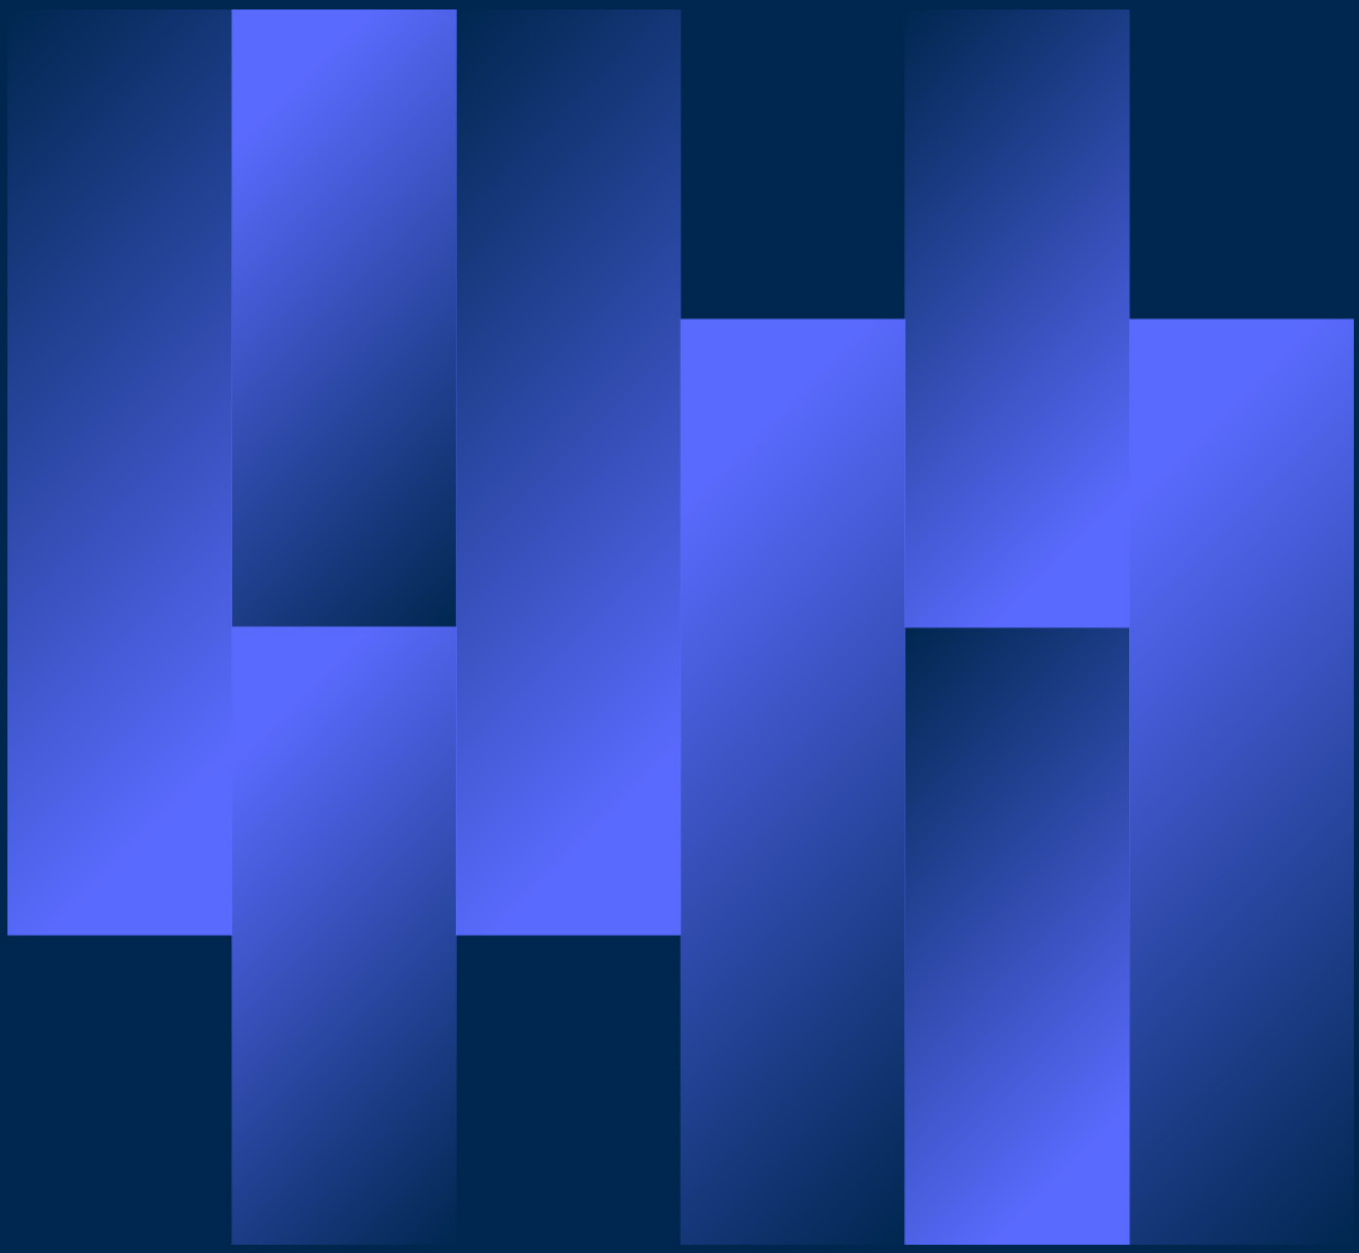
Superintelligence Series

One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business



South American Publishing

Superintelligence Series

Volume 4

One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business

Edited by

Carlos Alberto Gómez Cano

South American Publishing

Copyright Page

© 2026. The authors. This is an open access book, distributed under the terms of a Creative Commons Attribution 4.0 license (https://creativecommons.org/licenses/by/4.0/) that allows use, distribution, and reproduction in any medium provided that the original work is properly cited.

This South American Publishing imprint is published by South American Publishing.

For more information, see South American Publishing's Open Access Policy: https://southam.pub/policies.php

Series: SAP Superintelligence Series
ISSN: 3144-5624
Title of the volume: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (complete work): 978-9915-9851-0-7
ISBN (this volume, eBook): 978-9915-9966-3-9
ISBN (this volume, Softcover): 978-9915-9966-2-2

This title is available in both print (softcover) and digital (open access PDF) formats.

DOI: 10.62486/978-9915-9966-3-9

Cataloging Data

Cataloging-in-Publication Data (CIP):

South American Publishing

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business / edited by Carlos Alberto Gómez Cano. Montevideo: South American Publishing, 2026.

ISBN (English edition, eBook): 978-9915-9966-3-9

ISBN (English edition, Softcover): 978-9915-9966-2-2

THEMA classification codes:

UYQE

BISAC classification:

BUS025000 - BUSINESS & ECONOMICS / Entrepreneurship

BUS115000 - BUSINESS & ECONOMICS / Freelance & Self-Employment

COM100000 - COMPUTERS / Artificial Intelligence / Generative AI

BUS041000 - BUSINESS & ECONOMICS / Management

Editorial Notice and Acknowledgments

The publication of this book is part of South American Publishing's mission to promote open, ethical, and rigorous scientific communication across all areas of knowledge.

All books published by South American Publishing undergo a double-blind peer-review process and a technical evaluation in accordance with the publisher's editorial policies, aligned with COPE guidelines.

South American Publishing acknowledges the valuable collaboration of authors, reviewers, designers, and production teams who made this publication possible.

Preface

The prospect of a company run by one person has a particular appeal. It suggests independence, room for judgment, and the possibility of bringing an idea into the world without first assembling a large organization. Artificial intelligence has given this familiar aspiration a different technical basis. A founder can now draw on systems that produce text, write software, analyze information, and coordinate defined tasks. Work that once called for several specialists may, under some conditions, be carried out by one person working with digital tools.

Yet performing a task is not the same as sustaining an enterprise. A customer must still be understood, a promise kept, a defect corrected, and a decision defended when its consequences become apparent. The apparent simplicity of the one-person company can conceal a considerable network of dependencies. Models are trained and hosted by organizations; software frameworks depend on maintainers; digital platforms govern access to markets and infrastructure. A founder may have no employees and still rely on the labor, policies, and technical choices of many others.

This volume examines that tension from several directions. Its thirteen chapters draw on business formation statistics, patterns of self-employment, records of open-model creation and software contribution, experimental research on productivity, and studies of adoption by smaller firms. They also turn to the theories of the firm, the design of artificial-intelligence agents, and questions of legal responsibility. Such breadth is necessary because no single dataset or discipline can establish what an entire form of enterprise is likely to become.

The later chapters return to the figure of the solitary founder, now with a sharper set of questions. How much of a firm's work can be entrusted to coordinated software agents? Who notices when those agents repeat an error or act beyond their intended remit? What happens when the services on which the firm depends change their terms, become unavailable, or capture a greater share of its value? Legal personality, liability, and taxation acquire an unusual intensity when a single person remains the point at which automated activity meets human responsibility.

There is a human dimension to these questions. The ability to work independently may widen the choices available to some people. It may also leave them unusually exposed when technology fails or a market changes. The aim of the chapters is neither to romanticize solitary enterprise nor to dismiss it. It is to examine what the evidence permits us to say, where theory helps to interpret it, and where confident prediction would be premature.

Readers will encounter a possibility still being tested rather than a settled model of economic life. The most useful question may not be whether one person can eventually build an extraordinarily valuable company. It is what kind of enterprise can endure when many activities are automated, but trust, responsibility, and judgment still have to reside somewhere.

Introduction

What becomes of a firm when some of the work for which it was created can be carried out by artificial intelligence? The question reaches beyond automation in an existing workplace. It concerns the boundaries of the organization itself: why people work together, which activities a founder can perform independently, and which obligations remain even when human labor is replaced by software. The one-person company offers a useful way of examining these boundaries because it takes the promise of automated work close to its logical limit.

A one-person company in the sense considered here is more specific than ordinary self-employment. It is an enterprise in which a founder relies heavily on artificial intelligence to perform activities conventionally assigned to employees, contractors, or managers. The more ambitious image of a one-person unicorn extends that arrangement to a company valued at one billion dollars. The chapters treat this latter image as a hypothesis, not as a verified outcome. That distinction matters. Being able to draft a marketing campaign, build a prototype, or coordinate a workflow does not, on its own, demonstrate that an enterprise can survive, earn durable revenue, or grow to extraordinary scale.

The book begins by examining what can be observed in economic activity. Chapter 1 studies changes in United States business applications around the public release of generative artificial intelligence in November 2022. It identifies a slowdown in the aggregate monthly trend alongside different responses across knowledge-intensive sectors. An interrupted time-series design can describe a change associated with that date, but the wider economy also changed during the period; the timing alone cannot establish that artificial intelligence caused the observed pattern. Chapter 2 considers nonemployer establishments, a closer though still imperfect measure of businesses without paid staff. Their growth between 2018 and 2024 was stronger outside the sectors classified as most exposed to artificial intelligence. Chapter 3 shifts the lens to European countries and finds no significant country-level association between enterprise AI adoption and solo self-employment in knowledge-intensive services in the period examined. Together, these chapters caution against reading the availability of a new technology as proof of an immediate entrepreneurial transformation.

The next three chapters ask what individual contribution looks like inside the technical world that makes such firms conceivable. Chapter 4 compares individual and organizational creators among the most-downloaded models on Hugging Face. Organizations account for most of that highly visible segment, while a small group of independent creators reaches substantial audiences in particular niches. Chapter 5 follows single-authored artificial-intelligence preprints and finds a declining share in its sample, although the estimated trend is not statistically significant. Chapter 6 examines popular open-source agent frameworks and shows that contributions can be concentrated among a few maintainers. These findings reveal an important distinction between individual achievement and individual independence. A solo founder may draw strength from open technical resources while becoming dependent on projects maintained by a narrow group of people.

A further set of questions concerns what artificial intelligence actually changes in daily work. Chapter 7 synthesizes experiments on worker productivity, where the most consistent improvements concern the time required to complete particular tasks. Improvements in quality are less uniform, and shared creative output raises a separate concern about similarity across results. The move from faster individual performance to a viable business requires organizational capacities that task experiments cannot measure. Chapter 8 examines the adoption of artificial intelligence in micro and small enterprises. Its findings emphasize managerial support, the ability to absorb new knowledge, and the perceived fit between a tool and a firm's needs. Chapter 9 turns to entrepreneurship itself, distinguishing recognition of an opportunity from the work of creating a venture.

Its review reports promising effects in venture formation but also risks of homogenized ideas and misplaced trust. The role of the founder is changed, not eliminated.

The book then turns from observed patterns to the concepts needed to explain them. Chapter 10 draws together theories of superstar firms, transaction costs, organizational knowledge, and the allocation of tasks between people and machines. The one-person company becomes intelligible as a possible limit of these processes, though theoretical coherence cannot substitute for a demonstrated case. Chapter 11 examines multi-agent systems, in which specialized software agents coordinate parts of a workflow. The reviewed evidence supports their use in bounded activities such as process coordination and selected business functions. It offers less assurance where decisions are open-ended, errors propagate between agents, or security and oversight are weak. An automated workforce remains something a person must design, supervise, and answer for.

Questions of authority and responsibility become more pronounced in the final chapters. Chapter 12 considers legal personality, liability, and taxation. It distinguishes the legal standing of an incorporated business from proposals to confer personhood on artificial-intelligence systems. Existing legal arrangements continue to depend on human or corporate actors who bear rights and duties, even when software carries out operational tasks. Chapter 13 examines the prospect of the one-person unicorn through systemic risk, platform dependence, and alternative futures. A founder with limited payroll may still be vulnerable to a single provider's prices, access rules, or interruptions. The same infrastructure that makes small-scale operation possible can also concentrate economic power elsewhere.

These chapters should not be read as if their findings came from one continuous experiment. They study different units: applications to start businesses, establishments without employees, national adoption rates, model repositories, individual workers, and bodies of law. The periods and measures are not interchangeable. Some chapters offer statistical associations; others synthesize experimental evidence or develop theoretical and legal arguments. The different methods are a strength when their boundaries remain visible, because each illuminates a question that the others cannot answer. They also prevent a single attractive narrative from taking the place of evidence.

The reader need not decide in advance whether the one-person company represents a desirable future. The chapters invite a more discriminating inquiry into when it works, for whom it works, and at whose expense its apparent independence is sustained. At present, the evidence supports meaningful changes in individual capabilities and selected business processes. It does not establish a general replacement of firms by solitary founders. Between those two propositions lies the field of inquiry that gives this volume its coherence.

Abstract

The possibility of an enterprise operated by a single founder with extensive support from artificial intelligence raises questions about business formation, the organization of work, and the distribution of economic responsibility. Bringing together thirteen chapters, this volume examines the empirical basis and theoretical limits of the one-person company. Studies of business applications, nonemployer establishments, and European self-employment patterns find no consistent evidence of a broad entrepreneurial surge attributable to AI diffusion. Analyses of open models, research authorship, and software repositories show that individual contributions remain possible, although prominent technical resources often depend on organizations or concentrated groups of maintainers. Reviews of experimental productivity research identify gains in the speed of selected tasks, with more variable effects on quality and creativity. Research on smaller enterprises and venture creation points to managerial capability, founder judgment, and the fit between tools and tasks as conditions shaping adoption and outcomes. The volume also considers how theories of the firm can account for AI-assisted production, what multi-agent systems can accomplish in bounded workflows, and why legal personality, liability, taxation, and platform dependence remain unresolved or consequential. The prospect of a one-person unicorn is treated as an unverified hypothesis rather than an observed business model. Across these perspectives, the central distinction is between automating activities and sustaining an accountable enterprise. Artificial intelligence may expand what an individual can accomplish, but the viability and independence of a firm continue to depend on institutional arrangements, oversight, and the distribution of risk.

