Innovation: And How to Strike Back*. Harper Collins, New York, NY.

Fairbairn, D. R. (1982). The impact of information technology on employment. *Science and Public Policy*, 9(5), 236–239.

Fang, Y. H. (2019). An app a day keeps a customer connected: Explicating loyalty to brands and branded applications through the lens of affordance and service-dominant logic. *Information & Management*, 56(3), 377–391.

Farquhar, S., Kossen, J., Kuhn, L., & Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017), 625–630.

Faull, H. (2023). Airbnb's seed and Series A fundraising rounds. *FindVC*, 2 May 2023: https://findvc.co.uk/airbnbs-seed-and-series-a-fundraising-rounds/ (accessed 30 January 2024).

Feinberg, R. M., & Kuehn, D. (2020). Does a guaranteed basic income encourage entrepreneurship? Evidence from Alaska. *Review of Industrial Organization*, 57(3), 607–626.

Felin, T., & Holweg, M. (2024). Theory is all you need: AI, human cognition, and decision making. *Strategy Science*, 9(4), 346–371.

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.

Felten, E., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. SSRN working paper: http://doi.org/10.2139/ssrn.4414065.

Felzmann, H., Villaronga, E. F., Lutz, C., & Tamò-Larrieux, A. (2019). Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society*, 6(1), 1–14.

Fischer, T., & Sitkin, S. B. (2023). Leadership styles: A comprehensive assessment and way forward. *Academy of Management Annals*, 17(1), 331–372.

Fjelland, R. (2020). Why general artificial intelligence will not be realized. *Humanities and Social Sciences Communications*, 7(1), 1–9.

Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, 25(8–9), 909–928.

Floridi, L. (2024a). The ethics of artificial intelligence: Exacerbated problems, renewed problems, unprecedented problems. *American Philosophical Quarterly*, 61(4), 301–307.

Floridi, L. (2024b). Hypersuasion – On AI's persuasive power and how to deal with it. *Philosophy & Technology*, 37, article 64.

Foss, N. J., & Klein, P. G. (2020). Entrepreneurial opportunities: Who needs them? *Academy of Management Perspectives*, 34(3), 366–377.

Foss, N. J., Klein, P. G., Lien, L. B., Zellweger, T., & Zenger, T. (2021). Ownership competence. *Strategic Management Journal*, 42(2), 302–328.

Fossen, F. M., McLemore, T., & Sorgner, A. (2024). Artificial intelligence and entrepreneurship. *Foundations and Trends® in Entrepreneurship*, 20(8), 781–904.

Francis (2024). *Dilexit Nos*. Encyclical letter on the human and divine love of the heart of Jesus Christ. Vatican.

Frazier, K., & Krishna, D. (2025). The perils of premature AI regulation and the case for an AI moratorium. *National Law Review*, AI & The Law Newsletter, 9 July 2025: https://natlawreview.com/article/perils-premature-ai-regulation-and-case-ai-moratorium (accessed 20 October 2025).

Freiling, J., Gersch, M., & Goeke, C. (2008). On the path towards a competence-based theory of the firm. *Organization Studies*, 29(8–9), 1143–1164.

Friesenbichler, K. S., & Kügler, A. (2022). Servitization across countries and sectors: Evidence from world input-output data. *Economic Systems*, 46(3), 101014.

Frost, A. K., Frost, S. L., & Johnsen, C. G. (2025). Status entrepreneurship: The entrepreneurial pursuit of social distinction. *Journal of Business Venturing Insights*, 24, e00563.

Fu, Z., Niu, X., & Maher, M. L. (2023). Deep learning models for serendipity recommendations: A survey and new perspectives. *ACM Computing Surveys*, 56(1), article 19.

Fukuyama, F. (1992). *The End of History and the Last Man*. Free Press, New York, NY.

Fulcher, J. (2004). *Capitalism: A Very Short Introduction*. Oxford University Press, New York, NY.

Furman, J., & Seamans, R. (2019). AI and the Economy. *Innovation Policy and the Economy*, 19(1), 161–191.

Ganotakis, P. (2012). Founders' human capital and the performance of UK new technology based firms. *Small Business Economics*, 39, 495–515.

Gans, J. S. (2025a). *The Microeconomics of Artificial Intelligence*. MIT Press, Cambridge, MA.

Gans, J. S. (2025b). Working within reason. *Mess and Magic*, 8 June 2025: https://joshuagans.substack.com/s/mess-and-magic (accessed 8 June 2025).

Gans, J. S., Stern, S., & Wu, J. (2019). Foundations of entrepreneurial strategy. *Strategic Management Journal*, 40(5), 736–756.

Gao, B., Wang, Y., Xie, H., Hu, Y., & Hu, Y. (2023). Artificial intelligence in advertising: Advancements, challenges, and ethical considerations in targeting, personalization, content creation, and ad optimization. *Sage Open*, 13(4), 1–20.

Gao, T., Wang, S., Chen, B., & Yang, L. (2024). The impact of big tech corporate venture capital investments on innovation: Evidence from the equity investment market. *China Economic Review*, 83, 102111.

Garicano, L. (2025). The AI Becker problem: Who will train the next generation? *Silicon Continent*, 23 July 2025: www.siliconcontinent.com/p/the-ai-becker-problem (accessed 4 August 2025).

Gautier, A., & Maitry, R. (2024). Big tech acquisitions and product discontinuation. *Journal of Competition Law & Economics*, 20(3), 246–263.

Gimmon, E., & Levie, J. (2010). Founder's human capital, external investment, and the survival of new high-technology ventures. *Research Policy*, 39(9), 1214–1226.

Giraudo, E., Giudici, G., & Grilli, L. (2019). Entrepreneurship policy and the financing of young innovative companies: Evidence from the Italian Startup Act. *Research Policy*, 48(9), 103801.

Girotra, K., Meincke, L., Terwiesch, C., & Ulrich, K. T. (2023). Ideas are dimes a dozen: Large language models for idea generation in innovation. SSRN working paper: http://doi.org/10.2139/ssrn.4526071.

Gjerløv-Juel, P., & Guenther, C. (2019). Early employment expansion and long-run survival: Examining employee turnover as a context factor. *Journal of Business Venturing*, 34(1), 80–102.

Gofman, M., & Jin, Z. (2024). Artificial intelligence, education, and entrepreneurship. *Journal of Finance*, 79(1), 631–667.

Goleman, D. (1995). *Emotional Intelligence*. Bantam Books, New York, NY.

Golosov, M., Graber, M., Mogstad, M., & Novgorodsky, D. (2024). How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *Quarterly Journal of Economics*, 139(2), 1321–1395.

Gompers, P., & Lerner, J. (2001). The venture capital revolution. *Journal of Economic Perspectives*, 15(2), 145–168.

Gompers, P. A., Gornall, W., Kaplan, S. N., & Strebulaev, I. A. (2020). How do venture capitalists make decisions? *Journal of Financial Economics*, 135(1), 169–190.

Gong, Z., Zhong, P., & Hu, W. (2019). Diversity in machine learning. *IEEE Access*, 7, 64323–64350.

Goodfellow, I., Bengio, Y., Courville, A. (2018). *Deep Learning*. MIT Press, Cambridge, MA.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680.

Google (2026). Joint statement from Google and Apple. *Google Company Announcement*, 12 January 2026: https://blog.google/company-news/inside-google/company-announcements/joint-statement-google-apple/ (accessed 13 January 2026).

Gosmar, D., & Dahl, D. A. (2025). Hallucination mitigation using agentic AI natural language-based frameworks. arXiv preprint arXiv: 2501.13946.

Goulart, G. D. S., Stocchi, L., & Porto, R. B. (2026). Understanding market dynamics and performance of mobile apps. *Australasian Marketing Journal*, 34(1), 15–27.

Goya-Martinez, M. (2016). The emulation of emotions in artificial intelligence: Another step into anthropomorphism. In S. Tettegah and S. U. Noble (eds.), *Emotions, Technology, and Design* (pp. 171–186). Academic Press (Elsevier), Cambridge, MA.

Graham, P. (2012). How to get startup ideas. Personal blog: www.paulgraham.com (accessed 4 January 2025).

Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360–1380.

Gray, S. M., Howell, T., & Sackett, E. (2024). Talking past each other: Construal level, utilitarian motives, and entrepreneurial team formation. *Organization Science*, 35(5), 1745–1769.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4), 501–522.

Grilli, L. (2014). High-tech entrepreneurship in Europe: A heuristic firm growth model and three '(un-) easy pieces' for policy-making. *Industry and Innovation*, 21(4), 267–284.

Grilli, L. (2022). Entrepreneurship and new product development: Exploring the 'advantage of youth' and 'business acumen' views. *Journal of Product Innovation Management*, 39(5), 662–685.

Grilli, L. (2025). Institutional determinants of venture capital activity. In D. Cumming and B. Hammer (eds.), *The Palgrave Encyclopedia of Private Equity* (pp. 600–603). Palgrave Macmillan, Cham.

Grilli, L., Jensen, P. H., Murtinu, S., & Park, H. D. (2020). A close look at the contingencies of founders' effect on venture performance. *Industrial and Corporate Change*, 29(4), 997–1020.

Grilli, L., Latifi, G., & Mrkajic, B. (2019). Institutional determinants of venture capital activity: An empirically driven literature review and a research agenda. *Journal of Economic Surveys*, 33(4), 1094–1122.

Grilli, L., Mariotti, S., & Marzano, R. (2024). Artificial intelligence and shape-shifting capitalism. *Journal of Evolutionary Economics*, 34(2), 303–318.

Grilli, L., Mrkajic, B., & Giraudo, E. (2023). Industrial policy, innovative entrepreneurship, and the human capital of founders. *Small Business Economics*, 60(2), 707–728.

Grilli, L., & Pedota, M. (2024). Creativity and artificial intelligence: A multilevel perspective. *Creativity and Innovation Management*, 33(2), 234–247.

Grint, K. (2010). *Leadership: A Very Short Introduction*. Oxford University Press, Oxford.

Guerreiro, J., & Loureiro, S. M. C. (2023). I am attracted to my cool smart assistant! Analyzing attachment-aversion in AI-human relationships. *Journal of Business Research*, 161, 113863.

Guilford, J. P. (1984). Varieties of divergent production. *Journal of Creative Behavior*, 18(1), 1–10.

Gunkel, D., & Coghlan, S. (2025). Cut the crap: A critical response to 'ChatGPT is bullshit'. *Ethics and Information Technology*, 27(2), article 23.

Guriev, S., & Papaioannou, E. (2022). The political economy of populism. *Journal of Economic Literature*, 60(3), 753–832.

Guthrie, S. E. (2025). Anthropomorphism. *Encyclopædia Britannica*, 18 January 2025: www.britannica.com/topic/anthropomorphism (accessed 22 January 2025).

Hagiu, A., & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics*, 54(4), 638–667.

Hagiu, A., & Wright, J. (2025). Artificial intelligence and competition policy. *International Journal of Industrial Organization*, 103, 103134.

Hall, B. H. (2002). The financing of research and development. *Oxford Review of Economic Policy*, 18(1), 35–51.

Hall, P. A., & Soskice, D. (2001). An introduction to varieties of capitalism. In P. A. Hall and D. Soskice (eds.), *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage* (pp. 1–68). Oxford University Press, Oxford.

Hallonsten, O. (2023). *Empty Innovation: Causes and Consequences of Society's Obsession with Entrepreneurship and Growth*. Palgrave Macmillan, Cham.

Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2013). Who creates jobs? Small versus large versus young. *Review of Economics and Statistics*, 95(2), 347–361.

Hannigan, T. R., McCarthy, I. P., & Spicer, A. (2024). Beware of botshit: How to manage the epistemic risks of generative chatbots. *Business Horizons*, 67(5), 471–486.

Hanoch, Y. (2002). 'Neither an angel nor an ant': Emotion as an aid to bounded rationality. *Journal of Economic Psychology*, 23(1), 1–25.

Harari, Y. N. (2017). Reboot for the AI revolution. *Nature*, 550(7676), 324–327.

Harari, Y. N. (2024). *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Penguin Random House, New York, NY.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162, 1243–1248.

Hartmann, R. K., Spicer, A., & Krabbe, A. D. (2022). Toward an untrepreneurial economy? The entrepreneurship industry and the rise of the Veblenian entrepreneur. In R. N. Eberhart, M. Lounsbury and H. E. Aldrich (eds.), *Entrepreneurialism and Society: Consequences and Meanings* (pp. 19–49). Emerald Publishing Limited, Bingley.

Hassija, V., Chamola, V., Mahapatra, A., Singal, A., Goel, D., Huang, K., Scardapane, S., Spinelli, I., Mahmud, M., & Hussain, A. (2024). Interpreting black-box models: A review on explainable artificial intelligence. *Cognitive Computation*, 16(1), 45–74.

Hauser, O. P., Light, M., Shelmerdine, L., & Blumenau, J. (2025). Why evaluating the impact of AI needs to start now. *Nature*, 643, 910–912.

Hayek, F. A. (1944). *The Road to Serfdom*. Routledge & Kegan Paul, London.

Hellmann, T. (2000). Venture capitalists: The coaches of Silicon Valley. *The Silicon Valley Edge*, 113(6), 276–294.

Hellmann, T., Kazanci, K., & Sako, M. (2024). How AI is transforming venture building and venture capital, University of Oxford, Saïd Business School, *Oxford Answers*, 23 October 2024: www.sbs.ox.ac.uk/oxford-answers/how-ai-transforming-venture-building-and-venture-capital (accessed 20 June 2025).

Henrekson, M. (2007). Entrepreneurship and Institutions. IFN Working Paper No. 707, Research Institute of Industrial Economics (IFN).

Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), 1–10.

Hirsch, F. (1976). *Social Limits to Growth*. Harvard University Press, Cambridge, MA.

Hodgson, G. M. (1998). Evolutionary and competence-based theories of the firm. *Journal of Economic Studies*, 25(1), 25–56.

Hofstede, G. H., Hofstede, G. J., & Minkov, M. (2010). *Cultures and Organizations – Software of the Mind: Intercultural Cooperation and Its Importance for Survival*. McGraw-Hill, New York, NY.

Hojjati, H., Ho, T. K. K., & Armanfard, N. (2024). Self-supervised anomaly detection in computer vision and beyond: A survey and outlook. *Neural Networks*, 172(April), 106106.

Hoque, F., Davenport, T., & Nelson, E. (2025). Why AI demands a new breed of leaders. *MIT Sloan Management Review*, 9 April 2025: https://sloanreview.mit.edu/article/why-ai-demands-a-new-breed-of-leaders/ (accessed 22 July 2025).

Horowitz, B. (2025). *Keynote Speech at Columbia Business School*. Columbia Business School Youtube Channel: www.youtube.com/watch?v=JFUSTctUxGM (accessed 3 October 2025).

Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from homo silicus? National Bureau of Economic Research (Working Paper No. 31122).

Houston, M. (2025). How to build a passive income business for financial freedom. *Forbes*, 23 February 2025: https://bit.ly/4jLHbhK (accessed 24 August 2025).

Huang, J. (2025). *Keynote Speech at CES 2025*. NVIDIA Youtube Channel: www.youtube.com/watch?v=k82RwXqZHY8 (accessed 10 January 2025).

Hughes, L., Dwivedi, Y. K., Malik, T., Shawosh, M., Albashrawi, M. A., Jeon, I., Dutot, V., et al. (2025). AI agents and agentic systems: A multi-expert analysis. *Journal of Computer Information Systems*, 65(4), 489–517.

Hull, A. D., Liew, J. K. S., Palaoro, K. T., Grzegorzewski, M., Klipstein, M., Breuer, P., & Spencer, M. (2022). Why the United States must win the artificial intelligence (AI) race. *Cyber Defense Review*, 7(4), 143–158.

Hume, D. (1739, reprinted in 2000). *A Treatise of Human Nature*. Oxford University Press, New York, NY.

Hume, J. (2025). How AI-powered innovation can democratize breakthrough science. *World Economic Forum*, 25 June 2025: www.weforum.org/stories/2025/06/ai-innovation-democratizes-breakthrough-science/ (accessed 14 January 2026).

Hunt, R. A., & Kurdoglu, R. S. (2025). Font of innovation or algorithmic deforestation? The ecosystem impacts of artificial intelligence in entrepreneurship. *Journal of Business Venturing Insights*, 24, e00575.

Igami, M. (2020). Artificial intelligence as structural estimation: Deep Blue, Bonanza, and AlphaGo. *Econometrics Journal*, 23(3), S1–S24.

International Energy Agency (2024). *Electricity 2024*. IEA Publications, Paris.

Ioramashvili, C., Feldman, M., Guy, F., & Iammarino, S. (2024). Gathering round Big Tech: How the market for acquisitions concentrates the digital sector. *Cambridge Journal of Regions, Economy and Society*, 17(2), 293–306.

Isaac, M., Tan, E., & Metz, C. (2025). A.I. researchers are negotiating $250 million pay packages. Just like N.B.A. stars. *New York Times*, 31 July 2025: www.nytimes.com/2025/07/31/technology/ai-researchers-nba-stars.html (accessed 04 September 2026).

Jablonka, K. M., Schwaller, P., Ortega-Guerrero, A., & Smit, B. (2024). Leveraging large language models for predictive chemistry. *Nature Machine Intelligence*, 6(2), 161–169.

Jacobides, M. G., Brusoni, S., & Candelon, F. (2021). The evolutionary dynamics of the artificial intelligence ecosystem. *Strategy Science*, 6(4), 412–435.

Jacobides, M. G., Cennamo, C., & Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, 39(8), 2255–2276.

Jacobides, M. G., & Tae, C. J. (2015). Kingpins, bottlenecks, and value dynamics along a sector. *Organization Science*, 26(3), 889–907.

Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), 577–598.

Jameson, F. (1991). *Postmodernism, or the Cultural Logic of Late Capitalism*. Duke University Press, Durham, NC.

Janeway, W. H., Nanda, R., & Rhodes-Kropf, M. (2021). Venture capital booms and start-up financing. *Annual Review of Financial Economics*, 13(1), 111–127.

Jarrahi, M. H., Lutz, C., & Newlands, G. (2022). Artificial intelligence, human intelligence and hybrid intelligence based on mutual augmentation. *Big Data & Society*, 9(2), 1–6.

Jia, J., Jin, G. Z., & Wagman, L. (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, 40(4), 661–684.

Johnson, G. (2022). Economic research on privacy regulation: Lessons from the GDPR and beyond. National Bureau of Economic Research (Working Paper No. 30705).

Jones, R. (2017). The Internet is loving this Japanese artist's surreal pigeon shoes. *Yahoo FN*, 30 May 2017: https://bit.ly/4qTvuI3 (accessed 21 February 2025).

Jovanovic, B., & Rousseau, P. L. (2005). General purpose technologies. In P. Aghion and S. Durlauf (eds.), *Handbook of Economic Growth* (Vol. 1, pp. 1181–1224). Elsevier, Amsterdam.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, NY.

Kakutani, M. (2018). *The Death of Truth*. Tim Duggan Books, New York, NY.

Kamepalli, S. K., Rajan, R., & Zingales, L. (2020). Kill zone. National Bureau of Economic Research (Working Paper No. 27146).

Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., & Amodei, D. (2020). Scaling laws for neural language models. arXiv preprint arXiv:2001.08361.

Kaplan, S. N., Sensoy, B. A., & Strömberg, P. (2009). Should investors bet on the jockey or the horse? Evidence from the evolution of firms from early business plans to public companies. *Journal of Finance*, 64(1), 75–115.