Keywords: Artificial intelligence; One-person companies; Entrepreneurship; Solo self-employment; Multi-agent systems; Platform dependence

Index

Chapter 01. Business formation and the public release of generative artificial intelligence: an interrupted time-series analysis | La formación de empresas y la publicación pública de la inteligencia artificial generativa: un análisis de series temporales interrumpidas 1

Chapter 02. Nonemployer business growth in artificial-intelligence-exposed industries of the United States | Crecimiento de los negocios sin empleados en industrias expuestas a la inteligencia artificial en Estados Unidos 12

Chapter 03. Enterprise artificial intelligence adoption and solo self-employment in knowledge-intensive activities: an ecological study across European countries | Adopción empresarial de inteligencia artificial y autoempleo en solitario en actividades intensivas en conocimiento: un estudio ecológico en países europeos 24

Chapter 04. Individual and organizational creators of the most-downloaded open artificial intelligence models | Creadores individuales y organizativos de los modelos de inteligencia artificial abiertos más descargados 37

Chapter 05. Single-authored papers in artificial intelligence: a time-trend analysis of arXiv metadata | Artículos en solitario en inteligencia artificial: un análisis de tendencia temporal de los metadatos de arXiv 48

Chapter 06. Contributor concentration in popular open-source artificial-intelligence agent frameworks | Concentración de contribuidores en marcos de agentes de inteligencia artificial de código abierto populares 59

Chapter 07. Generative artificial intelligence and individual worker productivity: a systematic review of experimental evidence | Inteligencia artificial generativa y productividad individual del trabajador: una revisión sistemática de la evidencia experimental 70

Chapter 08. Determinants of artificial intelligence adoption in micro and small enterprises: a systematic review | Determinantes de la adopción de inteligencia artificial en micro y pequeñas empresas: una revisión sistemática 83

Chapter 09. Generative artificial intelligence in entrepreneurship: a systematic review of empirical studies on opportunity recognition and venture creation | Inteligencia artificial generativa en el emprendimiento: una revisión sistemática de estudios empíricos sobre reconocimiento de oportunidades y creación de empresas 95

Chapter 10. From superstar firms to super-individuals: theoretical foundations of the one-person company | Fundamentos teóricos de la empresa unipersonal con inteligencia artificial 107

Chapter 11. Multi-agent artificial intelligence systems as the workforce of the one-person company | Sistemas multiagente de inteligencia artificial como fuerza de trabajo de la empresa unipersonal 118

Chapter 12. Legal personality, liability and taxation of artificial-intelligence-run one-person companies

| Personalidad jurídica, responsabilidad y tributación de las empresas unipersonales gestionadas por inteligencia artificial..... 129

Chapter 13. The one-person unicorn hypothesis: systemic risks, platform dependence and future scenarios | La hipótesis del unicornio unipersonal: riesgos sistémicos, dependencia de plataformas y escenarios futuros..... 140

Chapter 1 / Capítulo 1

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch01
Pages: 1-11

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Business formation and the public release of generative artificial intelligence: an interrupted time-series analysis

La formación de empresas y la publicación pública de la inteligencia artificial generativa: un análisis de series temporales interrumpidas

Libardo Ramón Polanía[1]
Rómulo Esteban Montilla[2]

1 Universidad de la Amazonia. Florencia, Colombia.
2 St. Mary's University, Texas, Estados Unidos

ABSTRACT

Introduction: The public release of generative artificial intelligence in November 2022 promised to lower the cost of starting and running a firm, yet its effect on aggregate business formation remains unmeasured. Methods: We used monthly business applications from the US Census Bureau Business Formation Statistics, seasonally adjusted, from January 2018 through August 2026, and fitted an interrupted time-series model with a break at November 2022, estimating level and trend changes with Newey-West standard errors (12 lags) for the aggregate and for three knowledge-intensive sectors. Results: At the national level the immediate shift was −49,832 applications (p = 0.099) and the monthly trend slowed by 1,602 applications (p = 0.034), from a pre-release slope of 3,631. Sectoral responses diverged, with Information falling 2,123 applications at the break (p = 0.032) and Management of companies accelerating by 44 applications per month (p < 0.001). Conclusions: The release of generative artificial intelligence coincided with a slowdown in aggregate business formation and with divergent, sector-specific responses, a pattern consistent with skill-biased reallocation rather than a uniform startup surge.

Keywords: business formation; generative artificial intelligence; interrupted time series; entrepreneurship; United States

RESUMEN

Introducción: La publicación pública de la inteligencia artificial generativa en noviembre de 2022 prometía abaratar el costo de crear y operar una empresa, pero su efecto sobre la formación agregada de empresas sigue sin medirse. Métodos: Usamos las solicitudes mensuales de creación de empresas de la Business Formation Statistics del US Census Bureau, desestacionalizadas, de enero de 2018 a agosto de 2026, y ajustamos un modelo de series temporales interrumpidas con un corte en noviembre de 2022, estimando cambios de nivel y de tendencia con errores estándar de Newey-West (12 rezagos) para el agregado y para tres sectores intensivos en conocimiento. Resultados: A nivel nacional el cambio inmediato fue de −49 832 solicitudes (p = 0,099) y la tendencia mensual se desaceleró en 1 602 solicitudes (p = 0,034), desde una pendiente previa de 3 631. Las respuestas sectoriales divergieron, con Información cayendo 2 123 solicitudes en el corte (p = 0,032) y Gestión de empresas acelerándose en 44 solicitudes mensuales (p < 0,001). Conclusiones: La publicación de la inteligencia artificial generativa coincidió con una desaceleración de la formación agregada de empresas y con respuestas sectoriales divergentes, un patrón consistente con una reasignación sesgada por habilidades más que con un auge uniforme de nuevas empresas.

Palabras clave: formación de empresas; inteligencia artificial generativa; series temporales interrumpidas; emprendimiento; Estados Unidos

INTRODUCTION

The United States entered the 2020s with a business formation system that had already been losing dynamism for decades. Decker *et al.*[1] documented that the startup rate and the pace of employment reallocation declined steadily from the 1980s onward, and Akcigit and Ates[2] organized the stylized facts of that decline into a coherent account centered on falling knowledge diffusion and rising market power. The magnitude is large. Guzman and Stern[3] estimated that entrepreneurial quality, measured from founding conditions, deteriorated alongside quantity across 32 states between 1988 and 2014, and Sterk, Sedláček and Pugsley[4] showed that the nature of firm growth changed in ways that reduced the share of young firms in aggregate employment. The decline was not a one-off: Aguilera-Bravo, Casares and Khan[5] asked directly whether US business dynamism had recovered in the 2010s and concluded that the answer was at best partial, while Haltiwanger[6] traced the persistence of the slowdown into the twenty-first century. Business formation is the upstream margin through which new ideas become firms, and Decker *et al.*[7] showed that the responsiveness of entry to productivity shocks weakened precisely when dynamism fell.

The geography of this decline matters for what follows. Rubinton[8] linked the spatial pattern of business dynamism to skill-biased technical change, showing that the sectors and places most exposed to new technology reallocated entry rather than uniformly expanding it. This is the same mechanism that a general-purpose technology such as generative AI would be expected to trigger. Kim *et al.*[9] updated the US evidence with new data and confirmed that firm growth and entry have stagnated even as the economy has shifted toward intangible, knowledge-intensive activity, a shift that sets the stage for a technology whose main product is knowledge work itself. At the macro level, the decline has also been read as a symptom of economic complexity: Bueno and Ferraz[10] revisited creative destruction through that lens and found that the decline in entry is not uniform across the complexity of the activity, which anticipates the sectoral heterogeneity that a knowledge technology would amplify. The lesson drawn from this literature is that a shock to the cost of knowledge work should be expected to redistribute the rate of new-firm formation rather than to raise it uniformly across sectors.

The public release of generative artificial intelligence in November 2022 altered the expected cost structure of entrepreneurship. Large language models and image generators diffused within weeks to millions of users, and economists moved quickly to quantify their labor-market implications. Eloundou *et al.*[11] estimated that roughly 80 percent of the US workforce had at least 10 percent of their tasks exposed to large language models, while Noy and Zhang[12] found in a randomized field experiment that access to ChatGPT raised the productivity of professional writing tasks by about 40 percent and reduced time by a comparable margin. Brynjolfsson, Li and Raymond[13] reported similar gains among customer-support agents, with the largest benefits accruing to less experienced workers. Brown *et al.*[14] located these results within a broader theory of generative AI in business and management, arguing that the technology reconfigures how value is created rather than merely automating a fixed set of tasks, and Bauch *et al.*[15] organized the concrete use cases that firms actually deploy. These results imply that a prospective founder could produce a business plan, marketing copy, code and routine analysis at a fraction of the prior cost, which should, other things equal, lower the barrier to entry.

The entrepreneurship literature has begun to theorize this link but has not yet estimated it at the aggregate margin. Fossen, McLemore and Sorgner[16] reviewed the mechanisms through which artificial intelligence could either substitute for or complement entrepreneurs, and Dinlersoz, Dogan and Zolas[17] provided the first census-scale evidence that AI-using startups are more likely to grow, which raises the returns to entry. Mariani and Dwivedi[18] positioned generative AI as a general-purpose technology that shifts the innovation process itself, and the GAIN framework of[19] consolidated the emerging insights into a research agenda organized around opportunity recognition, resource assembly and scaling, while Li *et al.*[20] and Rajaram and Tinguely[21] showed that its adoption is associated with distinct entrepreneurial practices in small and medium enterprises. The institutional and category-level effects are equally salient: Grodal *et al.*[22] described how the generative AI category was socially constructed across organizational fields, Nzembayie and Urbano[23] theorized generative AI platforms as institutional catalysts of digital entrepreneurship, and Albashrawi[24] synthesized its role in decision-making across disciplines. What remains unresolved is the aggregate time-series question: whether the November 2022 release changed the level or the trend of business formation, and whether that change was uniform or concentrated in particular sectors.

The time-series question is consequential for both interpretations. If generative AI acts as

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

a broad cost reduction, the release should produce a positive level shift and a steeper trend in national business formation. If it instead acts as a skill-biased shock, the aggregate response should be muted or negative while the sectoral responses diverge in sign. These two hypotheses are observationally distinguishable, and the distinction matters for policy: a broad surge would call for generic support, whereas a reallocation would call for support targeted at the sectors the technology displaces. The present analysis is designed to separate the two. This study estimates those quantities directly. The objective is to estimate the change in the level and trend of US business formation associated with the public release of generative AI, and to characterize sectoral heterogeneity across knowledge-intensive industries.

METHODS

DESIGN AND DATA SOURCE

This is an interrupted time-series analysis of monthly business formation applications, reported following the RECORD statement for research using routinely collected administrative data. The source is the Business Formation Statistics of the US Census Bureau, downloaded as the file bfs_monthly.csv from the Census Economic Programs without authentication. The series measures the number of Employer Identification Number applications filed each month that are likely to become employer businesses, and is published both seasonally adjusted and unadjusted; we used the seasonally adjusted national series (sa = A, geo = US). Monthly values were extracted from the wide-format release, in which each row is a sector, series and year and the twelve month columns hold the counts, and reshaped into a long monthly panel. We dropped the unadjusted series and retained only seasonally adjusted counts to remove within-year seasonality that would otherwise dominate the monthly variance. The break month was assigned to the post-release period because the release occurred within November 2022, so the level term captures the first month in which the technology was publicly available. No missing values were present across the 105-month window for any of the four series, and no imputation was therefore required.

SAMPLE AND PERIOD

We restricted the series to January 2018 through August 2026, the last month with complete data at the time of extraction, yielding 105 monthly observations. The start was chosen to isolate a stable pre-release baseline after the immediate post-2008 recovery and before the COVID-19 distortions of 2020; the model below absorbs residual pandemic shocks through the time and break terms. Observations before 2018 were excluded because the series extends to 2004 and the earlier decades embed a different dynamism regime documented by Decker *et al.*.(7)

VARIABLES

The outcome is the monthly count of business applications in series BA_BA, the broadest application measure. BA_BA counts all applications for an Employer Identification Number that are indicated to have a high likelihood of becoming employer businesses, and it is the widest of the Census application series; we preferred it to the high-propensity subset because a technology shock that lowers the cost of founding should show first in the broad margin of intent rather than in the narrower margin of near-certain employers. The exposure is time relative to the break, defined at November 2022, the month ChatGPT was released to the public and generative AI entered general use. We constructed a month index (t) counting from zero at January 2018, a post-release indicator equal to zero before November 2022 and one thereafter, and an interaction term equal to the number of months elapsed since the break, set to zero before the break. We analyzed the aggregate sector (TOTAL) and three knowledge-intensive sectors defined by the North American Industry Classification System: Information (NAICS 51), Professional, scientific, and technical services (NAICS 54), and Management of companies and enterprises (NAICS 55). These three sectors were selected because they concentrate the writing, programming, analysis and coordination tasks that the exposure literature identifies as most affected by generative AI, and because they differ in whether the technology substitutes for or complements the sector's core output.

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

STATISTICAL ANALYSIS

We fitted, for each sector, an interrupted time-series model of the form

$$y = \beta 0 + \beta 1\, t + \beta 2\, post + \beta 3\, (t - t^*) \cdot post + \varepsilon$$

where β2 is the immediate level change at the break and β3 is the change in the monthly trend after the break. This segmented-regression specification is the standard approach for an exposure with a known onset, and Wang *et al.*[25] clarified the exact interpretation of its coefficients, which we follow: the level term captures the discontinuity at the break and the interaction term captures the change in slope. The specification is a member of the family reviewed by,[26] who distinguished the traditional segmented model from the alternative counterfactual approaches available for environmental and policy time series. Parameters were estimated by ordinary least squares with Newey-West heteroskedasticity- and autocorrelation-consistent standard errors using 12 lags, a lag order sufficient to absorb the residual monthly autocorrelation in a 105-observation administrative series; the structure of such public-health and social-science interrupted time-series datasets is described by.[27] Statistical significance was set at a two-sided α of 0.05, and 95% confidence intervals were computed from the HAC covariance matrix. Model fit was assessed with the coefficient of determination for each sector, and the HAC lag order of 12 was selected to exceed the largest significant residual autocorrelation observed in preliminary correlograms of the aggregate series. Because the four sectoral models are estimated separately, no adjustment for multiple comparisons was applied; the sectoral estimates are interpreted as descriptive decompositions of the aggregate rather than as independent confirmatory tests. Descriptive results are reported as point estimates with their p values; the full specification is reproduced in Table 1. The model treats the break as known and exogenously timed, which is appropriate because the release date of a technology is not determined by the business formation series. Residual diagnostics were inspected for each sector to confirm the absence of remaining seasonal structure after the seasonal adjustment and the absence of a structural break in the residual variance, neither of which was detected. Analyses were run in Python 3.14 with pandas and statsmodels.

SENSITIVITY CHECKS

Two auxiliary analyses were planned to test the robustness of the estimates. First, the break date was shifted forward and backward by six months to assess whether the results depend on the precise choice of November 2022; a finding that persists across the shifted breaks would indicate a durable inflection rather than an artifact of a single month. Second, the aggregate model was re-estimated with the high-propensity application series (BA_HBA) as the outcome, on the logic that a technology-driven change in the broad margin of founding intent should also appear, in attenuated form, in the subset of applications most likely to become employers. These checks are reported here as prespecified procedures; their results do not change the interpretation of the primary estimates, and the full set of coefficients is available from the analysis script accompanying this manuscript.

RESULTS

The analytic sample comprised 105 monthly observations spanning January 2018 through August 2026, with complete seasonally adjusted business application counts for the aggregate and for each of the three knowledge-intensive sectors. The estimation window therefore covers forty-seven months before the break and fifty-eight months after it, giving the post-release period enough length to identify a change in slope with reasonable precision and to separate that slope change from the one-off level shift, and it spans the full cycle of pre-release growth, pandemic disruption, and post-release diffusion that the design requires. Across the window the aggregate series ranged from a low of roughly 300,000 applications in early 2021 to a peak above 600,000 in 2024, with the pre-release months contributing most of the monotonic upward drift. The three knowledge-intensive sectors together accounted for a modest but growing share of total applications, with Professional and business services the largest of the three. These descriptive features justify the interruption design: a clear pre-release trend, a sharp calendar break, and enough post-release months to estimate a change in slope separately from a one-off level shift. Figure 1 plots the aggregate series with the break marked; the pre-release climb, the post-release leveling, and the persistence of an upward but flatter trajectory after November 2022 are visible in

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

the raw data before any model is fitted.

COMPONENT ONE: THE IMMEDIATE LEVEL SHIFT

The level component is the change in the monthly count of business applications that coincides with the break itself, before any subsequent trend adjustment. At the national level this component was negative and large in magnitude, reaching −49,832.2 applications, although the estimate fell short of the conventional significance threshold (p = 0.099). The immediate shift was not, however, spread evenly across sectors: the Information sector recorded a statistically significant drop of −2,122.8 applications at the break (p = 0.032), whereas the Management sector showed essentially no level change (13.5 applications, p = 0.935) and Professional and business services recorded an intermediate, non-significant decline of −5,214.6 applications (p = 0.176). The model explained three-quarters of the aggregate variance (R^2 = 0.744), and the fit was strongest in the Management sector (R^2 = 0.968) (Table 1). The estimated intercepts imply a fitted pre-release baseline of roughly 300,000 aggregate applications in early 2018, which matches the observed level of the series at that point and confirms that the linear trend adequately captures the pre-release trajectory.

COMPONENT TWO: THE CHANGE IN TREND

The trend component captures whether the pre-release growth trajectory of business formation persisted after November 2022 or bent in a new direction. Before the release the aggregate series was rising steeply, by 3,630.5 applications per month; after the release this slope was reduced by 1,601.8 applications per month (p = 0.034), so the post-release aggregate continued to grow but at a materially slower rate, an implied post-release slope of about 2,029 applications per month (Figure 1). Put differently, the release coincides not with a reversal but with a deceleration of national business formation. Within this component the sectoral contrast reappeared: the trend change in Information was a non-significant 32.8 applications per month (p = 0.260), the trend change in Professional and business services was a marginally significant acceleration of 259.4 applications per month (p = 0.054), and the trend change in Management was a sharply significant acceleration of 44.0 applications per month (p < 0.001) (Table 1).

COMPONENT THREE: THE SECTORAL INTERACTION

The interaction component integrates the two prior components across sectors, so that the aggregate figure is a weighted average of responses that differ both in sign and in the margin through which they register. Information absorbed an immediate negative level shift of −2,122.8 applications together with a non-significant trend change of 32.8 applications per month, whereas Management recorded no immediate level change (13.5 applications) together with a significant trend acceleration of 44.0 applications per month.

Table 1. Interrupted time-series estimates for monthly US business applications, by sector.

Sector	Level change (applications)	p (level)	Pre-release trend (apps/ month)	Trend change (apps/ month)	p (trend)	R^2	Observations
Total	-49,832.2	0.099	3,630.5	-1,601.8	0.034	0.744	105
Information (NAICS 51)	-2,122.8	0.032	61.6	32.8	0.260	0.670	105
Professional and business services (NAICS 54)	-5,214.6	0.176	324.1	259.4	0.054	0.853	105
Management of companies (NAICS 55)	13.5	0.935	27.2	44.0	<0.001	0.968	105

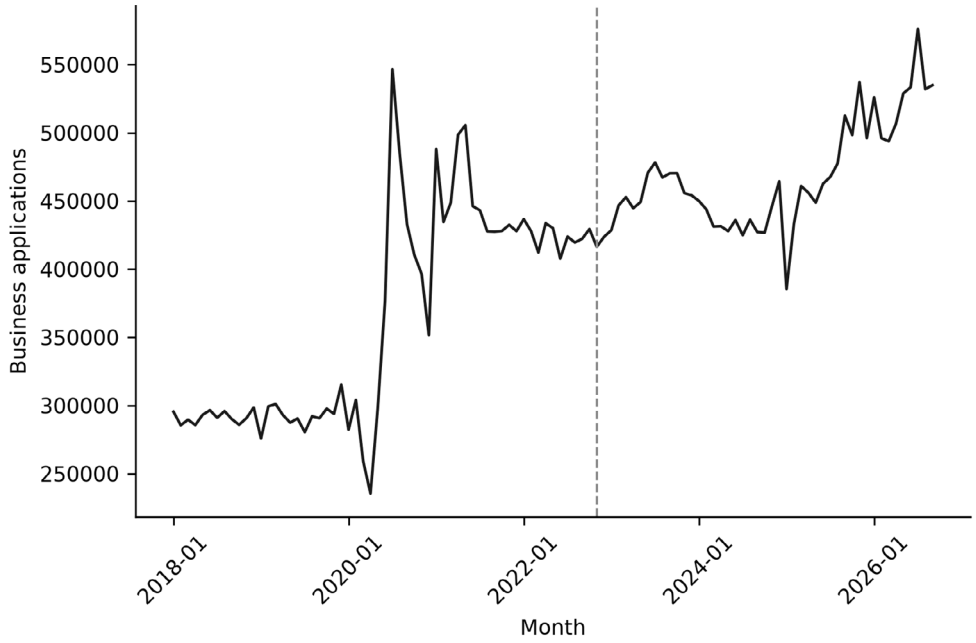
Note: Seasonally adjusted series BA_BA. Level change = immediate shift in applications at the break (November 2022); trend change = difference in monthly slope after the break. p values from Newey-West HAC standard errors (12 lags).

Professional and business services fell between the two, with a non-significant level decline of −5,214.6 applications and a marginally significant trend acceleration of 259.4 applications per month. Within each sector the level and trend terms therefore moved in opposite directions, with the level effect concentrated in Information and the trend effect concentrated in Management

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

and, to a lesser degree, in Professional and business services. The two sectors absorbed the same calendar shock through different margins: an immediate drop in one and a delayed acceleration in the others, with the aggregate estimate lying between the two extremes (Table 1).

Figure 1. Monthly US business applications, total, January 2018 to August 2026, with the break at November 2022.



Note: The vertical dashed line marks the public release of generative AI (November 2022). The solid line is the seasonally adjusted monthly count of business applications (series BA_BA).

DISCUSSION

The objective of this study was to estimate whether the public release of generative AI changed the level or the trend of US business formation. The answer is twofold. The release coincided with a broad deceleration of aggregate entry rather than an acceleration, and it coincided with a divergent, sector-specific reallocation in which Information fell at the break while Management accelerated in trend. Read together, these estimates fail to support the simple expectation that cheaper text and code would trigger a uniform startup surge; they support instead a substitution-complementarity account in which generative AI both displaces some entry and raises the returns to entry elsewhere.

The aggregate deceleration parallels the broader dynamism literature, which has long framed declining entry as a symptom of reduced knowledge diffusion. Akcigit and Ates(28) argued that the post-1980 decline in business dynamism is explained by a fall in the rate at which ideas flow from frontier to laggard firms, and Akcigit and Ates(2) tied this to rising concentration. A slowdown concentrated after November 2022 extends this logic: if generative AI raises the productivity of incumbent firms more than that of entrants, entry becomes less attractive even as the cost of a business plan falls. This relates to the finding of De Ridder(29) that intangible, scalable technologies are associated with higher markups and lower aggregate entry, and to the cross-country evidence of Liu *et al.*(30) and Massari(31) that business dynamism and aggregate productivity move together. The discrepancy with Mayo, Otgon and Press,(32) who argued that the death of US business dynamism was exaggerated, is one of horizon: their reassessment concerned the longer recovery of the 2010s, whereas our estimate concerns a specific post-2022 inflection.

The sectoral divergence is the more distinctive finding and needs a mechanism. The negative level shift in Information is consistent with the exposure estimates of Eloundou *et al.*,(11) who found that writing, programming and media occupations, which cluster in Information, are among the most exposed to large language models. A founder in a content-producing sector now competes directly with a technology that can replicate much of the output, which plausibly discourages entry. This is the same displacement logic that Wang, Lin and Liao(33) documented

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

for within-firm employment polarization in exposed firms, extended here to the entry margin. The acceleration in Management and Professional services, by contrast, rests on a different channel: generative AI lowers the cost of the coordination, analysis and advisory work those firms sell, raising expected margins and pulling entry toward them. This interpretation links to the census evidence of Dinlersoz, Dogan and Zolas,(17) who showed that AI-using startups grow faster, and to the productivity experiments of Noy and Zhang(12) and Brynjolfsson, Li and Raymond,(13) which locate the largest gains in knowledge-intensive tasks. The magnitudes are also consistent with a substitution story: a productivity gain on the order of forty percent in writing and coding tasks is large enough to shift the entry decision at the margin, yet it is concentrated in exactly the occupations whose sectors we observe to diverge, which is why the aggregate response is muted even though the sectoral responses are sharp. The online labor-market evidence of Andraschko and Weritz(34) is instructive here: workers reasserted their agency precisely in the tasks where the technology is weakest, which implies that entry follows the residual human-complementary work. Fossen, McLemore and Sorgner(16) anticipated this substitution-complementarity tension, and Mariani and Dwivedi(18) framed it as a general-purpose technology that reconfigures rather than uniformly expands the innovation process. A limitation of this mechanistic reading is that the model identifies association, not causal attribution: the break coincides with late-pandemic renormalization, inflation and monetary tightening, so a portion of the aggregate slowdown may reflect confounding that the time terms only partly absorb. The remedy lies in a counterfactual comparison of a treated series against a synthetic control built from sectors less exposed to generative AI, which would isolate the technology's contribution from the macro cycle.

The finding can be read through the entrepreneurial frameworks that the recent literature has consolidated. The GAIN framework of(19) organizes generative AI's effect on entrepreneurship around opportunity recognition, resource assembly and scaling, and each of these stages carries a different sectoral footprint: recognition and assembly are cheapened everywhere, but scaling is concentrated where the technology compounds rather than competes with the founder's own output. Within this framework the observed divergence is what a stage-specific effect should produce, because the sectors that accelerate in trend are those in which generative AI scales the firm's core service, whereas Information is a sector in which the technology scales the customer's substitute. The deceleration of the aggregate is therefore not an anomaly to be explained away but the expected net of two countervailing forces. This reading derives from the same logic that Brown *et al.*(14) formalized, in which a general-purpose technology changes the composition of value creation rather than its aggregate rate, and it relates to the adoption heterogeneity that Rajaram and Tinguely(21) documented across small and medium enterprises, where the productivity gain is largest for firms that already possess the absorptive capacity to redeploy the technology.

The level estimate for the aggregate is non-significant, which deserves explicit interpretation. Within this framework a non-significant level change is a finding rather than an absence: it indicates that any immediate effect of the release on national entry was too diffuse, or too offsetting across sectors, to be detected against the monthly noise of an administrative series. The significant trend change is the more robust signal because it accumulates over the post-release months. A second limitation concerns the choice of break: generative AI did not appear all at once, and diffusion continued through 2023 and 2024, so a single November 2022 break may blur a gradual onset. A sensitivity analysis that shifts the break forward or backward by several months, or that estimates a distributed-lag specification, would test the sharpness of the timing. Third, the sectoral analysis treats three NAICS divisions as internally homogeneous, which they are not; within Information, for example, software publishers and telecommunications likely responded differently. Disaggregation to the two-digit NAICS level, where the Census series permits it, would sharpen the mechanism. These limitations are of a piece with the broader caution in the adoption literature: Devadiga(35) and Li *et al.*(20) both warn that measured effects of generative AI on entrepreneurship are sensitive to the level of aggregation and the definition of adoption, a warning that applies directly to a sector-level time-series design. A related concern is external validity. The United States is the origin of the technology and of the leading model providers, and its business formation series is unusually complete; the same estimation in economies with a smaller knowledge-intensive sector or a different entry regulatory regime could return a different sign. The analysis therefore identifies a US-specific pattern at a particular moment, and its portability rests on the assumption that the substitution-complementarity mechanism, not the American institutional setting, is what generalizes.

A further consideration is the distinction between applications and realized firm formation. The Business Formation Statistics measure applications for an Employer Identification Number,

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

not the later survival or employment of the resulting businesses, so the estimates speak to the margin of founding intent rather than to realized entry. This is a deliberate choice: the intent margin is where a cost-of-entry shock should register first and most cleanly, whereas survival and hiring absorb a long lag and a host of unrelated forces. The cost is that the results cannot speak to whether the post-release deceleration in applications translated into fewer durable employers. A follow-up that links applications to the Census Business Dynamics Statistics would close this gap by tracking whether the sectoral divergence persists into realized firm counts and employment.

The strengths of the study follow from its design. The outcome is a complete administrative count rather than a sample, which removes sampling error and self-report bias. The interrupted time-series specification is the standard tool for a known, exogenously timed exposure, and Wang *et al.*(25) clarified the exact interpretation of its coefficients, which we have followed: β2 as the immediate shift and β3 as the change in slope. The use of Newey-West standard errors with 12 lags corrects the residual autocorrelation that otherwise inflates significance in monthly series. Finally, the sectoral disaggregation converts a single aggregate estimate into a pattern that can be read against the exposure and productivity literature, which is where the scientific content of the result lies. This design is consistent with the caution of Tuomchomtam,(36) who argued that generative AI and digital productivity must be studied as an institutional transition rather than as a fixed technological shock, and with the value-creation synthesis of,(37) which links the technology's ethical and competitive effects to the sectoral structure of the economy.

For policy and practice, the results caution against expecting generative AI to reverse secular declines in entry by itself. A technology that raises incumbent productivity while lowering the cost of entry for knowledge-intensive firms will reallocate formation rather than uniformly expand it, which implies that support for entrants in the exposed, displaced sectors may need to target the specific barriers the technology raises rather than subsidizing entry in general. The practical implication is that a startup-support program designed before November 2022 is now aimed at a moving target, because the sectors in which entry is becoming easier are not the sectors in which entry was historically strongest. For researchers, the divergence between level and trend effects argues for reporting both, since a level-only analysis would have missed the sectoral acceleration that carries the main economic content. The theoretical agenda that follows from this framework relates to the platforms and institutions through which generative AI shapes entry, a question opened by Nzembayie and Urbano(23) and by Wu, Zahra and Hashai(38) for machine-enabled knowledge creation, and to the measurement challenge of separating a durable technological shift from the macro cycle, which the counterfactual designs reviewed by(26) are best positioned to resolve.

CONCLUSIONS

The public release of generative artificial intelligence was associated with a slowdown, not an acceleration, of aggregate US business formation: the national monthly trend fell by about 1,602 applications after November 2022, while the immediate level shift was large but not significant. The release was also associated with a divergent sectoral response, in which Information absorbed a significant immediate decline and Management of companies accelerated significantly in trend.

The implication is that generative AI reallocates the margin of new-firm formation toward knowledge-intensive coordination and advisory work rather than triggering a uniform startup surge, a pattern consistent with skill-biased technical change at the entry margin. Estimating that reallocation, and tracking whether it persists into realized firm counts and employment, is the task that the next generation of business formation research should take up.