Kato, B. (2025). World's first AI-designed shoe is 3D-printed to your foot's specifications: 'The best fit you can possibly get'. *New York Post*, 22 January 2025: https://bit.ly/4sNaYuv (accessed 14 January 2026).

Kaufman, J. C., & Glăveanu, V. P. (2021). An overview of creativity theories. In J. C. Kaufman and R. J. Sternberg (eds.), *Creativity: An Introduction* (pp. 17–30). Cambridge University Press, Cambridge.

Kaufman, J. C., & Sternberg, R. J. (eds.) (2021). *Creativity: An Introduction*. Cambridge University Press, Cambridge.

Kearney, M. S., & Mogstad, M. (2019). Universal basic income (UBI) as a policy response to current challenges. Report by Aspen Institute Economic Strategy Group: www.economicstrategygroup.org/publication/universal-basic-income-ubi-as-a-policy-response-to-current-challenges/ (accessed 28 November 2025).

Keating, G. (2014). Five myths about Netflix. *Washington Post*, 21 February 2014: https://bit.ly/4pIhsYT (accessed 10 December 2024).

Kerr, S. P., Kerr, W. R., & Xu, T. (2018). Personality traits of entrepreneurs: A review of recent literature. *Foundations and Trends® in Entrepreneurship*, 14(3), 279–356.

Keynes, J. M. (1930). Economic possibilities for our grandchildren. In J. M. Keynes, *Essays in Persuasion* (pp. 321–332). Palgrave Macmillan, London.

Khan, L. M. (2017). Amazon's antitrust paradox. *Yale Law Journal*, 126, 710–805.

Kier, A. S., & McMullen, J. S. (2018). Entrepreneurial imaginativeness in new venture ideation. *Academy of Management Journal*, 61(6), 2265–2295.

Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114.

Kirkpatrick, D. (2010). *The Facebook Effect: The Inside Story of the Company that Is Connecting the World*. Simon & Schuster, New York, NY.

Kirzner, I. M. (1973). *Competition and Entrepreneurship*. University of Chicago Press, Chicago.

Klein Teeselink, B. (2025). Generative AI and labor market outcomes: Evidence from the United Kingdom. SSRN working paper: http://doi.org/10.2139/ssrn.5516798.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Hart, Schaffner and Marx, Boston, MA.

Kocka, J. (2016). *Capitalism: A Short History*. Princeton University Press, Princeton, NJ.

Koehn, N. F. (1997). Josiah Wedgwood and the first industrial revolution. In T. K. McCraw (eds.), *Creating Modern Capitalism* (pp. 17–48). Harvard University Press, Boston, MA.

Koivisto, M., & Grassini, S. (2023). Best humans still outperform artificial intelligence in a creative divergent thinking task. *Scientific Reports*, 13(1), 13601.

Kolling, T., Baisch, S., Schall, A., Selic, S., Rühl, S., Kim, Z., Rossberg, S., et al. (2016). What is emotional about emotional robotics? In T. Tettegah and Y. E. Garcia (eds.), *Emotions, Technology, and Health* (pp. 85–103). Academic Press, San Diego, CA.

Korinek, A. (2024). Economic policy challenges for the age of AI. National Bureau of Economic Research (Working paper No. 32980).

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X. H., Beresnitzky, A. V., Braunstein, I. & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint arXiv:2506.08872.

Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44(6), 1425–1452.

Kranzberg, M. (1986). Technology and history: 'Kranzberg's Laws'. *Technology and Culture*, 27(3), 544–560.

Krarup, T., & Horst, M. (2023). European artificial intelligence policy as digital single market making. *Big Data & Society*, 10(1), 1–14.

Krause, A., & Hübotter, J. (2025). Probabilistic Artificial Intelligence. arXiv preprint arXiv:2502.05244.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1106–1114.

Kulveit, J., Douglas, R., Ammann, N., Turan, D., Krueger, D., & Duvenaud, D. (2025). Gradual disempowerment: Systemic existential risks from incremental AI development. arXiv preprint arXiv:2501.16946.

Kurzweil, R. (2024). *The Singularity Is Nearer*. Viking, New York, NY.

Lahart, J. (2024). Will AI help or hurt workers? One 26-year-old found an unexpected answer. *Wall Street Journal*, 29 December 2024: www.wsj.com/economy/will-ai-help-hurt-workers-income-productivity-5928a389 (accessed 12 July 2025).

Lahart, J. (2025). MIT says it no longer stands behind student's AI research paper. *Wall Street Journal*, 16 May 2025: www.wsj.com/tech/ai/mit-says-it-no-longer-stands-behind-students-ai-research-paper-11434092 (accessed 8 August 2025).

Lahoti, P., Blumm, N., Ma, X., Kotikalapudi, R., Potluri, S., Tan, Q., Srinivasan, H., Packer, B., Beirami, A., Beutel, A., & Chen, J. (2023). Improving diversity of demographic representation in large language

models via collective-critiques and self-voting. In H. Bouamor, J. Pino, and K. Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 10383–10405). Association for Computational Linguistics, Singapore. https://aclanthology.org/2023.emnlp-main.643/.

Lam, N. (2022). Explanations in AI as claims of tacit knowledge. *Minds and Machines*, 32(1), 135–158.

Larson, B. Z., Moser, C., Caza, A., Muehlfeld, K., & Colombo, L. A. (2024). Critical thinking in the age of generative AI. *Academy of Management Learning & Education*, 23(3), 373–378.

Larson, E. J. (2021). *The Myth of Artificial Intelligence*. Belknap Press, London.

Lauer, D., & Horvath, D. (2025). AI-enhanced pitching training: The impact of idealized future self-avatars. Paper presented at the *Interdisciplinary European Conference on Entrepreneurship Research (IECER)*, 12-14 September 2025, University of Reykjavík, Reykjavík.

Lazar, M., Miron-Spektor, E., Chen, G., Goldfarb, B., Erez, M., & Agarwal, R. (2022). Forming entrepreneurial teams: Mixing business and friendship to create transactive memory systems for enhanced success. *Academy of Management Journal*, 65(4), 1110–1138.

LBC (2025). '*Godfather of AI' Predicts It Will Take Over the World*. LBC YouTube Channel: www.youtube.com/watch?v=vxkBE23zDmQ&t=513s (accessed 3 December 2025).

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.

Lee, B. C., & Chung, J. (2024). An empirical investigation of the impact of ChatGPT on creativity. *Nature Human Behaviour*, 8(10), 1906–1914.

Legnani, M., Magistretti S., & Dell'Era, C. (2024). Augmenting design thinking: Unveiling the role of generative AI in innovation management. Paper presented at the *35° Scientific Annual Meeting of the Italian Association of Management Engineering (AiIG)*, 10–11 October 2024, Palermo.

LePage, L. (2021). The importance of realism, character, and genre: How theatre can support the creation of likeable sociable robots. *International Journal of Social Robotics*, 13(6), 1427–1441.

Lerner, J. (2002). When bureaucrats meet entrepreneurs: The design of effective 'public venture capital' programmes. *Economic Journal*, 112(477), F73–F84.

Lerner, J. (2009). *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed – and What to Do about It*. Princeton University Press, Princeton, NJ.

Lerner, J. (2020). Government incentives for entrepreneurship. National Bureau of Economic Research (Working Paper No. c14426).

Levine, S. S., & Jain, D. (2023). How network effects make AI smarter. *Harvard Business Review*, 14 March 2023: https://hbr.org/2023/03/how-network-effects-make-ai-smarter (accessed 15 January 2026).

Levy, J. (2017). Capital as process and the history of capitalism. *Business History Review*, 91(3), 483–510.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.

Li, D. (2026). *Recent Developments in AI and the Labor Market*. AEA Recent Developments Lectures, American Economic Association Webcast: www.aeaweb.org/webcasts/2026/developments-in-ai-and-labor (accessed 10 January 2026).

Li, L., Liu, Y., Jin, Y., Cheng, T. E., & Zhang, Q. (2024). Generative AI-enabled supply chain management: The critical role of coordination and dynamism. *International Journal of Production Economics*, 277, 109388.

Li, W., Li, Z., Li, W., Zhang, Y., & Li, A. (2025). Mapping the empirical literature of the GDPR's (In-)effectiveness: A systematic review. *Computer Law & Security Review*, 57, 106129.

Li, Y., & Zahra, S. A. (2012). Formal institutions, culture, and venture capital activity: A cross-country analysis. *Journal of Business Venturing*, 27(1), 95–111.

Liang, L., Zhuang, H., Zou, J., & Acuna, D. E. (2024). The complementary contributions of academia and industry to AI research. arXiv preprint arXiv:2401.10268.

Lichtinger, G., & Hosseini Maasoum, S. M. (2025). Generative AI as seniority-biased technological change: Evidence from US resume and job posting data. SSRN working paper: http://doi.org/10.2139/ssrn.5425555.

Lin, C. Y., Tseng, C. K., Teng, W. C., Lee, W. C., Kuo, C. H., Gu, H. Y., Chung, K. L., & Fahn, C. S. (2009). The realization of robot theater: Humanoid robots and theatric performance. 2009 *International Conference on Advanced Robotics Proceedings*, Munich, 2009, 1–6. https://ieeexplore.ieee.org/document/5174809/authors#authors.

Luchini, S., Kaufman, J., & Beaty, R. (2025). Generative AI does not erase individual differences in human creativity. PsyArXiv preprint: https://doi.org/10.31234/osf.io/jszrn_v1.

Lundvall, B. Å., & Rikap, C. (2022). China's catching-up in artificial intelligence seen as a co-evolution of corporate and national innovation systems. *Research Policy*, 51(1), 104395.

Lyotard, J.-F. (1984). *The Postmodern Condition: A Report on Knowledge*. Manchester University Press, Manchester.

Ma, L., Xu, X., He, Y., & Tan, Y. (2024). Learning to adopt generative AI. SSRN working paper: http://doi.org/10.2139/ssrn.4990170.