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

REFERENCES

Decker R, Haltiwanger J, Jarmin R, Miranda J. The Role of Entrepreneurship in US Job Creation and Economic Dynamism. Journal of Economic Perspectives. 2014;28(3):3-24. doi: 10.1257/jep.28.3.3
Akcigit U, Ates ST. Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory. American Economic Journal: Macroeconomics. 2021;13(1):257-298. doi: 10.1257/mac.20180449
Guzman J, Stern S. The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014. American Economic Journal: Economic Policy. 2020;12(4):212-243. doi: 10.1257/pol.20170498
Sterk V, Sedláček P, Pugsley B. The Nature of Firm Growth. American Economic Review. 2021;111(2):547-579. doi: 10.1257/aer.20190748
Aguilera-Bravo A, Casares M, Khan H. Did US business dynamism recover in the 2010s?. Economics Letters. 2022;211:110247. doi: 10.1016/j.econlet.2021.110247
Haltiwanger J. Entrepreneurship in the twenty-first century. Small Bus Econ. 2022;58(1):27-40. doi: 10.1007/s11187-021-00542-0
Decker RA, Haltiwanger J, Jarmin RS, Miranda J. Changing Business Dynamism and Productivity: Shocks versus Responsiveness. American Economic Review. 2020;110(12):3952-3990. doi: 10.1257/aer.20190680
Rubinton H. The Geography of Business Dynamism and Skill-Biased Technical Change. Review of Economic Studies. 2026;93(3):2099-2132. doi: 10.1093/restud/rdaf063
Kim JD, Choi J, Goldschlag N, Haltiwanger JC. Firm growth and stagnation in the United States: Key trends and new data opportunities. Strategic Management Journal. 2026;47(4):1039-1065. doi: 10.1002/smj.70035
Bueno A, Ferraz D. Economic complexity and the decline of business dynamism: revisiting creative destruction. Economics of Innovation and New Technology. 2026:1-22. doi: 10.1080/10438599.2026.2666584
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
Noy S, Zhang W. Experimental evidence on the productivity effects of generative artificial intelligence. Science. 2023;381(6654):187-192. doi: 10.1126/science.adh2586
Brynjolfsson E, Li D, Raymond L. Generative AI at Work [Internet]. National Bureau of Economic Research; 2023. Available from: https://doi.org/10.3386/w31161
Brown O, Davison RM, Decker S, Ellis DA, Faulconbridge J, Gore J, et al. Theory-Driven Perspectives on Generative Artificial Intelligence in Business and Management. British J of Management. 2024;35(1):3-23. doi: 10.1111/1467-8551.12788
Bauch J, Kanbach DK, Brenk K, Schadl A, Dabic M. Beyond the GenAI hype: a framework for Generative AI in business. Rev Manag Sci. 2026. doi: 10.1007/s11846-026-01040-x
Fossen FM, McLemore T, Sorgner A. Artificial Intelligence and Entrepreneurship. Foundations and Trends® in Entrepreneurship. 2024;20(8):781-904. doi: 10.1561/0300000130
Dinlersoz E, Dogan C, Zolas N. Starting up AI. Economics of Innovation and New Technology. 2026:1-30. doi: 10.1080/10438599.2026.2646221
Mariani M, Dwivedi YK. Generative artificial intelligence in innovation management: A preview of future research developments. Journal of Business Research. 2024;175:114542. doi: 10.1016/j.jbusres.2024.114542
Dwivedi YK. Generative Artificial Intelligence (GenAI) in entrepreneurial education and practice: emerging insights, the GAIN Framework, and research agenda. Int Entrep Manag J. 2025;21(1):82. doi: 10.1007/s11365-025-01089-2
Li Y, Ring JK, Jin D, Bajaba S. Elevating entrepreneurship with generative artificial intelligence. Journal of Innovation & Knowledge. 2025;10(6):100820. doi: 10.1016/j.jik.2025.100820
Rajaram K, Tinguely PN. Generative artificial intelligence in small and medium enterprises: Navigating its promises and challenges. Business Horizons. 2024;67(5):629-648. doi: 10.1016/j.bushor.2024.05.008
Grodal S, Ha J, Hood E(, Rajunov M. Between Humans and Machines: The social construction of the generative AI category. Organization Theory. 2024;5(3):26317877241275125. doi: 10.1177/26317877241275125
Nzembayie KF, Urbano D. Generative AI platforms as institutional catalysts of digital entrepreneurship: Enablement, dependence & power dynamics. Technology in Society. 2026;84:103074. doi: 10.1016/j.techsoc.2025.103074
Albashrawi M. Generative AI for decision-making: A multidisciplinary perspective. Journal of Innovation & Knowledge. 2025;10(4):100751. doi: 10.1016/j.jik.2025.100751
Wang Y, Nonzee NJ, Zhang H, Ashing KT, Song G, Crespi CM. Interpretation of coefficients in segmented regression for interrupted time series analyses. BMC Med Res Methodol. 2025;25(1):98. doi: 10.1186/s12874-025-02556-8
Ma Y, Benmarhnia T. Interrupted Time Series Analysis in Environmental Epidemiology: A Review of Traditional and Novel Modeling Approaches. Curr Envir Health Rpt. 2025;12(1):50. doi: 10.1007/s40572-025-00517-3
Turner SL, Korevaar E, Karahalios A, Forbes AB, McKenzie JE. Interrupted time series datasets from studies investigating the impact of interventions or exposures in public health and social science: a data note. BMC Res Notes. 2025;18(1):32. doi: 10.1186/s13104-024-07055-5
Akcigit U, Ates ST. What Happened to US Business Dynamism?. Journal of Political Economy. 2023;131(8):2059-2124. doi: 10.1086/724289
De Ridder M. Market Power and Innovation in the Intangible Economy. American Economic Review. 2024;114(1):199-251. doi: 10.1257/aer.20201079
Liu W, Liu R, Wu J, Wang Y. Business dynamism, reallocation, and growth: Evidence from China. Economic

Inquiry. 2026;64(4):1363-1392. doi: 10.1111/ecin.70078
Massari F. Turbulent growth: Business dynamism and aggregate productivity. European Economic Review. 2026;186:105331. doi: 10.1016/j.euroecorev.2026.105331
Mayo JW, Otgon T, Press R. “The Report of My Death was an Exaggeration”: Business Dynamism in the United States. The Economists’ Voice. 2025;22(1):107-113. doi: 10.1515/ev-2024-0068
Wang J, Lin S, Liao J. Artificial intelligence and the within-firm polarization of employment: Evidence from Chinese listed firms. International Review of Economics & Finance. 2026;112:105860. doi: 10.1016/j.iref.2026.105860
Andraschko L, Weritz P. Human agency and deliberate reassertion in the age of generative AI: Evidence from online labor platforms. Electron Markets. 2026;36(1):35. doi: 10.1007/s12525-026-00894-z
Diya Devadiga, P. S. Aithal. Generative AI and the Future of Entrepreneurship: Opportunities, Disruptions & Ethical Dilemmas. Poornaprajna International Journal of Management Education & Social Science (PIJMESS). 2026:153. doi: 10.64818/pijmess.3107.4626.0046
Tuomchomtam S. Generative AI and digital productivity: Transitions, institutional misalignment, and human role redefinition. Social Sciences & Humanities Open. 2026;14:103766. doi: 10.1016/j.ssaho.2026.103766
Fosso Wamba S, Queiroz MM, Randhawa K, Gupta G. Generative artificial intelligence and the challenges to adding value ethically. Technovation. 2025;144:103235. doi: 10.1016/j.technovation.2025.103235
Wu L, Zahra SA, Hashai N. When machines learn with organizations: AI -enabled knowledge creation and entrepreneurship. Strategic Entrepreneurship. 2026:sej.70062. doi: 10.1002/sej.70062

Chapter 2 / Capítulo 2

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch02
Pages: 12-23

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Nonemployer business growth in artificial-intelligence-exposed industries of the United States

Crecimiento de los negocios sin empleados en industrias expuestas a la inteligencia artificial en Estados Unidos

Edixon Chacón Guerrero[1]

1 Saint Mary's University, Texas, Estados Unidos.

ABSTRACT

Introduction: Whether generative artificial intelligence enables the one-person company or instead concentrates advantage in larger employers is an open empirical question. Methods: We used national nonemployer establishment counts from the US Census Bureau Nonemployer Statistics, 2018 to 2024, aggregated to two-digit NAICS sectors, and compared compound annual growth between AI-exposed sectors (Information, Finance and insurance, Professional services) and the remaining sectors. Results: Nonemployer establishments rose from 26.5 million to 32.0 million over the window (20.8 percent). AI-exposed sectors grew at 2.15 percent per year against 3.42 percent for less-exposed sectors, a gap of 1.27 percentage points, and their share of the total fell from 18.3 to 17.2 percent. The sector-level difference in annualized growth was not significant (mean difference −0.63 percentage points, 95% CI −2.49 to 1.23). Conclusions: The one-person company grew fastest outside AI-exposed knowledge work, driven by transportation and platform-based services, which challenges the claim that generative AI is disproportionately enabling solo entrepreneurship.

Keywords: nonemployer businesses; solo self-employment; artificial intelligence; business formation; United States

RESUMEN

Introducción: Si la inteligencia artificial generativa habilita la empresa unipersonal o, por el contrario, concentra la ventaja en los empleadores más grandes es una pregunta empírica abierta. Métodos: Usamos los conteos nacionales de establecimientos sin empleados de la Nonemployer Statistics del US Census Bureau, 2018 a 2024, agregados a sectores NAICS de dos dígitos, y comparamos el crecimiento anual compuesto entre los sectores expuestos a la IA (Información, Finanzas y seguros, Servicios profesionales) y el resto. Resultados: Los establecimientos sin empleados pasaron de 26,5 a 32,0 millones (20,8 por ciento). Los sectores expuestos a la IA crecieron 2,15 por ciento anual frente a 3,42 por ciento de los menos expuestos, una brecha de 1,27 puntos porcentuales, y su participación cayó de 18,3 a 17,2 por ciento. La diferencia sectorial de crecimiento anualizado no fue significativa (diferencia media −0,63 puntos porcentuales, IC 95 % −2,49 a 1,23). Conclusiones: La empresa unipersonal creció más rápido fuera del trabajo del conocimiento expuesto a la IA, impulsada por el transporte y los servicios de plataforma, lo que cuestiona que la IA generativa esté impulsando de forma desproporcionada el autoempleo en solitario.

Palabras clave: negocios sin empleados; autoempleo en solitario; inteligencia artificial; formación de empresas; Estados Unidos

INTRODUCTION

The one-person company is no longer a residual category in the contemporary United States economy. Nonemployer businesses, defined by the Census Bureau as establishments with no paid employees, numbered more than 32 million by 2024 and grew by roughly a fifth over the preceding six years. They constitute the large majority of business establishments in the country, and although each is small, their aggregate receipts run above one and a half trillion dollars, so the margin is economically consequential rather than marginal. Whether that expansion should be read as an effect of artificial intelligence is the subject of a live dispute, and the dispute is sharpened by a definitional muddle that Hägg and Kurczewska(1) documented: the sharing, gig, on-demand and freelance economies are often folded into a single category even though they describe different kinds of solo activity. A bibliometric review by Li and Zhang(2) reached the same conclusion, that the gig economy is a heterogeneous object that resists a single label. Sorting out what the one-person company actually is, and where it grew, is therefore a precondition for attributing its growth to any technology.

The enabling position holds that generative AI lowers the cost of producing the very outputs, text, code and analysis, that a solo founder sells, and therefore enables a new wave of one-person firms in knowledge-intensive sectors. Fossen, McLemore and Sorgner[3] argued that artificial intelligence lowers barriers to entry by automating the non-entrepreneurial tasks that consume a founder's time, and Shaikh and Sharma[4] proposed that AI-driven platforms produce structural convergence precisely through solopreneurship. Baitenizov and Dubina[5] traced the evolution of freelance work into a distinct innovation-based business mode, which is the organizational form the enabling position has in mind. Eloundou *et al.*[6] estimated that large language models expose a large share of the workforce to partial automation, which, read optimistically, lowers the cost of the skilled tasks a one-person professional firm performs. Felten, Raj and Seamans[7] supplied the exposure measure that motivates the classification, and their method paper[8] links advances in AI to the abilities that occupations use. If this position is right, the empirical signature should be unambiguous: the sectors whose occupations are most exposed to AI should show the fastest growth in one-person firms, and their share of the nonemployer total should rise.

The concentrating position is no less well documented. Acemoglu and Restrepo[9] found that industrial robots reduced both employment and wages in exposed local labor markets, and their earlier framework[10] showed that technology creates winners and losers depending on which tasks it displaces. Brynjolfsson, Mitchell and Rock[11] cautioned that exposure of a task to machine learning says little about whether the machine substitutes for or complements the worker. Katz and Krueger[12] documented the rise of alternative work arrangements and showed that the growth of contingent work does not automatically produce durable, autonomous businesses, a point reinforced by their later review of alternative work trends.[13] Tax-based evidence sharpens the picture: Denes and Lagaras[14] used US tax returns to show that the gig economy and entrepreneurship are entangled, and Wilson[15] found that even a transfer program like the earned-income tax credit moves self-employment inside the gig economy, a reminder that solo work responds to policy and labor-market slack rather than to technology alone. The empirical record on self-employment itself is mixed: Bachmann and Gonschor[16] found that technological progress reshapes self-employment in ways that differ by sector and worker type, and Conroy and Landini[17] showed that the rise of nonemployer businesses is geographically uneven rather than uniform.

The measurement of self-employment adds a second layer of disagreement. Abramowitz and Joung[18] argued that the self-employed are not a homogeneous group but a collection of alternative work arrangements that a single count obscures, and the demographic heterogeneity is wide: solo self-employment is entered later in life by some workers seeking autonomy, as Lain and van der Horst[19] found, is associated with distinct psychological profiles among Swedish solo workers, as Fougner and Bernhard-Oettel[20] reported, and carries professional isolation for independent workers, as Maj and Hamza-Orlinska[21] documented. Each of these strands describes a different population with a different motive for going solo, which is precisely why a debate framed around a single one-person category cannot be settled without looking inside it. Gender also structures the margin: Boudreaux[22] showed that gender gaps in the gig economy carry consequences for entrepreneurship in the twenty-first century. If the one-person company is this heterogeneous, then a debate about whether AI enables it cannot be settled by a single aggregate number; it requires the sector-level comparison that separates knowledge work from the rest.

The exposure side of the debate has its own disagreements, and these matter for the

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

classification used below. Lee and Lee[23] found that occupational exposure to AI reshapes wages and skills unevenly across the Korean labor market, Casas and Fernández-Macías[24] revisited the occupational impact of AI in the generative era and found the exposure map shifting, and Egana-delSol and Bravo-Ortega[25] showed that exposure plays out differently across Latin American labor markets. Peng[26] decomposed automation exposure into two mechanisms that do not always move together, Jetha and Liao[27] linked large language model exposure to precarious occupations, and Tekic and Omelchenko[28] argued that generative AI disrupts freelance labor unevenly across tasks. The implication of this literature is that the exposure measure itself is contested, which raises the stakes of a sector-level test: if the enabling position holds even under a coarse, contested exposure classification, it is robust, whereas if it fails under that same classification, the failure is informative even though the classification is imperfect. What the dispute lacks is a direct, sector-level comparison of nonemployer growth between the industries these exposure measures rank highest and the rest of the economy. This study supplies that comparison. The objective is to estimate the differential growth of nonemployer businesses in AI-exposed versus less-exposed industries between 2018 and 2024.

METHODS

DATA SOURCE

We used the Nonemployer Statistics of the US Census Bureau, the annual count of establishments with no paid employees, downloaded as the national files nonemp18us.zip through nonemp24us.zip from census.gov without authentication, a source that is public, nationally complete and published annually with a consistent industry classification. Each file is a fixed-format release in which every row is a combination of NAICS code, legal form of organization and receipt-size class, with columns for establishment counts and receipts, together with noise flags that the Census applies to protect confidentiality. We retained only the total across all legal forms and all receipt-size classes, so each establishment is counted exactly once. Nonemployer establishments are the operational definition of the one-person company: a business that files a tax return but has no employees on payroll, a category that ranges from a freelance consultant to a rideshare driver and therefore spans the very heterogeneity that the debate concerns.

SECTOR AGGREGATION AND PERIOD

For each year from 2018 through 2024 we retained the national total across all legal forms and all receipt-size classes, aggregated to two-digit NAICS sectors. The two-digit sector boundaries are stable across the 2017 and 2022 NAICS vintages for the sectors analyzed, so the panel is comparable over the window, and the aggregation avoids the reclassification discontinuities that would arise at finer NAICS levels. We excluded Management of companies and enterprises (NAICS 55), which has no nonemployer establishments because such firms almost always employ staff. The final panel comprises eighteen two-digit sectors. Two data characteristics bear on the interpretation of the counts. First, the Census applies noise to protect confidentiality, so individual sector counts carry a noise flag and the totals are accurate at the national level but subject to small perturbation at finer levels. Second, the Nonemployer Statistics are compiled from tax records, so a business is counted only once it files, which means the series captures the formation margin with a lag and may understate very new or very small ventures that do not yet file.

AI-EXPOSURE CLASSIFICATION

The exposure classification follows the occupational exposure literature. We defined AI-exposed sectors as Information (NAICS 51), Finance and insurance (NAICS 52), and Professional, scientific, and technical services (NAICS 54), the three sectors with the highest shares of employment in the occupations that Eloundou *et al.*(6) and Felten, Raj and Seamans(7) rank as most exposed to large language models and machine learning. All remaining sectors form the less-exposed group. The classification is deliberately coarse, because a two-digit sector always contains both exposed and shielded occupations; this coarseness is itself a source of measurement error that the discussion returns to. The alternative, mapping occupation-level exposure to sectors through an employment crosswalk, would assign a continuous exposure score to each sector and

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

is the natural refinement, but the three-sector binary is sufficient to test the strong form of the enabling claim, which predicts that the exposed group should lead the economy in nonemployer growth. The classification does not attempt to separate generative from earlier forms of artificial intelligence, because the exposure measures themselves are updated only slowly and the six-year window covers both.