Maarouf, A., Feuerriegel, S., & Pröllochs, N. (2025). A fused large language model for predicting startup success. *European Journal of Operational Research*, 322(1), 198–214.

Makridakis, S. (2017). The forthcoming Artificial Intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90(1), 46–60.

Malerba, F., & Orsenigo, L. (1995). Schumpeterian patterns of innovation. *Cambridge Journal of Economics*, 19(1), 47–65.

Manaresi, F., Menon, C., & Santoleri, P. (2021). Supporting innovative entrepreneurship: An evaluation of the Italian 'Start-up Act'. *Industrial and Corporate Change*, 30(6), 1591–1614.

Matricano, D. (2024). Designing effective policies for innovative start-ups: Lessons learned in Italy. *Journal of Business Venturing Insights*, 22, e00486.

Mayer, J. D., Salovey, P., & Caruso, D. R. (2004). Emotional intelligence: Theory, findings, and Implications. *Psychological Inquiry*, 15(3), 197–215.

McAfee, A., & Brynjolfsson, E. (2012). Big data: The management revolution. *Harvard Business Review*, 90(10), 60–68.

McClelland, D. C. (1965). Toward a theory of motive acquisition. *American Psychologist*, 20(5), 321–333.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133.

McLean, S., Read, G. J., Thompson, J., Baber, C., Stanton, N. A., & Salmon, P. M. (2023). The risks associated with Artificial General Intelligence: A systematic review. *Journal of Experimental & Theoretical Artificial Intelligence*, 35(5), 649–663.

McMullen, J. S., Plummer, L. A., & Acs, Z. J. (2007). What is an entrepreneurial opportunity? *Small Business Economics*, 28(4), 273–283.

Mei, Q., Xie, Y., Yuan, W., & Jackson, M. O. (2024). A Turing test of whether AI chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences*, 121(9), e2313925121.

Melumad, S., & Yun, J. H. (2025). Experimental evidence of the effects of large language models versus Web search on depth of learning. SSRN working paper: http://doi.org/10.2139/ssrn.5104064.

Menke, M., Xu, Q., & Gu, L. (2007). An analysis of the universality, flexibility, and agility of total innovation management: A case study of Hewlett–Packard. *Journal of Technology Transfer*, 32, 49–62.

Metz, C. (2015). WhatsApp's co-founder on how the App became a phenomenon. *Wired*, 26 October 2015: www.wired.com/2015/10/whatsapps-co-founder-on-how-the-iconoclastic-app-got-huge/ (accessed 15 January 2026).

Metz, C. (2016). In two moves, AlphaGo and Lee Sedol redefined the future. *Wired*, 16 March 2016: www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/ (accessed 15 January 2026).

Meyer, B. (2011). John McCarthy. *Communications of the ACM*, 28 October 2011: https://cacm.acm.org/blogcacm/john-mccarthy/ (accessed 15 January 2026).

Michelacci, C., & Silva, O. (2007). Why so many local entrepreneurs? *Review of Economics and Statistics*, 89(4), 615–633.

Mihet, R., Rishabh, K., & Gomes, O. M. D. C. (2026). Is it AI or data that drives firm market power? *Journal of Monetary Economics*, 157, 103878.

Minh, D., Wang, H. X., Li, Y. F., & Nguyen, T. N. (2022). Explainable artificial intelligence: A comprehensive review. *Artificial Intelligence Review*, 55, 3503–3568.

Minsky, M. (2006). *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon and Schuster, New York, NY.

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. Penguin Random House, New York, NY.

Mitchell, M. (2021). Why AI is harder than we think. arXiv preprint arXiv:2104.12871.

Mitchell, M. (2025). Artificial Intelligence learns to reason. *Science*, 387(6740), eadw5211.

Mitchell, M., Ghosh, A., Luccioni, A. S., & Pistilli, G. (2025). Fully autonomous AI agents should not be developed. arXiv preprint arXiv:2502.02649.

Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing*, 29(1), 1–16.

Mollick, E. (2024). *Co-intelligence: Living and Working with AI*. Penguin Random House, New York, NY.

Moose, I. A. (2025). *The Economics of Artificial Intelligence*. Edward Elgar Publishing, Cheltenham.

Motta, M., Peitz, M., & Schweitzer, H. (eds.) (2022). *Market Investigations: A New Competition Tool for Europe?* Cambridge University Press, Cambridge.

Mowery, D., & Rosenberg, N. (1979). The influence of market demand upon innovation: A critical review of some recent empirical studies. *Research Policy*, 8(2), 102–153.

Müller, V. C., & Bostrom, N. (2016). Future progress in artificial intelligence: A survey of expert opinion. In V. C. Müller (eds.), *Fundamental Issues of Artificial Intelligence* (pp. 553–571, Synthese Library, vol 376). Springer, Berlin.

Murtinu, S., & De Massis, A. (2025). Artificial intelligence machines as relational nonhuman actors in entrepreneurial teams. *Journal of Small Business Management*, 63(6), 2890–2916.

My First Million (2024). *I got rejected from YC (4x)…. now my side hustle is worth $1.16B*. My First Million YouTube Channel, Episode No. 658: www.youtube.com/watch?v=TOQtJch3kGk (accessed 31 January 2025).

Naddaf, M. (2025). AI chatbots are sycophants – Researchers say it's harming science. *Nature*, 647, 13–14.

Nambisan, S., Siegel, D., & Kenney, M. (2018). On open innovation, platforms, and entrepreneurship. *Strategic Entrepreneurship Journal*, 12(3), 354–368.

Narayanan, A., & Kapoor, S. (2024). *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press, Princeton, NJ.

National Academies of Sciences, Engineering, and Medicine (2024). *Artificial Intelligence and the Future of Work*. The National Academies Press, Washington, DC.

Nature (2016). Digital intuition. *Nature*, 529, 437.

Nature (2025). Bring us your LLMs. *Nature*, 645, 559.

Nave, G., Meincke, L., & Terwiesch, C. (2024). ChatGPT decreases idea diversity in brainstorming. PsyArXiv preprint: https://doi.org/10.31234/osf.io/yqb29.

Nechushtai, E., Zamith, R., & Lewis, S. C. (2024). More of the same? Homogenization in news recommendations when users search on Google, YouTube, Facebook, and Twitter. *Mass Communication and Society*, 27(6), 1309–1335.

Neiman, B., & Vavra, J. (2023). The rise of niche consumption. *American Economic Journal: Macroeconomics*, 15(3), 224–264.

Neirotti, P., & Paolucci, E. (2007). Assessing the strategic value of information technology: An analysis on the insurance sector. *Information & Management*, 44(6), 568–582.

Nicholas, T. (2019). *VC: An American History*. Harvard University Press, Cambridge, MA.

Nicolaou, N., & Shane, S. (2009). Can genetic factors influence the likelihood of engaging in entrepreneurial activity? *Journal of Business Venturing*, 24(1), 1–22.

Nikkhoo, I. (2025). 3 ways AI is transforming venture capital investment. *Crunchbase News*, 17 April 2025: https://news.crunchbase.com/ai/transforming-vc-investment-nikkhoo-navigate/ (accessed 15 January 2026).

Nikolenko, S. I. (2021). *Synthetic Data for Deep Learning*. Springer Nature, London.

Norbäck, P. J., & Persson, L. (2023). Why big data can make creative destruction more creative – But less destructive. IFN Working Paper No. 1454, Research Institute of Industrial Economics (IFN).

North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge.

North, D. C. (1991). Institutions. *Journal of Economic Perspectives*, 5(1), 97–112.

North, D. C. (1994). Economic performance through time. *American Economic Review*, 84(3), 359–368.

O'Donnell, P., Leger, M., O'Gorman, C., & Clinton, E. (2024). Necessity entrepreneurship. *Academy of Management Annals*, 18(1), 44–81.

Obschonka, M., Grégoire, D. A., Nikolaev, B., Ooms, F., Lévesque, M., Pollack, J. M., & Behrend, T. S. (2025). Artificial intelligence and entrepreneurship: A call for research to prospect and establish the scholarly AI frontiers. *Entrepreneurship Theory and Practice*, 49(3), 620–641.

OECD (2025). Intellectual property issues in artificial intelligence trained on scraped data. *OECD Artificial Intelligence Papers*, no. 33.

Orwig, W., Bellaiche, L., Spooner, S., Vo, A., Baig, Z., Ragnhildstveit, A., Schacter, D. L., Barr, N., & Seli, P. (2026). Using AI to generate visual art: Do individual differences in creativity predict AI-assisted art quality? *Creativity Research Journal*, 38(1), 242–253.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge.

Otis, N., Clarke, R., Delecourt, S., Holtz, D., & Koning, R. (2024). The uneven impact of generative AI on entrepreneurial performance. Harvard Business School Working Paper Series no. 24-042.

Packard, M. D., & Clark, B. B. (2020). On the mitigability of uncertainty and the choice between predictive and nonpredictive strategy. *Academy of Management Review*, 45(4), 766–786.

Packard, M. D., Clark, B. B., & Klein, P. G. (2017). Uncertainty types and transitions in the entrepreneurial process. *Organization Science*, 28(5), 840–856.

Padilla, J. (2024). The 'crisis' of antitrust economics. *Oxford Review of Economic Policy*, 40(4), 706–717.

Pannone, A., Raj, A., Ravichandran, H., Das, S., Chen, Z., Price, C. A., Sultana, M., & Das, S. (2024). Robust chemical analysis with graphene chemosensors and machine learning. *Nature*, 634, 572–578.

Parker, S. C. (2009). *The Economics of Entrepreneurship*. Cambridge University Press, Cambridge.

Pasquinelli, M. (2016). Abnormal encephalization in the age of machine learning. *E-flux Journal*, no. 75: https://images-eflux.b-cdn.net/assets/af0f5664-d80d-4633-aa7b-d52daf9c15e8 (accessed 4 December 2023).

Pasquinelli, M. (2023). *The Eye of the Master: A Social History of Artificial Intelligence*. Verso Books, London.