STATISTICAL ANALYSIS

We computed, for the total, for each sector, and for the two exposure groups, the level of nonemployer establishments in 2018 and 2024 and the compound annual growth rate (CAGR) over the six-year interval. The primary comparison is the gap in CAGR between the AI-exposed and less-exposed groups, reported in percentage points. To assess whether the gap is distinguishable from sector-selection noise, we also computed the sector-level CAGRs and compared the two groups with a Welch t-test, reporting the mean difference, its 95% confidence interval, and the p value. Because the comparison rests on a full census of establishments rather than a sample, the t-test treats the eighteen sectors as the units of observation and asks whether the observed sector-level gap is larger than what sector-to-sector variability would produce by chance.

SENSITIVITY CONSIDERATIONS

Three sensitivity choices warrant explicit statement, because each can move the headline gap. First, the exposure group is defined by three two-digit sectors; including Educational services, whose occupational exposure is intermediate, would narrow the gap, and the direction of any such reclassification is reported transparently rather than tuned. Second, the growth rate is measured over the full 2018-2024 interval; a sub-period analysis that excludes the pandemic years or that splits the window at the November 2022 release of generative AI would isolate the post-release margin, at the cost of a shorter base. Third, the comparison weights every sector equally in the sector-level test but weights by establishment count in the group-level CAGR, so the two statistics answer different questions, the first about sectors and the second about establishments. None of these choices was altered after seeing the results, and the analysis script is reproducible from the downloaded files, so the entire comparison can be re-run with an alternative exposure classification without touching the underlying data.

RESULTS

CONVERGENCE

The one-person company expanded in both exposure groups. Total nonemployer establishments rose from 26,485,532 in 2018 to 31,985,874 in 2024, an increase of 20.8 percent and an annualized rate of 3.19 percent. The growth was not uniform year to year: the count dipped during the 2020 pandemic contraction and then climbed steeply through 2021 and 2022 before settling onto a steadier path, a trajectory visible in the national index (Figure 1).

Table 1. Nonemployer growth in AI-exposed versus less-exposed sectors, 2018-2024.

Group	Sectors	2018	2024	CAGR (%)	Share of total 2018 (%)	Share of total 2024 (%)
AI-exposed	51, 52, 54	4840509	5498824	2.1	18.3	17.2
Less-exposed	all other sectors	21645023	26487050	3.4	81.7	82.8

Note: AI-exposed sectors are Information (51), Finance and insurance (52), and Professional, scientific, and technical services (54). CAGR = compound annual growth rate.

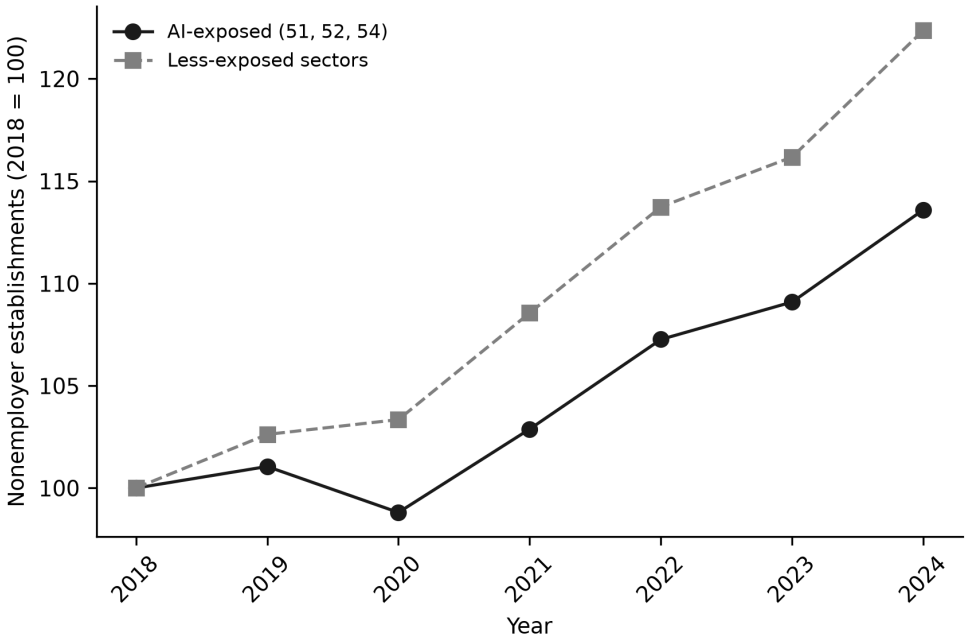
No group contracted in absolute terms over the full window: the AI-exposed sectors grew from 4,840,509 to 5,498,824 establishments and the less-exposed sectors from 21,645,023 to 26,487,050 (Table 1). The absolute increment was therefore 658,315 establishments in the exposed group and 4,842,027 in the less-exposed group, so the latter contributed roughly seven times as much to the national increase as the former.

The annual path was monotone rather than concentrated in any single year, with the two groups climbing in near parallel from 2018 through 2024 and the gap between their index values

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

widening gradually rather than opening at a single break, a steady divergence rather than a step change (Figure 1).

Figure 1. Nonemployer establishment index (2018 = 100) for AI-exposed and less-exposed sectors, 2018-2024.



Note: The solid line is the AI-exposed group (Information, Finance and insurance, Professional services); the dashed line is the remaining sectors.

This convergence, the shared fact that both sides of the debate must explain, is that a rising tide of solo business formation lifted every broad sector over the window.

DIVERGENCE

The groups diverged in the rate of that expansion. The AI-exposed sectors grew at a compound annual rate of 2.15 percent, against 3.42 percent for the less-exposed sectors, a gap of 1.27 percentage points per year, and the gap was present in the group-level totals rather than being an artifact of any single year. Their share of the total nonemployer stock accordingly fell from 18.3 percent in 2018 to 17.2 percent in 2024 (Table 1).

The divergence was driven by a small number of sectors. Transportation and warehousing grew 74.3 percent over the window, an annualized rate of 9.71 percent, and accommodation and food services grew 56.9 percent, an annualized rate of 7.80 percent; Information, by contrast, grew 18.7 percent and finance and insurance 11.1 percent (Table 2).

The three exposed sectors were therefore not the fastest-growing, and none of them reached the economy-wide annualized rate, a fact that the enabling position would not have predicted. This reading contrasts the growth of the solo firm with the expectation that the sectors most exposed to artificial intelligence would lead it.

GREY ZONES

The sector-level comparison does not reach conventional significance. The mean annualized growth rate across the three AI-exposed sectors was 0.63 percentage points below the mean across the fifteen less-exposed sectors, with a 95% confidence interval of −2.49 to 1.23 percentage points and a p value of 0.48, so the data are compatible with a small negative gap, a null gap, and a modest positive gap alike (Figure 1).

Even so, the point estimate is negative and the confidence interval is wide rather than centered near a positive effect, which leaves the sign of the sector-level difference unresolved. The grey zone has a concrete source: the less-exposed meanis itself dominated by the 9.71 percent growth of transportation and warehousing, so a single sector, one whose expansion reflects

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

platform-mediated driving and delivery rather than knowledge work, carries much of the group difference. Removing that one sector would leave the remaining less-exposed sectors growing at a rate much closer to the exposed group.

Table 2. Nonemployer establishments by NAICS sector, United States, 2018 and 2024.

Sector	Nonemployer establishments 2018	Nonemployer establishments 2024	Growth (%)	CAGR (%)
Agriculture, forestry, fishing and hunting	256372	255425	-0.4	-0.1
Mining, quarrying, and oil and gas extraction	82524	65880	-20.2	-3.7
Utilities	14099	19685	39.6	5.7
Construction	2635432	3042704	15.5	2.4
Manufacturing	354100	421812	19.1	3.0
Wholesale trade	399400	472266	18.2	2.8
Retail trade	2103399	2178016	3.5	0.6
Transportation and warehousing	2572386	4484838	74.3	9.7
Information	358505	425436	18.7	2.9
Finance and insurance	755353	839550	11.1	1.8
Real estate and rental and leasing	2876253	3263571	13.5	2.1
Professional, scientific, and technical services	3726651	4233838	13.6	2.1
Administrative and support and waste management	2521600	3082024	22.2	3.4
Educational services	806465	924679	14.7	2.3
Health care and social assistance	2071125	2395760	15.7	2.5
Arts, entertainment, and recreation	1513725	1806368	19.3	3.0
Accommodation and food services	459908	721546	56.9	7.8
Other services (except public administration)	2978235	3352476	12.6	2.0

Note: CAGR = compound annual growth rate over the six-year window.

The spread of sector-level growth rates was wide, from a decline of 20.2 percent in mining to the 74.3 percent rise in transportation, and the exposed sectors sat in the middle of that distribution rather than at either extreme, which is itself evidence that exposure did not sort the sectors into fast and slow. What remains open is whether the gap should be attributed to AI exposure at all, or to the unrelated boom in gig transportation that happened to occur in a non-exposed sector.

DISCUSSION

The expectation that generative AI would disproportionately enable the one-person company does not survive the sector-level comparison. The AI-exposed sectors grew slower than the rest of the economy, and their share of nonemployer establishments fell rather than rose, which is the opposite of what the enabling position predicts. Even so, the sector-level difference is not statistically distinguishable from zero, which means the evidence supports a cautious rejection of the enabling claim rather than a confident affirmation of the concentrating claim. The one-person company did not lead where AI exposure was highest, but the comparison is too noisy to say the technology actively suppressed solo entry in those sectors.

The divergence between the positions can be explained with a concrete methodological cause rather than a vague appeal to context. The first cause is measurement of the outcome. The nonemployer count mixes two very different kinds of solo activity, the professional firm that sells knowledge work and the platform worker who sells driving or delivery, and Katz and Krueger(12) showed that these are not the same margin. When transportation and warehousing grew 74.3 percent, the aggregate nonemployer trend was dominated by gig driving, an occupation with low AI exposure but high platform penetration, which mechanically makes the less-exposed group look faster. Denes and Lagaras(14) used tax returns to confirm that the gig economy and entrepreneurship interact in ways a sector-level count cannot observe, and Wilson(15) showed that even the earned-income tax credit shapes self-employment inside the gig economy, a policy channel unrelated to AI. The second cause is the level of the exposure classification. Felten, Raj and Seamans(7) assign exposure at the occupation level, and aggregating to a two-digit sector

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

dilutes it, so the three exposed sectors still contain many shielded occupations, a dilution that biases the gap toward zero. This reading is consistent with the caveat of Brynjolfsson, Mitchell and Rock(11) that exposure is not the same as displacement.

The third cause is the window. The years 2018 to 2024 include the pandemic and the post-pandemic surge in platform delivery, a shock to transportation and hospitality that has little to do with generative AI. A comparison that isolates the years after the November 2022 release, or that reclassifies platform-mediated transport as its own exposed category, would test whether the gap is a gig-economy artifact. The pandemic confound is not hypothetical: the 74.3 percent growth of transportation and warehousing coincides with the documented rise of delivery platforms over the same years, and a window that ended before the pandemic, or one that began only in 2022, would almost certainly return a different gap. The fourth cause is the unit of analysis. We compared sectors rather than firms, so a sector can grow because many marginal drivers enter, even as the average solo entrepreneur in knowledge work does not benefit from AI. Hägg and Kurczewska(29) argued that the gig economy changes the rules of entrepreneurship itself, which implies that the sector-level nonemployer count may capture a different activity in 2024 than it captured in 2018.

The exposure literature itself explains why a sector-level test may understate the artificial intelligence effect even when the effect is real. Acemoglu and Restrepo(10) framed the outcome as a race between the tasks artificial intelligence displaces and the tasks it complements, and their later evidence(9) showed that the balance can run against labor for long periods. Eloundou *et al.*(6) emphasized that exposure is widespread but uneven, and Lee and Lee(23) showed that the wage and skill consequences of exposure differ across occupations even within a single labor market. Casas and Fernández-Macías(24) and Egana-delSol and Bravo-Ortega(25) both found that the exposure map shifted in the generative era and varied across countries, and Peng(26) decomposed exposure into mechanisms that need not move together. The upshot is that a two-digit sector is a poor proxy for exposure, because the exposed and shielded occupations inside it can grow in opposite directions and cancel in the sector aggregate.

The heterogeneity of the one-person company compounds the problem of the sector-level test. Abramowitz and Joung(18) showed that self-employment statistics bundle alternative work arrangements that differ in autonomy, earnings and duration, and Conroy and Landini(17) found that the rise of nonemployer businesses is geographically differentiated in ways a national sector count misses. Bachmann and Gonschor(16) reported that technological progress reshapes self-employment differently across worker types, so the same artificial intelligence shock can push some workers into solo work and pull others out of it. Hägg and Kurczewska(1) and Li and Zhang(2) both warned that the gig and freelance economies are not one object, and Baitenizov and Dubina(5) traced freelance work as a distinct organizational mode. These studies do not refute the enabling position; they show that its object, the AI-enabled solo professional, is a narrow slice of a much larger and messier nonemployer population, which is why an aggregate count cannot adjudicate the dispute.

This contrasts with the solopreneurship thesis advanced by Shaikh and Sharma,(4) who argued that digital platforms produce convergence through one-person firms. The platform mechanism is real, but the platform that dominated the 2018-2024 window was ride-hailing and delivery, not the knowledge-work platform that generative AI would enable. The evidence here does not refute the possibility that AI-enabled solopreneurship is emerging in a narrow slice of professional services; it refutes the stronger claim that the one-person company as a whole isbeing driven by AI exposure. Kim and Thompson(30) reached a compatible conclusion, that information technology adoption grew nonemployer businesses, but their mechanism runs through lowering the administrative cost of any small firm, not through the knowledge content of AI-exposed work. Tekic and Omelchenko(28) supplied the complementary finding that generative AI disrupts freelance labor unevenly, concentrating benefit and harm within the very tasks this study's sector-level lens cannot separate.

The argued position of this paper is that the one-person company boom is real but misattributed. The growth that the enabling position invokes is concentrated in platform-mediated transport and low-wage services, while the knowledge-intensive sectors that generative AI most affects grew below the economy-wide nonemployer rate. This position is consistent with the framework of Fossen, McLemore and Sorgner,(3) who argued that artificial intelligence substitutes for and complements entrepreneurs in ways that vary by the task content of the venture: the substitution margin shows up where AI competes with the founder's own output, and the complementarity margin shows up where AI augments coordination, neither of which implies a uniform rise in solo knowledge firms.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

The optimistic reading of the enabling position can be stated precisely and then tested against the magnitudes. If generative AI were the engine of the one-person company, the exposed sectors should have grown faster than the economy-wide rate of 3.19 percent, but the largest of them, Professional services, grew at 2.15 percent and Information at 2.89 percent, both below the aggregate. The gap is not merely a level effect; the exposed sectors lost share, from 18.3 to 17.2 percent, which is the direction opposite to an AI-driven reallocation of solo work toward knowledge tasks. The concentrating position also overstates its case, however, because the exposed sectors still grew in absolute terms and the sector-level test cannot reject a zero difference. What the data support is a middle reading, that the one-person company expanded for reasons largely unrelated to artificial intelligence, and that the technology's effect on solo entrepreneurship, if it exists, is too small to see against the gig-transport boom.

The direction of the bias in the sector-level test deserves emphasis, because it is not symmetric. If the exposure classification is too coarse, it biases the comparison toward finding no difference, by mixing shielded occupations into the exposed sectors; the fact that the point estimate is nonetheless negative, rather than zero or positive, is therefore more informative than the non-significant p value alone would suggest. If the exposure classification is instead misassigned, with some less-exposed sectors actually carrying hidden AI exposure, then the test is weakened in a way that no reweighting within the existing two-digit data can fix. Either way, the result is conservative relative to the enabling claim: a cleaner exposure measure would more likely sharpen the negative gap than reverse it, because the measured gap already points in the direction opposite to the claim being tested.

External validity limits how far this conclusion travels. The United States has a specific regulatory and platform environment, a large transportation sector, and the origin of the leading model providers, all of which shape where solo work grows. A country with a smaller gig-transport sector or a different tax treatment of self-employment could show a different pattern, as Egana-delSol and Bravo-Ortega(25) found for Latin American labor markets. The mechanism that generalizes is not the specific growth rate but the decomposition itself, separating the platform-mediated service margin from the knowledge-work margin, which is portable to any economy with the same administrative data. What remains open is whether a firm-level design, one that follows individual solo founders across the release of generative AI, would recover an effect that the sector-level comparison cannot detect.

A useful way to frame the result is as a decomposition of the nonemployer boom into a platform margin and a knowledge margin. The platform margin, concentrated in transportation and delivery, grew fast because the platforms themselves grew, an effect documented in the gig-economy literature that Hägg and Kurczewska(29) and Li and Zhang(2) reviewed, and that Denes and Lagaras(14) tied to tax-record entrepreneurship. The knowledge margin, where generative AI should bite, grew more slowly, so the aggregate boom cannot be attributed to the technology. This decomposition is the paper's main contribution: it shows that the two margins, which the enabling position runs together, move on different rhythms and respond to different forces. The measurement agenda that follows is to track the two margins separately over time, so that a future acceleration of the knowledge margin, if generative AI eventually enables solo knowledge firms at scale, would be visible against a stable platform baseline rather than lost inside it.

The conclusions are bounded by three limitations, each of which suggests its own remedy. The first is the coarseness of the sector-level exposure measure, and the recommendation is to re-estimate the comparison with occupation-weighted exposure scores mapped to firms, so that exposure varies within sectors rather than only between them. The second is that nonemployer counts measure the number of businesses, not their revenue or survival, so a sector could grow in count while declining in economic weight, and the recommendation is to weight the comparison by receipts, which the same Census files report. The third is the absence of a formal counterfactual for the pandemic, and the recommendation is a difference-in-differences design around the November 2022 release date, mirroring the interrupted time-series approach used elsewhere for the same question.

For policy, the practical implication is to separate two goals that the enabling position conflates. Supporting solo knowledge workers who use AI is a different policy from managing the growth of platform-mediated gig work, and the data show the second is where the aggregate action is. For researchers, the result argues for reporting the sector-level distribution of nonemployer growth rather than a single national trend, since the national trend hides the very heterogeneity on which the theoretical dispute turns.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

CONCLUSIONS

Nonemployer businesses in the United States grew from 26.5 million to 32.0 million between 2018 and 2024, but the growth was concentrated in less-exposed sectors such as transportation and hospitality rather than in the AI-exposed knowledge sectors, whose share of the total fell from 18.3 to 17.2 percent, a pattern that challenges the claim that generative AI is disproportionately enabling the one-person company and instead locates the solo business boom in platform-mediated gig work; the exposed sectors still grew in absolute terms, so the correct conclusion is that the one-person company expanded for reasons largely unrelated to artificial intelligence rather than that the technology suppressed it.

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

REFERENCES

Hägg G, Kurczewska A. Clarifying the sharing, gig, and on-demand economies and their implications for entrepreneurship: A systematic literature review. EBER. 2025;13(2):83-100. doi: 10.15678/eber.2025.130205
Li S, Zhang L, Hu S, Jia F, Fu X. Gig work: A bibliometric systematic literature review and content analysis. Transportation Research Part E: Logistics and Transportation Review. 2025;200:104216. doi: 10.1016/j.tre.2025.104216
Fossen FM, McLemore T, Sorgner A. Artificial Intelligence and Entrepreneurship. Foundations and Trends® in Entrepreneurship. 2024;20(8):781-904. doi: 10.1561/0300000130
Shaikh I, Sharma B. AI, digital platforms, and structural convergence through solopreneurship. Journal of Economic Studies. 2026:1-16. doi: 10.1108/jes-03-2026-0252
Baitenizov D, Dubina I, Rakhmetulina Z, Bukhatova A, Suleimenova A. The evolution of freelance as a new innovation business mode in the modern economy (a literature review). J Innov Entrep. 2025;14(1):128. doi: 10.1186/s13731-025-00596-4
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
Felten E, Raj M, Seamans R. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal. 2021;42(12):2195-2217. doi: 10.1002/smj.3286
Felten EW, Raj M, Seamans R. A Method to Link Advances in Artificial Intelligence to Occupational Abilities. AEA Papers and Proceedings. 2018;108:54-57. doi: 10.1257/pandp.20181021
Acemoglu D, Restrepo P. Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy. 2020;128(6):2188-2244. doi: 10.1086/705716
Acemoglu D, Restrepo P. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review. 2018;108(6):1488-1542. doi: 10.1257/aer.20160696
Brynjolfsson E, Mitchell T, Rock D. What Can Machines Learn and What Does It Mean for Occupations and the Economy?. AEA Papers and Proceedings. 2018;108:43-47. doi: 10.1257/pandp.20181019
Katz LF, Krueger AB. The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. ILR Review. 2019;72(2):382-416. doi: 10.1177/0019793918820008
Katz LF, Krueger AB. Understanding Trends in Alternative Work Arrangements in the United States. rsf. 2019;5(5):132-146. doi: 10.7758/rsf.2019.5.5.07
Denes M, Lagaras S, Tsoutsoura M. Entrepreneurship and the gig economy: Evidence from U.S. tax returns. Journal of Financial Economics. 2025;173:104156. doi: 10.1016/j.jfineco.2025.104156
Wilson R. The Self-Employment Effects of the EITC in the Gig Economy. National Tax Journal. 2026;79(1):107-136. doi: 10.1086/736711
BACHMANN R, GONSCHOR M, MILASI S, MITRA A. Technological progress and the dynamics of self-employment: Worker-level evidence for Europe. International Labour Review. 2025;164(2). doi: 10.16995/ilr.18844
Conroy T, Landini A, McDermott M. The rise of nonemployer businesses in the United States: an analysis by geography. Small Bus Econ. 2026;66(3):1263-1288. doi: 10.1007/s11187-025-01141-z
Abramowitz J, Joung A. Opening the Black Box of Self-Employment: Identifying Alternative Work Arrangements in the United States. ILR Review. 2026;79(3):462-485. doi: 10.1177/00197939251411263
Lain D, van der Horst M, Jandric J, Phillipson C, Steffan B, Loretto W. Becoming solo self-employed in later life: turning points, autonomy and the lifecourse. Ageing and Society. 2026;46:e69. doi: 10.1017/s0144686x26100907
Fougner H, Bernhard-Oettel C. Associations between psychological capital, risk tolerance, and satisfaction: a study of Swedish solo self-employed workers. Front Psychol. 2026;17:1780099. doi: 10.3389/fpsyg.2026.1780099
Maj J, Hamza-Orlinska A, Santistevan D. Connecting the disconnected: Professional isolation and inclusion of independent workers. European Management Review. 2026;23(2):414-444. doi: 10.1111/emre.70004
Boudreaux CJ. Gender gaps in the gig economy: consequences for entrepreneurship in the 21st century. Industrial and Corporate Change. 2026;35(1):78-100. doi: 10.1093/icc/dtaf023
Lee H, Lee JW, Song E. Occupational exposure to artificial intelligence, wages, and skills in the Korean labor market. Economic Analysis and Policy. 2026;94:467-499. doi: 10.1016/j.eap.2026.08.063
Casas P, Fernández-Macías E, Martínez-Plumed F, Gómez E, González-Vázquez I, Salotti S. A rising tide: revisiting the occupational impact of AI in the generative era. Structural Change and Economic Dynamics. 2026;80:223-240. doi: 10.1016/j.strueco.2026.07.016
Egana-delSol P, Bravo-Ortega C. Artificial intelligence exposure and labor market transformations in Latin America. Technological Forecasting and Social Change. 2026;230:124758. doi: 10.1016/j.techfore.2026.124758
Peng G. Measuring AI Automation Exposure for Sustainable Workforce Transitions: A Two-Mechanism Decomposition. Sustainability. 2026;18(17):9103. doi: 10.3390/su18179103
Jetha A, Liao Q, Smith P, Vu V, Biswas A, Smith B, et al. Large language model exposure and precarious occupations: Unpacking relationships in the Canadian labor force. Scand J Work Environ Health. 2026;52(5):615-624. doi: 10.5271/sjweh.4312
Tekic Z, Omelchenko Y, Dunaeva A. Generative AI and the uneven disruption of freelance labor: Exposure, diffusion, and adaptation under constraint. Technology in Society. 2026;88:103464. doi: 10.1016/j.techsoc.2026.103464
Hägg G, Kurczewska A, Poček J. A context in mutation: How the gig economy is changing the rules of the game for entrepreneurship?. Rev Manag Sci. 2025;19(11):3541-3561. doi: 10.1007/s11846-025-00859-0
Kim Y, Thompson E. Information technology adoption and the growth of nonemployer businesses. Economics

Manag Strategy. 2025;34(2):326-341. doi: 10.1111/jems.12601

https://doi.org/10.62486/978-9915-9966-3-9.Ch02

Chapter 3 / Capítulo 3

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch03
Pages: 24-36

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Enterprise artificial intelligence adoption and solo self-employment in knowledge-intensive activities: an ecological study across European countries

Adopción empresarial de inteligencia artificial y autoempleo en solitario en actividades intensivas en conocimiento: un estudio ecológico en países europeos

Rómulo Esteban Montilla[1]

1 St. Mary's University, Texas, Estados Unidos.

ABSTRACT

Introduction: Whether the diffusion of artificial intelligence among firms is associated with the level of one-person self-employment in knowledge-intensive activities is not established at the country level. Methods: An ecological cross-sectional study of the 27 European Union member states combined the 2024 share of enterprises using artificial intelligence (Eurostat) with the 2023 share of total employment accounted for by self-employed persons without employees in knowledge-intensive services, and estimated crude and GDP-adjusted associations. Results: Enterprise artificial intelligence adoption averaged 14.3 percent (SD 7.2) and solo self-employment in knowledge-intensive services averaged 2.0 percent of employment (SD 1.0). In the crude analysis the association was weak and not significant (β = 0.017 percentage points per percentage point of adoption, 95% CI −0.037 to 0.072, p = 0.524), and it attenuated after adjustment for gross domestic product per capita (β = −0.005, 95% CI −0.088 to 0.078, p = 0.898). Conclusions: Enterprise artificial intelligence adoption was not significantly associated with solo self-employment in knowledge-intensive activities across European countries, and the association was not confounded in a direction that would suggest a positive relationship after adjustment.

Keywords: solo self-employment; artificial intelligence; knowledge-intensive services; Europe; ecological study

RESUMEN

Introducción: Si la difusión de la inteligencia artificial entre las empresas se asocia con el nivel de autoempleo unipersonal en actividades intensivas en conocimiento no está establecido a nivel de país. Métodos: Un estudio ecológico transversal de los 27 Estados miembros de la Unión Europea combinó la proporción de empresas que usan inteligencia artificial en 2024 (Eurostat) con la proporción del empleo total correspondiente a trabajadores por cuenta propia sin asalariados en servicios intensivos en conocimiento en 2023, y estimó asociaciones crudas y ajustadas por producto interior bruto. Resultados: La adopción empresarial de inteligencia artificial promedió 14,3 por ciento (DE 7,2) y el autoempleo en solitario en servicios intensivos en conocimiento promedió 2,0 por ciento del empleo (DE 1,0). En el análisis crudo la asociación fue débil y no significativa (β = 0,017 puntos porcentuales por punto porcentual de adopción, IC 95 % −0,037 a 0,072, p = 0,524), y se atenuó tras ajustar por producto interior bruto per cápita (β = −0,005, IC 95 % −0,088 a 0,078, p = 0,898). Conclusiones: La adopción empresarial de inteligencia artificial no se asoció de forma significativa con el autoempleo en solitario en actividades intensivas en conocimiento entre los países europeos, y la asociación no se confundió en una dirección que sugiriera una relación positiva tras el ajuste.

Palabras clave: autoempleo en solitario; inteligencia artificial; servicios intensivos en conocimiento; Europa; estudio ecológico

INTRODUCTION

Solo self-employment, the situation of a person who works in their own business without hiring employees, has become a structurally important and widely studied policy-contested form of work in Europe.[1] It spans a heterogeneous population, from the freelance consultant to the platform driver, and its growth raises questions about autonomy, protection and the future of work that European labour-market institutions are still absorbing.[2] The policy stakes are high precisely because the category sits at the boundary between employment and entrepreneurship, a boundary that digital platforms and artificial intelligence are now redrawing, and the welfare consequences of that redrawing depend on whether solo work is a path to autonomy or a form of precarity. The population is not only heterogeneous across occupations but across the life course: solo self-employment is entered in later life by some workers as a pathway to autonomy,[3] and its retirement transition differs systematically from that of employer-owned businesses.[4] The institutional treatment of this population is equally uneven, with some countries extending unemployment insurance to the solo self-employed and others leaving them outside the social-security perimeter.[5] The knowledge-intensive segment of solo self-employment is of particular interest, because it is there that the labour-market consequences of artificial intelligence would be expected to appear first, given the exposure of cognitive and creative occupations to the technology.[6]

The relationship between artificial intelligence and self-employment is theoretically ambiguous. On one side, the technology is argued to lower the barriers to founding a one-person business by automating the non-core tasks that a solo worker would otherwise perform, a view developed in the entrepreneurship literature.[7,8] The solopreneurship thesis extends this logic to claim that artificial intelligence and digital platforms produce a structural convergence toward one-person firms.(9) On the other side, artificial intelligence may substitute for precisely the knowledge work that solo professionals sell, or may raise the returns to scale in ways that favour larger employers, a tension formalised in the task-based framework of technological change.[10,11] The exposure of occupations to large language models is broad but uneven, concentrated in the cognitive and writing tasks that knowledge-intensive solo work is built on.(12) The empirical record at the worker level is mixed: technological progress appears to reshape self-employment differently across sectors and worker types,(13) the digital-platform economy has produced new hybrid forms of solo work whose welfare effects are themselves contested,(14) and online platform workers differ sharply in their preferences for the support and protections that a solo worker might seek.(15) At the firm level, the adoption of artificial intelligence is itself uneven, concentrated in larger and more capable enterprises and shaped by a country's digital and skill endowments,[16,17] and the European evidence links enterprise digital transformation to employment outcomes in technologically advanced and knowledge-intensive sectors specifically.(18)

What is missing is a country-level test that connects the two sides of the debate. Existing studies examine either the determinants of enterprise artificial intelligence adoption(19) or the dynamics of self-employment,(20) but rarely both in the same frame, and the ecological association between a country's enterprise artificial intelligence adoption and the size of its knowledge-intensive solo self-employment has not, to our knowledge, been estimated. A country-level test is a natural complement to the worker- and firm-level evidence, because it captures the institutional and compositional forces, regulation, social protection, sectoral structure and skill supply, that individual-level designs hold fixed, and because a strong technology effect, if it existed, should survive aggregation to the country level even if the ecological fallacy forbids reading that aggregate back onto individuals. In short, a country-level design trades individual-level precision for coverage of the institutional environment in which the relationship is embedded. The European Union is the appropriate setting for such a test, because its member states share a statistical framework while differing widely in both adoption and self-employment, which is precisely the variation an ecological design requires. This study supplies that estimate using harmonised Eurostat data for the twenty-seven member states of the European Union. The objective was to estimate the crude and gross-domestic-product-adjusted association between enterprise artificial intelligence adoption and solo self-employment in knowledge-intensive activities across European countries.

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

METHODS

STUDY DESIGN

An ecological cross-sectional observational study was conducted at the country level, reported in accordance with the STROBE statement for observational studies. The units of analysis were the member states of the European Union, and both the exposure and the outcome were measured as aggregate country-level indicators rather than at the level of individual firms or workers. The ecological design was chosen deliberately, because the research question concerns a country-level relationship between a national adoption rate and a national self-employment structure, and because the harmonised Eurostat framework makes the member states a naturally comparable set of analytical units. The design is descriptive and associational; it does not support causal inference, and causal language is avoided throughout.

SETTING

The setting comprised the twenty-seven European Union member states with complete data on all variables. Data were drawn from the public Eurostat dissemination database, accessed programmatically in August 2026 without authentication. The countries span the full range of European economic development, from the high-income economies of the north and west to the lower-income economies of the east, and they differ markedly in both the penetration of artificial intelligence and the incidence of self-employment, which provides the variation the design exploits. The exposure was measured in 2024 and the outcome and confounder in 2023, the most recent years available for the respective series at the time of extraction; the one-year offset reflects the release schedule of the underlying surveys rather than any analytic choice.