Pecora NERA (2024). *Il CEO di Datapizza svela come ha creato una start up da +1M/anno – Pierpaolo D'Odorico* (in Italian). Pecora NERA YouTube Channel: www.youtube.com/watch?v=FvsmWw3MM6U (accessed 28 January 2025).

Pedota, M. (2024). Rethinking absorptive capacity in the age of artificial intelligence. *Academy of Management Proceedings*, 2024, 16946.

Penrose, E. (1959). *The Theory of the Growth of the Firm*. Oxford University Press, Oxford.

Peters, B. G. (2018). The challenge of policy coordination. *Policy Design and Practice*, 1(1), 1–11.

Picard, R. W. (1997). *Affective Computing*. MIT Press, Cambridge, MA.

Piccialli, F., Cuomo, S., Crisci, D., Prezioso, E., & Mei, G. (2020). A deep learning approach for facility patient attendance prediction based on medical booking data. *Scientific Reports*, 10(1), 14623.

Piva, E., Grilli, L., & Rossi-Lamastra, C. (2011). The creation of high-tech entrepreneurial ventures at the local level: The role of local competences and communication infrastructures. *Industry and Innovation*, 18(6), 563–580.

Polanyi, M. (1958). *Personal Knowledge: Towards a Post-critical Philosophy*. Routledge, London.

Polanyi, M. (1966). *The Tacit Dimension*. Doubleday & Co., New York, NY.

Prado, T. S., & Bauer, J. M. (2022). Big Tech platform acquisitions of start-ups and venture capital funding for innovation. *Information Economics and Policy*, 59, 100973.

Prainsack, B. (2019). Logged out: Ownership, exclusion and public value in the digital data and information commons. *Big Data & Society*, 6(1), 1–15.

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4), 515–526.

Protogerou, A., Caloghirou, Y., & Vonortas, N. S. (2017). Determinants of young firms' innovative performance: Empirical evidence from Europe. *Research Policy*, 46(7), 1312–1326.

Puranam, P. (2024). *Re-humanize: How to Build Human-centric Organizations in the Age of Algorithms*. Penguin Random House, New York, NY.

Purdy, M. (2024). What is agentic AI, and how will it change work? *Harvard Business Review*, 12 December 2024: https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work (accessed 5 February 2025).

Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon and Schuster, New York, NY.

Rady, J., Townsend, D., & Hunt, R. (2026). From algorithmic hallucinations to alien minds: Addressing the ideator's dilemma through entrepreneurial work. *Journal of Business Venturing*, 41(1), 106550.

Rai, A., Constantinides, P., & Sarker, S. (2019). Next generation digital platforms: Toward human-AI hybrids. *MIS Quarterly*, 43(1), iii–ix.

Raisch, S., & Fomina, K. (2025). Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review*, 50(2), 441–464.

Ramoglou, S., & McMullen, J. S. (2024). 'What is an opportunity?': From theoretical mystification to everyday understanding. *Academy of Management Review*, 49(2), 273–298.

Reback, G. L. (2002). Patently absurd. *Forbes*, 24 June 2002: www.forbes.com/asap/2002/0624/044_print.html (accessed 23 October 2023).

Resnik, P. (2025). Large language models are biased because they are large language models. *Computational Linguistics*, 51(3), 885–906.

Reuters (2025). China to rely on artificial intelligence in education reform bid. *Reuters*, 17 April 2025: http://bit.ly/4sBq7id (accessed 15 January 2026).

Rider, C., & Negro, G. (2015). Organizational failure and intraprofessional status loss. *Organization Science*, 26(3), 633–649.

Rieckmann, M. (2012). Future-oriented higher education: Which key competencies should be fostered through university teaching and learning? *Futures*, 44(2), 127–135.

Ries, E. (2011). *The Lean Startup*. Portfolio Penguin, London.

Riggio, R. E., Eaton, L. G., & Funder, D. C. (2020). Skill in social situations: The essence of savoir-faire. In R. J. Sternberg, and A. Kostić (eds.), *Social Intelligence and Nonverbal Communication* (pp. 333–357). Palgrave Macmillan, Cham.

Rikap, C. (2024a). Intellectual monopolies as a new pattern of innovation and technological regime. *Industrial and Corporate Change*, 33(5), 1037–1062.

Rikap, C. (2024b). Varieties of corporate innovation systems and their interplay with global and national systems: Amazon, Facebook, Google and Microsoft's strategies to produce and appropriate artificial intelligence. *Review of International Political Economy*, 31(6), 1735–1763.

Rocha, V., Carneiro, A., & Varum, C. (2018). Leaving employment to entrepreneurship: The value of co-worker mobility in pushed and pulled-driven startups. *Journal of Management Studies*, 55(1), 60–85.

Rocha, V., & Grilli, L. (2024). Early-stage start-up hiring: The interplay between start-ups' initial resources and innovation orientation. *Small Business Economics*, 62(4), 1641–1668.

Rodrik, D. (2015). *Economics Rules: The Rights and Wrongs of the Dismal Science*. WW Norton & Company, New York, NY.

Rohlfs, J. H. (2003). *Bandwagon Effects in High-Technology Industries*. MIT Press, Cambridge, MA.

Roli, A., Jaeger, J., & Kauffman, S. A. (2022). How organisms come to know the world: Fundamental limits on artificial general intelligence. *Frontiers in Ecology and Evolution*, 9, 806283.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386–408.

Rothaermel, F. T., & Deeds, D. L. (2006). Alliance type, alliance experience and alliance management capability in high-technology ventures. *Journal of Business Venturing*, 21(4), 429–460.

Runco, M. A., & Jaeger, G. J. (2012). The standard definition of creativity. *Creativity Research Journal*, 24(1), 92–96.

Russell, S., & Norvig, P. (2014). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson, London.

Rutherford, M. W., & Buller, P. F. (2007). Searching for the legitimacy threshold. *Journal of Management Inquiry*, 16(1), 78–92.

Ryu, S., Yu, Y., & Seo, H. (2024). Can untrained neural networks detect anomalies? *IEEE Transactions on Industrial Informatics*, 20(4), 6477–6488.

Sadowski, J. (2019). When data is capital: Datafication, accumulation, and extraction. *Big Data & Society*, 6(1), 1–12.

Sagnard, A. (2024). Influencer marketing course prepares students for the future of digital marketing. *Tulane University*, 20 November 2024: https://sopa.tulane.edu/news/influencer-marketing-course-prepares-students-future-digital-marketing (accessed 15 January 2026).

Saheb, T., & Saheb, T. (2023). Topical review of artificial intelligence national policies: A mixed method analysis. *Technology in Society*, 74, 102316.

Samoili, S., López Cobo, M., Gómez, E., De Prato, G., Martínez-Plumed, F., & Delipetrev, B. (2020). AI watch: Defining artificial intelligence. JRC Technical Report EUR 30117 EN, Publications Office of the European Union: https://data.europa.eu/doi/10.2760/382730 (accessed 27 January 2026).

Santarelli, E., & Vivarelli, M. (2002). Is subsidizing entry an optimal policy? *Industrial and Corporate Change*, 11(1), 39–52.

Sarasvathy, S. D. (2001). Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review*, 26(2), 243–263.

Sauermann, H. (2018). Fire in the belly? Employee motives and innovative performance in start-ups versus established firms. *Strategic Entrepreneurship Journal*, 12(4), 423–454.

Saviotti, P. P. (2002). Variety, growth and demand. In A. McMeekin, M. Tomlinson, K. Green and V. Walsh (eds.), *Innovation by Demand* (pp. 41–55). Manchester University Press, Manchester.

Scardapane, S. (2024). Alice's adventures in a differentiable wonderland--Volume I, A tour of the land. arXiv preprint arXiv:2404.17625.

Schank, R. (2006). Once upon a time in AI. *Artificial Intelligence*, 170(18), 1254–1255.

Schilling, M. A. (2019). *Quirky*. Public Affairs – Hachette Book Group, New York, NY.

Schindehutte, M., Morris, M., & Allen, J. (2006). Beyond achievement: Entrepreneurship as extreme experience. *Small Business Economics*, 27, 349–368.

Schmid, J., Sytsma, T., & Shenk, A. (2024). Evaluating natural monopoly conditions in the AI foundation model market. Report by RAND Corporation: https://bit.ly/45byNSs (accessed 15 January 2026).

Schmodde, L., & Wehner, M. C. (2024). Integrating emotion regulation, emotional intelligence, and emotion-focused coping in the entrepreneurial context: A review and research agenda. *International Small Business Journal*, 42(8), 984–1015.

Schumpeter, J. A. (1911/1934). *The Theory of Economic Development*. Harvard University Press, Cambridge, MA.

Schumpeter, J. A. (1939). *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*. McGraw-Hill, New York, NY.

Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers Publishers, New York, NY.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–457.

Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792), 706–710.

Sequoia Capital (2025). *AI's Trillion-Dollar Opportunity: Sequoia AI Ascent 2025 Keynote*. Sequoia Capital YouTube Channel: www.youtube.com/watch?v=v9JBMnxuPX8 (accessed 19 May 2025).

Shah, S. K., Agarwal, R., & Echambadi, R. (2019). Jewels in the crown: Exploring the motivations and team building processes of employee entrepreneurs. *Strategic Management Journal*, 40(9), 1417–1452.

Shane, S. (2000). Prior knowledge and the discovery of entrepreneurial opportunities. *Organization Science*, 11(4), 448–469.

Shane, S. (2003). *A General Theory of Entrepreneurship: The Individual-Opportunity Nexus*. Edward Elgar Publishing, Cheltenham.

Shane, S. (2009). Why encouraging more people to become entrepreneurs is bad public policy. *Small Business Economics*, 33(2), 141–149.

Shane, S., Locke, E. A., & Collins, C. J. (2003). Entrepreneurial motivation. *Human Resource Management Review*, 13(2), 257–279.