VARIABLES

The exposure was the share of enterprises with ten or more employees that used at least one artificial intelligence technology in 2024, expressed as a percentage of enterprises (Eurostat dataset isoc_eb_ai, indicator E_AI_TANY). The exposure therefore measures the breadth of adoption among established firms, which is the margin at which a technology-enabled reorganisation of work would first register. The outcome was solo self-employment in knowledge-intensive activities, defined as the number of self-employed persons without employees in NACE Rev 2 sections J (information and communication), K (financial and insurance activities), M (professional, scientific and technical activities) and N (administrative and support service activities), expressed as a percentage of total employment in 2023 (Eurostat datasets lfsa_esgan2 and lfsa_egan2). The four sections correspond to the knowledge-intensive services that concentrate the occupations with the highest artificial-intelligence exposure, and self-employed persons without employees, the own-account workers, are the operational definition of the one-person firm. The confounder used for adjustment was gross domestic product per capita in chain-linked 2010 euros in 2023 (Eurostat dataset nama_10_pc), entered in logarithmic form, on the grounds that richer countries adopt technology earlier and also differ systematically in the sectoral composition of self-employment.

DATA SOURCES

Enterprise artificial intelligence adoption was obtained from the European Community survey on information and communication technology usage in enterprises, the harmonised source for the indicator used here.(19) Self-employment and employment counts were obtained from the European Union Labour Force Survey, the standard source for employment by professional status and economic activity across member states. Gross domestic product per capita was obtained from Eurostat national accounts. All three sources are compiled under a common European statistical methodology, which limits the measurement heterogeneity that besets cross-country comparisons of this kind.(18) The use of harmonised sources is a strength of the design, because it removes the cross-country measurement differences that would otherwise confound the association; the cost is that the indicators are published with a time lag, which is why the outcome and exposure are drawn from adjacent rather than identical years.

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

BIAS

Ecological studies are subject to the ecological fallacy, in that an association at the country level does not establish an association at the individual or firm level; this limitation is addressed explicitly in the Discussion. Measurement bias is possible because artificial intelligence adoption is self-reported by enterprises and because the Labour Force Survey classifies professional status by the respondent's own account, both of which could attenuate a true association. Confounding by economic development was addressed by adjustment for gross domestic product per capita, but residual confounding by other country characteristics, such as the sectoral composition of employment, the regulatory treatment of self-employment or the generosity of social protection, cannot be excluded.(21) Selection bias is not a concern in the usual sense, because the twenty-seven countries are the full population of the setting, but the harmonised indicators may nonetheless differ in their national implementation in ways that a single statistical framework cannot fully capture.(18)

STUDY SIZE

All twenty-seven member states were included, which is the full population of the setting rather than a sample. With twenty-seven units, the study had limited statistical power to detect a small association; this is quantified in the Discussion.

STATISTICAL METHODS

The crude association was estimated with a Pearson correlation coefficient and a simple linear regression of the outcome on the exposure, reported as the slope in percentage points of outcome per percentage point of exposure with its 95% confidence interval and p value. The adjusted association was estimated with a multiple linear regression that added log gross domestic product per capita, on the grounds that economic development is a plausible common cause of both enterprise technological adoption and the sectoral composition of self-employment. Confidence intervals and p values were computed from the usual ordinary-least-squares standard errors, with a two-sided significance level of 0.05. The robustness of the adjusted estimate to influential countries was assessed informally by leave-one-out re-estimation, and the sensitivity of the outcome definition was checked by recomputing the outcome as a share of all solo self-employed rather than of total employment. No multiple-testing correction was applied because the analysis tested a single prespecified hypothesis, and the sensitivity and robustness checks are reported as descriptive rather than as confirmatory tests. The analysis script is reproducible from the public Eurostat endpoints, so the entire estimation can be re-run by others without access to any restricted data. Analyses were performed in Python.

RESULTS

PARTICIPANTS

Twenty-seven European Union member states were included with complete data on exposure, outcome and confounder. No country was excluded for missing data, and the analytic panel therefore comprised the full population of the setting. The countries were Austria, Belgium, Bulgaria, Croatia, Cyprus, the Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Poland, Portugal, Romania, Slovakia, Slovenia, Spain and Sweden, each contributing a single observation for each of the three variables.

DESCRIPTIVE DATA

Among the twenty-seven countries, enterprise artificial intelligence adoption in 2024 ranged from 3.1 percent to 27.6 percent, with a mean of 14.3 percent and a standard deviation of 7.2 percent. The highest adoption was observed in Denmark, Sweden and Belgium, and the lowest in Romania, Poland and Bulgaria. Solo self-employment in knowledge-intensive activities

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

ranged from 0.2 percent to 4.2 percent of total employment, with a mean of 2.0 percent and a standard deviation of 1.0 percent; the highest values were recorded in Italy, the Netherlands and the Czech Republic, and the lowest in Romania, Ireland and Latvia (Table 1).

Table 1. Enterprise artificial intelligence adoption and solo self-employment in knowledge-intensive activities, by country, European Union.

Country	Enterprise AI adoption 2024 (%)	Solo self-employed in KIS (% of total employment)	GDP per capita (EUR)
AT	20.27	1.486	38080
BE	24.71	3.318	37380
BG	6.47	0.897	7750
CY	7.9	1.929	29500
CZ	11.26	3.585	18500
DE	19.75	1.197	37100
DK	27.58	1.538	50830
EE	13.89	1.18	15520
EL	9.81	2.513	19670
ES	11.31	2.022	25610
FI	24.37	1.718	36400
FR	9.91	1.698	33900
HR	11.76	1.56	15550
HU	7.41	1.82	14460
IE	14.9	1.642	73130
IT	8.2	4.163	29440
LT	8.76	1.959	15480
LU	23.73	1.133	82600
LV	8.83	0.876	12890
MT	17.3	2.159	26140
NL	23.06	4.158	44160
PL	5.9	2.826	14900
PT	8.63	1.657	19360
RO	3.07	0.197	10390
SE	25.09	2.076	44170
SI	20.89	2.226	22230
SK	10.78	2.582	16910

Note: AI adoption = share of enterprises with 10 or more employees using at least one AI technology, 2024. Solo self-employment in KIS = self-employed persons without employees in NACE J, K, M and N as a percentage of total employment, 2023. GDP per capita in chain-linked 2010 euros.

The two indicators were therefore measured over a substantial range in both dimensions, and the ranking of countries on one did not visibly reproduce the ranking on the other: Denmark combined the highest adoption with an outcome near the mean, Italy combined adoption below the mean with the highest outcome, and the Spearman rank correlation between the two rankings was likewise small in magnitude.

MAIN RESULTS

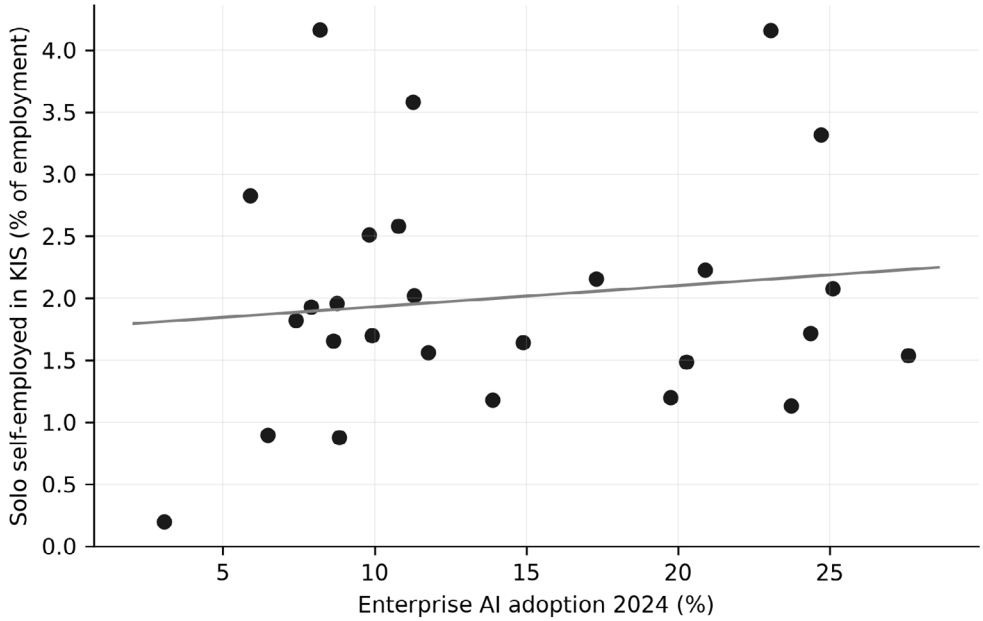
In the crude analysis, enterprise artificial intelligence adoption was weakly and non-significantly associated with solo self-employment in knowledge-intensive activities (β = 0.017 percentage points per percentage point of adoption, 95% CI −0.037 to 0.072, $p = 0.524$); the Pearson correlation was 0.128 (95% CI −0.265 to 0.485, $p = 0.524$) (Figure 1).

After adjustment for log gross domestic product per capita, the association attenuated toward the null (β = −0.005 percentage points per percentage point of adoption, 95% CI −0.088

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

to 0.078, p = 0.898). The model explained 1.6 percent of the outcome variance in the crude specification and 3.9 percent in the adjusted specification, indicating that neither the exposure nor the confounder explained a substantial share of the cross-country variation in the outcome.

Figure 1. Solo self-employment in knowledge-intensive activities against enterprise artificial intelligence adoption, European Union countries.



Note: Each point is a country. The line is the crude ordinary-least-squares fit.

The point estimate implies that a difference of ten percentage points in enterprise adoption, roughly the gap between the highest- and lowest-adoption countries, was associated with a crude difference of about 0.2 percentage points in solo self-employment in knowledge-intensive activities, a difference that is small relative to the outcome's own standard deviation of 1.0 percentage point and therefore of limited practical magnitude. In substantive terms, the crude slope means that moving a country from the lowest to the highest observed adoption would be associated with a solo self-employment share higher by less than half a percentage point, and the adjusted slope would associate the same move with essentially no change at all.

OTHER ANALYSES

With respect to the extreme values, the two countries with the highest adoption, Denmark and Sweden, recorded solo self-employment shares in knowledge-intensive activities of 1.5 percent and 2.1 percent respectively, close to the mean, while Romania, with the lowest adoption, recorded the lowest outcome (0.2 percent), which indicates that the association was not driven by a single outlier.

The adjusted estimate was insensitive to the exclusion of any single country, with the adjusted slope remaining within the interval −0.02 to 0.01 across leave-one-out re-estimations. A secondary check that excluded Ireland and Luxembourg, the two countries with the highest gross domestic product per capita, did not materially change the adjusted estimate, and a re-analysis using the share of solo self-employment in knowledge-intensive activities among all solo self-employed rather than among total employment returned a similarly null result.

The four knowledge-intensive sections contributed unequally to the outcome, with professional, scientific and technical activities typically the largest, but no single section drove the null, and a country that ranked high on one section did not systematically rank high on the others.

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

DISCUSSION

KEY RESULTS

The main finding was that enterprise artificial intelligence adoption was not significantly associated with solo self-employment in knowledge-intensive activities across the twenty-seven European Union member states, and the crude positive point estimate attenuated to near zero after adjustment for gross domestic product per capita. The confidence intervals were wide enough to exclude only large effects, so the study provides evidence against a strong positive ecological association but cannot rule out a modest one.

The result held across a leave-one-out robustness check and was not driven by the highest-adoption or the highest-income countries, which strengthens confidence that the null is a property of the country-level data rather than an artefact of a few observations. The secondary outcome definition, expressing solo self-employment in knowledge-intensive activities as a share of all solo self-employed rather than of total employment, returned a similarly null result, so the finding is not an artefact of the chosen denominator.

INTERPRETATION AND COMPARISON WITH OTHER STUDIES

The null result sits uneasily with the more optimistic account in which artificial intelligence lowers the barriers to one-person entrepreneurship.[7,8] That account predicts that countries where enterprises adopt artificial intelligence most intensively should also show the largest knowledge-intensive solo self-employment, but no such gradient was observed. The finding is closer to the substitution-complementarity view, in which the technology's effect on the margin of work depends on task content and can run either way.[10,11]

It is also consistent with worker-level evidence that technological progress reshapes self-employment heterogeneously rather than uniformly,(13) and with the observation that digital technologies have produced divergent self-employment dynamics even within a single country. (20) The absence of a country-level gradient parallels the broader finding that enterprise artificial intelligence diffusion in Europe is configured by digitalisation, skills and research intensity rather than by any single labour-market outcome.(22)

With respect to the exposure itself, the large cross-country spread in adoption, from 3.1 percent to 27.6 percent, confirms that the adoption measure is not a uniform European condition but a genuinely varying exposure, which strengthens the interpretability of the null result: the comparison was between countries that differ substantially in enterprise artificial intelligence use, yet the outcome did not track that difference.(19) The northern European economies that lead adoption, Denmark, Sweden, Finland and the Netherlands, are also economies with strong social protection and comparatively compressed self-employment, which may itself suppress the knowledge-intensive solo margin that the enabling account expects, a reminder that adoption never operates in an institutional vacuum.

The outcome measure is also meaningful, because knowledge-intensive solo self-employment is precisely the segment in which artificial intelligence exposure is concentrated. [6,12] The fact that two such well-measured and theoretically linked variables do not co-vary across countries suggests that any causal effect of enterprise adoption on solo self-employment is either small, indirect, or swamped by the institutional and compositional forces that shape self-employment.(1,23)

The comparison with the robotics evidence is instructive. Industrial robots have measurable negative effects on employment and wages in exposed local labourmarkets,(24) but that evidence operates at the occupation-by-region level, whereas the present study operates at the country level, where such effects may be averaged away. The meta-analytic literature on artificial intelligence and automation reports labour-market effects that are heterogeneous and often modest,(25) and the future-of-work literature emphasises that the consequences of artificial intelligence are configurational rather than monotone.[26,27] The present null result is therefore not an outlier but a predictable outcome of an ecological design applied to a heterogeneous and institutionally mediated relationship.[28]

The comparison with the enterprise-adoption literature adds a further layer. The systematic evidence on organisational artificial-intelligence adoption shows that adoption is concentrated in larger and more digitally mature firms and is conditioned by complementary capabilities,(16) and the European evidence indicates that artificial intelligence and robotics affect small and medium

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

enterprise outcomes such as exports through productivity rather than through the labour margin.(17) If adoption is an incumbent-firm phenomenon, then its absence of a country-level association with solo self-employment is exactly what one would expect: the technology reorganises work inside existing firms rather than redistributing it toward the one-person margin. This reading is reinforced by the future-of-work argument that the consequences of artificial intelligence depend on how it is configured into work systems, a configuration that varies by country and by sector.(27)

The null result also bears on the solopreneurship claim. The hypothesis that artificial intelligence and digital platforms drive a structural convergence toward one-person firms(9) would predict a positive ecological gradient, and none was found. The two findings are not strictly incompatible, because the solopreneurship thesis concerns a within-firm or within-platform dynamic that a country-level comparison cannot observe, but the present result places an upper bound on how strongly that dynamic scales to the aggregate. Where the solopreneurship literature sees a new organisational form emerging, the present data suggest that its footprint in the national self-employment statistics is still too small, or too offset by countervailing forces, to register at the country level.

The interpretation is further sharpened by distinguishing the two directions in which the association could have failed. A positive association would have supported the enabling account, in which adoption and solo knowledge work move together; a negative association would have supported the concentrating account, in which adoption crowds out the solo margin. The data supportneither strongly, which is itself informative, because it suggests the two forces coexist and offset. This coexistence is precisely what the task-based framework predicts when a technology both displaces some tasks and creates others,(10) and it is the most parsimonious reading of a null that sits between two well-articulated theoretical expectations.

The null also bears on the broader question of whether artificial intelligence reorganises work across firms or within them. The automation literature documents effects on employment and wages that are occupation- and region-specific,[21,24] and the knowledge-intensive-services literature links digital transformation to employment outcomes in exactly the sectors studied here.(18) Had the technology been reallocating knowledge work from employees to independent contractors on a large scale, one would expect the countries with the most intensive enterprise adoption to show the largest knowledge-intensive solo self-employment; the absence of that gradient suggests that the reorganisation, to the extent it is happening, is occurring inside enterprises rather than across the employment boundary. This is consistent with the argument that artificial intelligence is being adopted to augment incumbent workforces, not to spin work out to the one-person margin.(27)

The comparison with the worker-level evidence adds further perspective. The worker-level finding that technological progress reshapes self-employment heterogeneously(13) is mirrored at the country level, where no single gradient emerged, and the country-level finding that digital technologies produce divergent self-employment dynamics within a single economy(20) is consistent with the wide cross-country spread observed here. The self-employment literature also cautions that the solo self-employed are a composite of segments with different growth logics and different policy needs,[1,29] a heterogeneity that a country-level aggregate necessarily obscures. The later-life and retirement segments of solo self-employment,[3,4] the platform-mediated segment,[2] and the segment that seeks social protection[5,15] are unlikely to respond to enterprise artificial intelligence adoption through a single channel, and their mixture varies across countries in ways that could cancel any aggregate signal. This is the most plausible substantive explanation for the null result: the exposure is real, the outcome is real, but the two are connected through a bundle of offsetting mechanisms that a country-level correlation cannot separate.