Shane, S., & Venkataraman, S. (2000). The promise of entrepreneurship as a field of research. *Academy of Management Review*, 25(1), 217–226.

Shapiro, C. (2018). Antitrust in a time of populism. *International Journal of Industrial Organization*, 61, 714–748.

Shepherd, D. A., Souitaris, V., & Gruber, M. (2021). Creating new ventures: A review and research agenda. *Journal of Management*, 47(1), 11–42.

Shi, F., & Evans, J. (2023). Surprising combinations of research contents and contexts are related to impact and emerge with scientific outsiders from distant disciplines. *Nature Communications*, 14(1), 1641.

Shojaee, P., Mirzadeh, I., Alizadeh, K., Horton, M., Bengio, S., & Farajtabar, M. (2025). The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity. arXiv preprint arXiv:2506.06941.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631(8022), 755–759.

Siepel, J., Cowling, M., & Coad, A. (2017). Non-founder human capital and the long-run growth and survival of high-tech ventures. *Technovation*, 59, 34–43.

Signori, A., & Vismara, S. (2018). M&A synergies and trends in IPOs. *Technological Forecasting and Social Change*, 127, 141–153.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.

Skillman, P. (2025). *Thinking of AI through the Lens of AI and Health and What It Means for Management Scholars – Keynote Speech at the Human Tech Day 2 'Reframing Processes, Organizations, and Systems in the More-than-human Society'*. AV Streaming YouTube Channel: www.youtube.com/live/4E9iUyZQWes (accessed 27 January 2025).

Smil, V. (2023). *Invention and Innovation: A Brief History of Hype and Failure*. MIT Press, Cambridge, MA.

Smith, A. (1776). *The Wealth of Nations*. Oxford University Press, New York, NY.

Smith, K. (2024). The biology of smell is a mystery-AI is helping to solve it. *Nature*, 633(8028), 26–29.

Sobel, R. S. (2008). Testing Baumol: Institutional quality and the productivity of entrepreneurship. *Journal of Business Venturing*, 23(6), 641–655.

Sohrabpour, V., Oghazi, P., Toorajipour, R., & Nazarpour, A. (2021). Export sales forecasting using artificial intelligence. *Technological Forecasting and Social Change*, 163, 120480.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 354–374.

Spitzer, M. (2012) *Digitale Demenz (in German)*. Droemer, Munich.
Srnicek, N. (2025). *Silicon Empires: The Fight for the Future of AI*. Polity, Cambridge.
Stein, M. K., & Shollo, A. (2025). Microfoundations of rationality in the age of AI: On emotions, bodies and intelligence. *Information and Organization*, 35(3), 100583.
Sterk, V., Sedláček, P., & Pugsley, B. (2021). The nature of firm growth. *American Economic Review*, 111(2), 547–579.
Stinchcombe, F. (1965). Social structure and organizations. In J. March (eds.), *Handbook of Organizations* (pp. 142–193). Rand McNally, Chicago, IL.
Stocchi, L., Pourazad, N., Michaelidou, N., Tanusondjaja, A., & Harrigan, P. (2022). Marketing research on Mobile apps: Past, present and future. *Journal of the Academy of Marketing Science*, 50(2), 195–225.
Stokes, S. (2025). Ciroos is building AI teammates that fix tech issues faster. Here's the pitch deck it used to raise $21 million. *Business Insider*, 3 June 2025: www.businessinsider.com/ciroos-pitch-deck-ai-teammate-startup-21m-seed-funding-round-2025-5 (accessed 15 January 2015).
Storey, D. J., & Tether, B. S. (1998). Public policy measures to support new technology-based firms in the European Union. *Research Policy*, 26(9), 1037–1057.
Strozier, M. (2025). Pro take: Get ready, VCs. Solo founders are coming. *Wall Street Journal*, 25 February 2025: www.wsj.com/articles/pro-take-get-ready-vcs-solo-founders-are-coming-346e2aef (accessed 5 July 2025).
Suddaby, R., Bruton, G. D., & Si, S. X. (2015). Entrepreneurship through a qualitative lens: Insights on the construction and/or discovery of entrepreneurial opportunity. *Journal of Business Venturing*, 30(1), 1–10.
Taddy, M. (2019). The technological elements of artificial intelligence. In A. Agrawal, J. Gans and A. Goldfarb (eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 61–87). University of Chicago Press, Chicago, IL.
Tahmincioglu, E. (2004). More companies say, 'The DVD is in the mail'. *New York Times*, 26 December 2004: www.nytimes.com/2004/12/26/business/yourmoney/more-companies-say-the-dvd-is-in-the-mail.html (accessed 15 January 2026).
Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6), 285–305.
Terven, J., Cordova-Esparza, D. M., Ramirez-Pedraza, A., Chavez-Urbiola, E. A., & Romero-Gonzalez, J. A. (2024). Loss functions and metrics in deep learning. arXiv preprint arXiv:2307.02694.
The Economist (2017). The world's most valuable resource is no longer oil, but data. *The Economist*, 6 May 2017: https://bit.ly/4qtk5Pi (accessed 31 January 2024).

The Economist (2025a). Inheriting is becoming nearly as important as working. *The Economist*, 1 March 2025: https://bit.ly/49Bxcqn (accessed 15 January 2026).

The Economist (2025b). Who is winning in AI – China or America? *The Economist*, 2 September 2025: https://bit.ly/49p24f2 (accessed 15 January 2026).

The Nobel Prize (2024). LinkedIn post, 16 October 2024: https://bit.ly/3LBIAL0 (accessed 9 January 2025).

Thiel, P., & Masters, B. (2014). *Zero to One: Notes on Startups, or How to Build the Future*. Crown Currency, New York, NY.

Time (1948). Art: Potter to the Queen. *Time*, 26 April 1948: https://time.com/archive/6600828/art-potter-to-the-queen/ (accessed 16 March 2025).

Tognini, G. (2023). From bodybuilder to billionaire: How Gymshark founder Ben Francis built a sportswear unicorn. *Forbes*, 5 April 2025: https://bit.ly/4jRY5v6 (accessed 15 January 2026).

Townsend, D. M., Hunt, R. A., Rady, J., Manocha, P., & Jin, J. H. (2025). Are the futures computable? Knightian uncertainty and artificial intelligence. *Academy of Management Review*, 50(2), 415–440.

Tranchero, M., Brenninkmeijer, C. F., Murugan, A., & Nagaraj, A. (2024). Theorizing with large language models. National Bureau of Economic Research (Working Paper No. 33033).

Tully, S., Longoni, C., & Appel, G. (2025). Lower artificial intelligence literacy predicts greater AI receptivity. *Journal of Marketing*, 89(5), 1–20.

Tyssedal, J. J. (2026). What is a positional good? Recovering Hirsch's insights. *Economics & Philosophy*, 42(1), 111–132.

UiPath (2025). UiPath announces AI partnership with Google Cloud to transform medical processes. *UiPath News*, 9 April 2025: www.uipath.com/newsroom/uipath-announces-medical-summarization-agent-google-cloud (accessed 26 April 2025).

UNESCO (2022). *Knowledge-driven Actions: Transforming Higher Education for Global Sustainability*. United Nations Educational, Scientific and Cultural Organization, Paris.

Unger, J. M., Rauch, A., Frese, M., & Rosenbusch, N. (2011). Human capital and entrepreneurial success: A meta-analytical review. *Journal of Business Venturing*, 26(3), 341–358.

United States Copyright Office (2025). Copyright and artificial intelligence. Part 3: Generative AI training. Report by United States Copyright Office: https://bit.ly/4qztox9 (accessed 22 December 2025).

Vannuccini, S. (2025). Move fast and integrate things: The making of a European industrial policy for artificial intelligence. Manchester Institute of Innovation Research (MIoIR), University of Manchester, Working paper no. 2025-02.

Varian, H. (2010). Computer mediated transactions. *American Economic Review: Papers & Proceedings*, 100(2), 1–10.

Varian, H. (2019). Artificial intelligence, economics, and industrial organization. In A. Agrawal, J. Gans and A. Goldfarb (eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 399–419). University of Chicago Press, Chicago, IL.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017)*, 5998–6008.

Veblen, T. (1899). *The Theory of the Leisure Class*. Macmillan, New York, NY.

Verdict (2025). Perplexity AI achieves $18bn valuation with new funding. *Verdict*, 18 July 2025: www.verdict.co.uk/perplexity-ai-18bn-valuation-new-funding/ (accessed 20 July 2025).

Verganti, R. (2017). *Overcrowded: Designing Meaningful Products in a World Awash with Ideas*. MIT Press, Cambridge, MA.

Vismara, S., Latifi, G., Meinzinger, L., & Pass, A. (2026). Generative AI-powered venture screening: Can large language models help venture capitalists? *International Review of Financial Analysis*, 109, 104748.

Vomberg, A., Schauerte, N., Krakowski, S., Bogusz, C. I., Gijsenberg, M. J., & Bleier, A. (2023). The cold-start problem in nascent AI strategy: Kickstarting data network effects. *Journal of Business Research*, 168, 114236.

Walker, P., & Dowd, K. (2025). Founder ownership report 2025. Report by Carta: https://carta.com/data/founder-ownership/ (accessed 20 January 2026).

Walker, R., Jeffery, M., So, L., Sriram, S., Nathanson, J., Ferreira, J., & Feldmeier, J. (2017). *Netflix Leading with Data: The Emergence of Data-driven Video*. Kellogg School of Management Cases, Evanston, IL.

Wang, S., Zhang, Y., Xiao, Y., & Liang, Z. (2025). Artificial intelligence policy frameworks in China, the European Union and the United States: An analysis based on structure topic model. *Technological Forecasting and Social Change*, 212, 123971.

Wang, Z., Liu, A., Lin, H., Li, J., Ma, X., & Liang, Y. (2024). Rat: Retrieval augmented thoughts elicit context-aware reasoning in long-horizon generation. arXiv preprint arXiv:2403.05313.

Weber, M. (1922/1946). Science as a vocation. In H. H. Gerth and C. Wright Mills (eds.), *From Weber: Essays in Sociology* (pp. 129–156), Oxford University Press, New York, NY.

Weitzman, M. L. (1996). Hybridizing growth theory. *American Economic Review: Papers and Proceedings*, 86(2), 207–212.

Williamson, O. E. (2000). The new institutional economics: Taking stock, looking ahead. *Journal of Economic Literature*, 38(3), 595–613.

Winter, S. G. (2016). Tacit knowledge. In M. Augier and D. J. Teece (eds.), *The Palgrave Encyclopedia of Strategic Management* (pp. 1–3). Palgrave Macmillan, London.

Wired (2025). *Bill Gates on His Early Years & the Inspiration Behind 'Source Code: My Beginnings'*. Wired Webcast: https://bit.ly/49XFqKP (accessed 11 March 2025).

Wood, M. (2020). Raising capital for startups: 8 statistics that will surprise you. *Fundera*, 3 February 2020: www.fundera.com/resources/startup-funding-statistics (accessed 21 March 2021).

Woolcock, N. (2024). Video gamer looking for that extra edge? Hire a performance coach. *Times*, 14 December 2024: https://bit.ly/4pHUpNO (accessed 13 April 2025).

Wooldridge, M. (2020). *The Road to Conscious Machines: The Story of AI*. Penguin Random House, London.

World Economic Forum (2024). Navigating the AI frontier: A primer on the evolution and impact of AI agents. World Economic Forum White Paper: https://reports.weforum.org/docs/WEF_Navigating_the_AI_Frontier_2024.pdf (accessed 6 June 2025).

Wu, A., Higgins, M., Zhang, M., & Jiang, H. (2024). AI wars. Harvard Business School Case No. 723-434. Harvard Business School Publishing, Boston, MA.

Wu, A., & Yang, A. (2025). AI wars in 2025. Harvard Business School Case No. 725-484. Harvard Business School Publishing, Boston, MA.

Wu, B., Chen, H., & Shi, Y. (2025). Influence of artificial intelligence development on supply chain diversification. *Finance Research Letters*, 78, 107210.

Wu, V., & Gnanasambandam, C. (2017). A machine-learning approach to venture capital. *McKinsey Quarterly*, 26 June 2017: https://bit.ly/49orl95 (accessed 21 May 2025).

Xiao, S., & Hu, B. (2026). Exploring the relationship between ChatGPT use and students' higher-order thinking skills in tertiary education: A case study. *Education and Information Technologies*, 31, 3077–3095.

Xie, Y., & Wang, S. (2025). Generative artificial intelligence in entrepreneurship education enhances entrepreneurial intention through self-efficacy and university support. *Scientific Reports*, 15(1), 24079.

Y Combinator (2025). *Cursor CEO: Going beyond Code, Superintelligent AI Agents, and Why Taste Still Matters*. Y Combinator YouTube Channel: www.youtube.com/watch?v=oOylEw3tPQ8 (accessed 17 June 2025).

Yang, J. (2026). Unintended effects of Digital Market Act on new entry and investment. *Digital Policy, Regulation and Governance*, 28(4), 355–371.

Yusubova, A., Andries, P., & Clarysse, B. (2020). Entrepreneurial team formation and evolution in technology ventures: Looking beyond the top management team. *Journal of Small Business Management*, 58(5), 893–922.

Zhang, A. H. (2024). *High Wire*. Oxford University Press, New York, NY.

Zhao, H., & Seibert, S. E. (2006). The big five personality dimensions and entrepreneurial status: A meta-analytical review. *Journal of Applied Psychology*, 91(2), 259–271.

Zhou, L., Paul, S., Demirkan, H., Yuan, L., Spohrer, J., Zhou, M., & Basu, J. (2021). Intelligence augmentation: Towards building human-machine symbiotic relationship. *AIS Transactions on Human-Computer Interaction*, 13(2), 243–264.

Index

abduction, 80
absorptive capacity, 95
accessibility and openness (AI data and models), 49, 62–69, 84, 88–89, 96, 101, 165–167, 176, 192–193, 227, 229–232
acqui-hiring, 164, 192
acquisitions, 7–8, 33, 152, 179–186, 190–195, 211, 233–234
  killer, 194–196
a-criticality (AI), 55–56, 228, 239
action space (innovative entrepreneurship), 8, 101, 123, 165–166, 197
adverse selection, 32, 173–176, 179, 190, 217
AI Act (EU governance), 219, 223–227
AI engineers. *See* AI talent
AI entrepreneurial buddies, 119–123, 135–140, 142–144, 148, 154, 161–162, 166, 189, 209, 211, 215
AI giants, 7–8, 35, 44, 58, 68–69, 82, 159, 168, 174–186, 190–199, 206, 209–212, 217–218, 228, 232–236, 244. *See also* MAMAA
AI policy buddies. *See* AI entrepreneurial buddies
AI science (R&D), 7, 38–40, 52–55, 170–172, 177–178, 190–193
AI supply chains, 57–69, 161
AI talent, 8, 130, 173–176, 179, 190–193, 214
AI VC buddies. *See* AI entrepreneurial buddies
AI wars, interstate, 198–199, 241
AI(d) offering entrepreneurship, 82, 86–88, 96, 117–119, 148, 151, 163–164, 166, 244
AI-based start-ups, 7, 64, 96, 129–133, 155, 158–164, 168–186, 191–198, 211, 214, 217, 223–226
  creativity-centric, 7, 161–163, 169
  cross-AI function and cross-sectoral domain, 7, 159, 163–164, 169
  predictive analytics, 7, 159–161, 169
AI-Becker problem. *See* seniority-biased technological change (AI)
AI-served start-ups, 7, 96, 164–168, 177, 179, 190–191
  business intelligence-enhanced, 7, 164–167
  creativity-augmented, 7, 164–167
AI-vertical foreclosure, 167–168, 176
algorithmic aversion, 51, 81, 148, 153, 228, 239
algorithmic deforestation, 75, 230
AlphaFold, 38, 57, 76, 159, 167, 176
AlphaGo. *See* Go
anthropomorphization (AI), 13, 112–116
antitrust, 2, 199, 227
  Chicago School, 233
  market investigation tool, 228, 232–236
  New Brandeis Movement, 124, 199, 233
  rule of law, 235
  rule of reason, 235
  structure–conduct–performance paradigm (Harvard), 233
application programme interfaces (APIs), 61–64, 68, 194
Artificial Intelligence (AI). *See also* deep learning
  agentic, 5–6, 12, 41, 47–48 54, 56, 62, 64–65, 82, 108, 117–123, 126, 136, 144, 163–164, 169, 189, 207, 211

AGI, 41–42, 115
definition of, 11–12
explainable, 46, 51
foundation models, 60–65, 94, 163, 178
general-purpose technology, 8, 12, 198, 217
generative, 5–7, 12, 36–57, 59–64, 67–69, 73–76, 80–82, 90–97, 99–104, 115, 119–123, 130, 132, 136–138, 143, 158, 161–167, 169, 174, 178, 189, 207, 210–211
narrow vs. general, 41–42
neuro-symbolic, 46
predictive, 5–7, 12, 36–59, 61, 64, 66–69, 75–79, 82–89, 95–97, 99–102, 104, 132, 150–153, 158–161, 163–169, 178, 189, 210, 215, 230
subsymbolic, 36–37, 50
symbolic, 36–37, 50
weak vs. strong, 41–42

Baudrillard, Jean, 208, 244
Baumol, William J., 3, 14–16, 19–20 35, 107, 124, 127, 190, 204–205, 220, 228, 244
Bengio, Yoshua, 40
bias
algorithmic (AI), 43, 53–54, 152, 214
automation (AI), 51
home, 171
Big Tech, 1, 7, 35, 39, 120, 168, 176–186, 195, 199, 210. *See also* AI giants; MAMAA
bodily engagement, 108–112, 116–119, 123
Bowles, Paul, 5, 14, 124, 190, 212, 244
brittleness (AI), 49, 153, 214

capitalism
fair, 5, 14, 124, 190, 212
fenced. *See* fenced capitalism
good, 3, 35, 127, 190, 220
shapeshifting, 1, 11, 244
shrinking, 200
varieties, 11, 19, 212, 244
cloud, 43–44, 178–179, 194–195, 228, 236

competence-based view (of the firm/start-up), 24, 125–126, 138–140
convolutional neural networks (CNNs or ConvNets), 38–39, 44, 60, 76
coordination costs, 139, 145–147, 154
creativity, 6, 21–23, 73–82, 90–97, 161–167
convergent thinking, 22–23, 79–82, 90–96
divergent thinking, 6, 22–23, 79–82, 90–96
crowdfunding, 32, 93, 122, 149, 155–157, 214
Crunchbase, 131, 168, 171, 181, 196, 224

data commons, 229–232
data scraping, 84, 232
datafication, 59, 79, 153, 159, 198
deep learning. *See also* Artificial Intelligence (AI)
architectures, 38–40
digital statistical machine, 5, 44–57, 59, 67, 78, 95–96, 102, 114, 173, 189
human capital (founders'), 130–133, 142
inferential infrastructure, 43–44, 96, 178, 185, 190, 194–195, 210, 230
multi-layer neural networks, 36, 40, 44–49, 54, 86, 99, 178
venture capital (VC), 150–153, 156
deepfakes, 55
Delphi oracle. *See* hypersuasion
digital dementia, 56, 209
Digital Markets Act (EU governance), 199, 223, 233–234
Digital Omnibus (EU governance), 224, 226
digital servitization, 58, 146
drift
concept, 49, 87, 153, 214
data, 49, 153, 214

econometrics, 43–48
original estimates, 131, 172, 181, 201, 224–226
education (AI), 128, 131, 138, 227–228, 238–241
effectuation, 23

efficiency
  allocative, 14–15
  dynamic, 14–15, 17, 197, 231
  productive, 14–15, 85–89, 96
emotion AI, 111–112, 164
emotional intelligence, 110–111, 164
emotional robotics, 112
empathy circuit, 110–111
end of
  history, 244
  truth, 55
Engel's law, 207
entrepreneurial bricolage, 23
entrepreneurial idea
  definition of, 21–22
  ideator's dilemma, 102
  Kirznerian ideas. *See* Kirznerian entrepreneurship
  Not Thought Here syndrome, 6, 103–107
  Schumpeterian ideas. *See* Schumpeterian entrepreneurship
entrepreneurial opportunity. *See* entrepreneurial idea
entrepreneurial savoir-faire, 108–112, 116, 141–142, 148, 155, 240
entrepreneurial society, 15, 217. *See also* capitalism, good
entrepreneurial team size. *See* solopreneurship
entrepreneurship policy, 3–4, 8, 32–34, 212–220. *See also* institutions, reforms
external enabler (of entrepreneurship), 91

fenced capitalism, 8, 211–212
flow (psychology), 28, 104–105
freelancers. *See* gig economy

GAFAM. *See* MAMAA
gambler's ruin theory, 4
GDPR (EU governance), 221, 223, 226
generative adversarial networks (GANs), 38–39, 60
gig economy, 200–202, 209, 237
Go, 46, 76–79
  moves (37 & 78), 76–79, 81, 153, 230
grand societal challenges, 53, 228, 238–241

hallucinations (AI), 53–55, 102
Hamlet, Shakespeare, 14–16, 112, 244
Hassabis, Demis, 58, 77
heterogony of ends, 240
hikikomori entrepreneurship, 116–117
Hinton, Geoffrey, 13, 40, 42, 115
homogenization (AI), 57, 202–203, 239
human capital (founders'), 1, 4, 6–7, 23–28, 129–143, 157, 171–176, 199, 211, 217
  deep learning, 130–133, 142
  dynamic AI-enabled orchestration capability, 7, 135–138, 143, 148, 154, 166, 228, 240
  generic, 24–25, 129–135, 141–142, 211
  hard skills, 141–142
  imprinting, 26
  soft skills, 141–142
  specific, 6, 23–28, 95, 108, 117, 125, 129–135, 139, 141–142, 148, 170, 199, 208, 211, 217, 230
  wealth effect, 139
hybrid intelligence (human–AI), 6, 56, 75, 81, 133, 143, 216, 239
hypersuasion, 51, 55, 153

in silico samples, 121
incumbent advantage, 6–8, 65–69, 83–89, 96, 101, 165–167, 190, 198, 217–218, 227–232, 244
indirect network externalities, 48, 62, 64, 66, 168, 176
  data flywheel, 49
  data-driven learning effects, 48, 66–67, 69, 84, 137, 168, 178, 229
  feedback loops, 48–49, 57, 61–62, 66–67, 84, 115, 137, 177–178, 203
  winner-takes-all (or most), 62–64, 197–198
industrial policy, weaponization, 8, 241–242
initial coin offering, 32, 155, 157
innovation
  by meaning, 95, 208
  combinatorial, 90
  empty, 204–205
  ephemeral, 205
  thick. *See* thick innovation
  thin. *See* thin innovation
  toxic, 210

institutions, 1–3, 8, 15–16, 34, 96, 107, 124, 126, 150, 190, 205, 213–242
  definition of, 19
  formal, 19–20, 126, 212, 228
  informal, 19–20, 107, 126, 155, 202, 212, 229, 239
  institutional matrix. *See* North, Douglass C.
  reforms, 219–242
intellectual property rights (IPRs), 17, 67–69, 120, 228, 231–232
  patents, 17, 26, 58, 67, 120, 130, 156, 171, 182
IPOs, 33, 152, 185

Keynes, John Maynard, 206
Kirznerian entrepreneurship, 6, 82–89, 96, 101, 165, 167, 190, 203, 231
Knightian uncertainty, 6, 99–103, 106, 108, 118, 127–129, 143, 152, 214, 220
Knowledge spillover theory of entrepreneurship (KSTE), 25–26, 125, 170–172, 199
knowledge tacitness, 50–52, 103, 107, 153, 189
Kranzberg, Melvin
  six laws on technological development, 12–13, 190

labour market turmoil, 8, 13, 98, 115, 126–129, 174, 176, 179, 200–202, 206–207, 214, 217–218, 227–228, 236
Large Concept Models (LCMs), 54. *See also* Large Language Models (LLMs)
Large Language Models (LLMs), 38–39, 49, 62–64, 68, 74–75, 80, 90–95, 121, 150, 162, 178, 211, 239
Large Multimodal Models (LMMs), 38, 62–64. *See also* Large Language Models (LLMs)
Large Reasoning Models (LRMs), 38, 40, 50, 55. *See also* Large Language Models (LLMs)
leadership (entrepreneurial), 6, 29, 103–106, 131, 139–140, 145
LeCun, Yann, 40

M&As. *See* acquisitions
machine learning, 36
MAMAA, 7, 32, 35, 39, 44, 168, 174, 176–186, 191–198, 211. *See also* AI giants
market failures, 29–32, 213–214
Michelangelo
  The Creation of Adam, 113–115
middleware, 62, 64
Model Context Protocol (MCP), 64

necessity entrepreneurs, 207, 218. *See also* technological unemployment
need for achievement, 27–28, 106–108, 241
Nexis Uni, 183
North, Douglass C., 15, 19
  institutional matrix, 19, 34, 229

passive entrepreneurship, 117–119, 126, 140
physical AI, 114–116. *See also* robotics
PitchBook, 131, 151, 156, 168–169, 181
polarization (AI), 57, 202–203, 218, 239
Pope Francis, 59
post-essential consumption and economy, 207–209, 218. *See also* thin innovation
postmodern entrepreneurship, 8, 128, 135, 207–209, 211–212, 216, 218, 227–228, 240, 244
postmodernism (postmodernity). *See* postmodern entrepreneurship
precautionary principle (policy), 219, 227
processing units, 43, 52, 178

quantitative empirical evidence (original)
  descriptive statistics, 131, 147, 155–157, 168–169, 180–183, 196, 201, 224
  econometric analyses. *See* econometrics, original estimates
  geoparsing, 172
  textual analysis, 172, 183, 201, 222, 238
quantum computing, 54

regulation. *See* antitrust; antitrust, market investigation tool
regulatory crackdown (China governance), 199, 210
retrieval-augmented generation (RAG), 49, 54
retrieval-augmented thought (RAT), 54
robotics, 11, 48, 109, 112–116, 148, 156
Rodrik, Dani, 31, 126

safety-net programmes. *See* universal basic income (UBI)
Schumpeter, Joseph A., 16–19, 82, 90, 101, 108, 118, 124, 127–128, 182, 231, 235
  creative destruction, 16–17, 20, 82, 197, 226, 230–231
  entrepreneurial motivations, 16–17, 27, 173, 182–186, 235
  Mark I (SM I), 16–17, 124, 170, 173, 197, 226, 230–231, 243
  Mark II (SM II), 18–19
Schumpeterian entrepreneurship, 6, 82–83, 90–97, 101, 166, 190, 203, 216
Scopus, 171, 183, 238
self-efficacy, 27, 103, 106–108, 134, 141, 240
seniority-biased technological change (AI), 199–202, 208, 218
small-bore and short-haul entrepreneurship, 117, 154, 202
Smith, Adam, 124, 240
social capital (founders'), 1, 4, 6–7, 28–31, 144–149, 157
  legitimacy, 7, 30, 144, 147–149
  network reach, 7, 30, 144–147
social mobility, 17, 128–129, 143
social welfare, 2, 4, 32, 35, 86–88, 128, 180, 185, 191, 197, 205, 209, 212–215, 217–219, 226, 230–232
  definition of, 15
Socratic questioning, 239
SoleScribe Snack Shoe, 91–95, 97, 120–123, 136–138, 162, 166, 207–209
solopreneurship, 28, 117, 138–140, 143, 209, 213
sovereignty
  AI, 63, 198–199
  personal data, 229
Stanford University (AI), 39, 62, 168, 178, 222
Stargate project, 178
sycophancy (AI), 53, 239
symbiotic intelligence. *See* hybrid intelligence (human–AI)
synthetic data (AI training), 37–38, 49, 59–60, 153
System 2 (Kahneman)
  critical thinking, 56. *See also* a-criticality
systemic lock-in. *See* Systemic Lock-in Model
Systemic Lock-in Model, 8, 186, 190–198, 211–212, 217–218, 227–228, 232–236
systemizing mechanism, 83

talent trap, golden, 7, 173–176, 214
technological unemployment, 206–207, 227–228, 236. *See also* labour market turmoil
theatricality, 6, 108–112, 141
theory of mind, 111
thick innovation, 209, 216, 240, 244
thin innovation, 8, 203–212, 216, 218, 240. *See also* postmodern entrepreneurship
transformer, 38–40, 60, 62, 164
trust (AI), 51–52, 78, 95, 103, 107, 119, 148, 153, 189
Twain, Mark, 126, 157

unicorn, 106, 135, 139, 155, 157, 169
unintended consequences (policy), 8, 220–221, 223–228, 233–236
universal basic income (UBI), 228, 236–238

variational autoencoders (VAEs), 38–39
Vaswani, Ashish, 39, 60, 164
Veblen, Thorstein, 203–204
Veblenian entrepreneurship, 203–204, 216

venture capital (VC), 1, 4, 7, 26, 30–34, 149, 157, 169, 179–182, 185, 189–190, 193–194, 214–216, 234
  cloud, 194
  coaching, 7, 33, 149, 153–157, 215–216
  commoditization, 153
  deep learning, 151–153, 156
  democratization, 151
  dynamic AI-enabled orchestration capability, 154
  selection, 7, 33, 149–153, 157, 193–194, 215–216

watermarking, 55, 109
Weber, Max, 2–3, 243