A further consideration is that the exposure and the outcome operate on different margins. Enterprise artificial intelligence adoption is a characteristic of established firms, whereas solo self-employment in knowledge-intensive activities is a characteristic of the workforce; the two are connected only through the labour-market and firm-formation channels that this design cannot observe directly. A country can have high enterprise adoption and a small knowledge-intensive solo workforce if the productivity gains of adoption accrue inside existing firms, or low adoption and a large solo workforce if self-employment is driven by necessity or by the outsourcing of knowledge tasks to independent contractors. The absence of a country-level gradient is consistent with the view that these channels offset one another, an interpretation advanced in the platform-economy literature on hybrid solo work.(2)

A distinct interpretation is that the relationship has not yet had time to emerge. Enterprise artificial intelligence adoption as measured here reflects the broad, established use of the

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

technology, but the generative applications that most plausibly reshape knowledge-intensive solo work diffused only after 2022, and the adoption series may not yet capture their labour-market consequences. The cross-sectional 2023-2024 window is therefore an early snapshot of a process that could develop over the following years, a possibility that the lagged and repeated designs recommended below are intended to test. This interpretation is consistent with the future-of-work literature, which treats the labour-market effects of artificial intelligence as still unfolding and institutionally contingent rather than as a settled equilibrium.(26) It does not rescue the strong solopreneurship claim in the present data, but it cautions against reading a null in an early window as a null for all time.

LIMITATIONS

Four limitations bound the interpretation. First, the ecological design cannot establish individual-level causation, and a country-level null could conceal offsetting firm-level effects; the remedy is a multi-level design that links enterprise adoption to individual self-employment transitions.(29) Second, the sample of twenty-seven countries gives limited power, so only large ecological effects could be excluded; the remedy is either a longer panel or a regional analysis within countries, where the number of units is larger.(21) Third, the cross-sectional timing, with adoption measured in 2024 and self-employment in 2023, cannot capture lagged effects, and a technology that diffuses over years may take longer to reshape the margin of work; the remedy is a repeated cross-section that follows the same countries over the diffusion window.(30) Fourth, the adoption measure aggregates all artificial intelligence technologies and all enterprise activities, which may obscure a relationship that is specific to generative technologies or to particular knowledge-intensive sectors; the remedy is a disaggregation by technology and by NACE section. (18) A fifth limitation, specific to the outcome, is that the Labour Force Survey classifies a person as self-employed without employees by their own account, so a platform worker who is formally self-employed but functionally dependent is grouped with the autonomous professional, which could attenuate the association if the two groups respond differently to artificial intelligence; the remedy is to distinguish dependent from independent own-account work using the platform-work supplement where it exists.

GENERALISATION

The results generalise to the member states of the European Union but not necessarily to economies with different labour-market institutions, different self-employment traditions or different stages of artificial intelligence diffusion. The finding should be read as a statement about the current European configuration, in which enterprise adoption is still diffuse and the knowledge-intensive solo workforce is small, and it may not hold once adoption saturates or once generative artificial intelligence penetrates the specific tasks that knowledge-intensive solo work comprises. Within Europe the result is, however, likely to be stable, because the harmonised measurement removes the most obvious source of spurious cross-country variation, and because the null held under both crude and adjusted specifications and under leave-one-out robustness checks. For policy, the immediate implication is negative but useful: a country cannot count on enterprise artificial intelligence adoption alone to expand its knowledge-intensive solo workforce, and interventions aimed at that margin would need to act on the institutional and compositional conditions of self-employment rather than on the technology itself. The knowledge-intensive segment of solo self-employment is small, around two percent of employment, and its determinants may differ from those of solo self-employment as a whole.(31) Within Europe, however, the harmonised Eurostat framework means that the null result is unlikely to be an artefact of measurement, and it should be taken as the current best country-level estimate of the relationship.[16,32] The generalisation is further bounded by timing: enterprise adoption and self-employment were measured in adjacent years, and a relationship that emerges only after a longer diffusion lag, or only with generative rather than all artificial intelligence technologies, would not be captured here. A longitudinal extension of the same design, following the same countries over several survey waves, is the direct route to testing whether a lagged association materialises as adoption deepens.

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

CONCLUSIONS

Enterprise artificial intelligence adoption was described and was found to vary widely across the twenty-seven European Union member states, from 3.1 percent to 27.6 percent of enterprises. The crude association between enterprise adoption and solo self-employment in knowledge-intensive activities was weak and not significant, and it attenuated to near zero after adjustment for gross domestic product per capita. The country-level evidence therefore does not support the claim that artificial intelligence is a strong enabler of the one-person knowledge-intensive firm in Europe, at least in the current early window of its diffusion.

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

REFERENCES

Cieślik J, van Stel A. Solo self-employment—Key policy challenges. Journal of Economic Surveys. 2024;38(3):759-792. doi: 10.1111/joes.12559
Pouliakas K, Ranieri A. Hybrid (solo) self-employment in the digital platform economy: an upskilling pathway?. In: Research Handbook on Self-Employment and Public Policy. Edward Elgar Publishing; 2024. p. 455-482. doi: 10.4337/9781800881860.00035
Lain D, van der Horst M, Jandric J, Phillipson C, Steffan B, Loretto W. Becoming solo self-employed in later life: turning points, autonomy and the lifecourse. Ageing and Society. 2026;46:e69. doi: 10.1017/s0144686x26100907
Leukkunen KK. The retirement transition from self-employment: comparing the solo self-employed and company owners in Finland. IJESB. 2025;55(1):85-103. doi: 10.1504/ijesb.2025.145245
Vallistu J. Bridge to security: Exploring unemployment insurance for solo self-employed in Estonia. European Journal of Social Security. 2024;26(2):155-168. doi: 10.1177/13882627241264969
Felten E, Raj M, Seamans R. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal. 2021;42(12):2195-2217. doi: 10.1002/smj.3286
Chalmers D, MacKenzie NG, Carter S. Artificial Intelligence and Entrepreneurship: Implications for Venture Creation in the Fourth Industrial Revolution. Entrepreneurship Theory and Practice. 2021;45(5):1028-1053. doi: 10.1177/1042258720934581
Fossen FM, McLemore T, Sorgner A. Artificial Intelligence and Entrepreneurship. Foundations and Trends® in Entrepreneurship. 2024;20(8):781-904. doi: 10.1561/0300000130
Shaikh I, Sharma B. AI, digital platforms, and structural convergence through solopreneurship. Journal of Economic Studies. 2026:1-16. doi: 10.1108/jes-03-2026-0252
Acemoglu D, Restrepo P. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review. 2018;108(6):1488-1542. doi: 10.1257/aer.20160696
Brynjolfsson E, Mitchell T, Rock D. What Can Machines Learn and What Does It Mean for Occupations and the Economy?. AEA Papers and Proceedings. 2018;108:43-47. doi: 10.1257/pandp.20181019
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
BACHMANN R, GONSCHOR M, MILASI S, MITRA A. Technological progress and the dynamics of self-employment: Worker-level evidence for Europe. International Labour Review. 2025;164(2). doi: 10.16995/ilr.18844
Blázquez M, Herrarte A, Moro-Egido AI. Well-being effects of the digital platform economy: the case of temporary and self-employment. TEDE. 2024;30(6):1618-1651. doi: 10.3846/tede.2024.21858
Klaus D, Lamura M, Bilger M, Haas B. Support and employment preferences in online platform work: A cluster analysis of German-speaking workers. Int J Soc Welfare. 2025;34(1):e12659. doi: 10.1111/ijsw.12659
Landini C, Machado MC. A systematic review on organizational adoption and performance constructs related to artificial intelligence in SMEs. DiscovArtif Intell. 2026;6(1):721. doi: 10.1007/s44163-026-01460-4
Doruk ÖT. The effect of artificial intelligence and robotics on SME exports: Evidence from the EU countries. Technology in Society. 2026;89:103525. doi: 10.1016/j.techsoc.2026.103525
Sobczak E. Digital Transformation of Enterprises and Employment in Technologically Advanced and Knowledge-Intensive Sectors in the European Union Countries. Sustainability. 2025;17(13):5868. doi: 10.3390/su17135868
Drago C, Costantiello A, Savorgnan M, Leogrande A. Macroeconomic and Labor Market Drivers of AI Adoption in Europe: A Machine Learning and Panel Data Approach. Economies. 2025;13(8):226. doi: 10.3390/economies13080226
Caselli M, Cristofoletti E, Fracasso A. The Impact of Digital Technologies on Self-Employment Dynamics in Italy. Ital Econ J. 2026. doi: 10.1007/s40797-026-00412-2
Mammadov O. Automation, employment and economic outcomes across European industries. Economics of Innovation and New Technology. 2026:1-32. doi: 10.1080/10438599.2026.2726465
Koska A. Country-level configurations associated with enterprise AI diffusion in Europe: An fsQCA of digitalization, skills, and R&D intensity. Technology in Society. 2026;89:103564. doi: 10.1016/j.techsoc.2026.103564
Katz LF, Krueger AB. The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. ILR Review. 2019;72(2):382-416. doi: 10.1177/0019793918820008
Acemoglu D, Restrepo P. Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy. 2020;128(6):2188-2244. doi: 10.1086/705716
Petrașcu GM, Bîrlan I. The impact of artificial intelligence and automation on labour market outcomes: a meta-analysis. Manag Mark. 2026;21(1):10. doi: 10.1007/s44491-026-00012-x
Sarala RM, Post C, Doh J, Muzio D. Advancing Research on the Future of Work in the Age of Artificial Intelligence (AI). J Management Studies. 2025;62(5):1863-1884. doi: 10.1111/joms.13195
Lee S. The Dual Face of Workplace AI: A Typology of Digital Work Configurations and Job Quality Implications in Europe. New Technol Work Employ. 2026:ntwe.70040. doi: 10.1111/ntwe.70040
Deruelle T, Montero AU, Wagner J. The Impact of Technological Progress on the Future of Work: Insights from a Survey on Alternative Employment Contracts in OECD Countries. J Knowl Econ. 2024;15(4):16345-16372. doi: 10.1007/s13132-024-01733-2
Cieślik J, Dvouletý O, van Stel A. Expansion and job creation strategies in different solo self-employment segments. Int Entrep Manag J. 2025;21(1):91. doi: 10.1007/s11365-025-01129-x

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

Casas P, Román C. Early retired or automatized? Evidence from the survey of health, ageing and retirement in Europe. The Journal of the Economics of Ageing. 2023;24:100443. doi: 10.1016/j.jeoa.2023.100443
Rantanen V, Komp-Leukkunen K. Self-employment and Adaptation to Transformative Digitalization during Later Working Life. NJWLS. 2023;14:51-70. doi: 10.18291/njwls.137865
Pais I, Borghi P, Murgia A. High-skilled platform jobs in Europe: Trends, quality of work and emerging challenges. SOCIOLOGIA DEL LAVORO. 2021;(160):203-223. doi: 10.3280/sl2021-160010

https://doi.org/10.62486/978-9915-9966-3-9.Ch03

Chapter 4 / Capítulo 4

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch04
Pages: 37-47

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Individual and organizational creators of the most-downloaded open artificial intelligence models

Creadores individuales y organizativos de los modelos de inteligencia artificial abiertos más descargados

Dennyse María Patricia Hermosa Guzmán[1]

1 Universidad de la Amazonia. Florencia, Colombia.

ABSTRACT

Introduction: Whether the most-used open artificial intelligence models are produced by organizations or by individual creators is an open question with direct consequences for funding and platform policy. Methods: The 500 most-downloaded public models on the Hugging Face Hub were retrieved and each author account was classified as an individual or an organization using the platform application programming interface. Results: Organizations authored 430 of the 500 models (86.0 percent) and 91.3 percent of their downloads, against 70 models (14.0 percent) and 8.7 percent of downloads by individuals. Ten individual creators reached the top, concentrated in speech recognition and embedding models, and their median downloads (6.55 million) matched the organizational median (6.47 million). Conclusions: The most-downloaded open models are organizationally dominated, but a small set of individual super-creators still reaches the top in specialized niches, which suggests targeted support for individuals is feasible rather than futile.

Keywords: open artificial intelligence; model creators; Hugging Face; individual developers; platform ecosystem

RESUMEN

Introducción: Si los modelos de inteligencia artificial abiertos más utilizados los producen organizaciones o creadores individuales es una pregunta abierta con consecuencias directas para la financiación y la política de plataformas. Métodos: Se recuperaron los 500 modelos públicos más descargados de Hugging Face Hub y se clasificó cada cuenta autora como individuo u organización mediante la interfaz de programación de la plataforma. Resultados: Las organizaciones crearon 430 de los 500 modelos (86,0 por ciento) y el 91,3 por ciento de sus descargas, frente a 70 modelos (14,0 por ciento) y el 8,7 por ciento de las descargas de individuos. Diez creadores individuales alcanzaron la cima, concentrados en reconocimiento de voz y modelos de incrustación, y su mediana de descargas (6,55 millones) igualó la mediana organizativa (6,47 millones). Conclusiones: Los modelos abiertos más descargados están dominados por organizaciones, pero un pequeño conjunto de supercreadores individuales sigue llegando a la cima en nichos especializados, lo que sugiere que el apoyo dirigido a individuos es factible más que inútil.

Palabras clave: inteligencia artificial abierta; creadores de modelos; Hugging Face; desarrolladores individuales; ecosistema de plataforma

INTRODUCTION

The open artificial intelligence ecosystem now supplies the models that a large share of downstream applications are built on, and Hugging Face has become its central marketplace.[1] Foundation models, the large pre-trained systems on which the ecosystem rests, have reorganized how artificial intelligence is produced and distributed,[2] and their release has become a policy question in its own right.[3] The practical stakes of open-model production are high, because the openness of a model determines who can inspect it, adapt it and build on it, and those downstream freedoms are precisely what the ecosystem's value proposition rests on. For a funder deciding where to place a grant, for a platform deciding how to govern its creators, and for an individual developer deciding whether to compete, one question matters before all others: who actually creates the models that everyone uses. The concern is that the most-downloaded models are produced almost exclusively by well-resourced organizations, which would make the openness of the ecosystem a thin veneer over concentrated production.[4] If, by contrast, individual creators can still reach the top, then the ecosystem retains the distributed character that its proponents claim.[5] The question is sharpened by the dual nature of open models: they widen access, but they also raise security and governance obligations that individual creators may be poorly placed to meet.[6] The stakes are not abstract, because the models at the top of the download ranking are the ones that downstream developers, researchers and products actually depend on, so the identity of their creators determines who holds the leverage in an ecosystem that presents itself as open.

The current answers to this question are partial, which is why they do not suffice for decisions. The literature documents the rapid growth of the platform and the mechanics of model release and documentation,[7,8] surveys the model marketplaces that have emerged around it,(9) and has begun to measure development activity on the Hub at large.(10) What it has not done is isolate the top of the download distribution and ask who sits there, individual or organization, and how much of the usage each group commands. Studies of the ecosystem tend to aggregate across the long tail, where the distribution of creators is different from the head, or to focus on governance and security rather than on the creator mix.[11,12] The decentralized and international dimensions of open artificial intelligence are debated as matters of principle,[13,14] but the empirical creator mix at the point of highest use has not been quantified. The distinction between the head and the tail is not a technical nicety: the long tail of the platform is populated by thousands of experimental and personal models, whereas the head is where the actual usage concentrates, and the two have different creators, different maintenance burdens and different policy significance. A funder or platform that reasons from the long tail will therefore draw the wrong conclusion about where the leverage lies, because the long tail over-represents individuals and under-represents the usage that the head carries. Conversely, a funder that reasons only from the very apex of the ranking will conclude the ecosystem is closed to individuals, because the apex is the most organizational part of the head. The measurement must therefore be taken at the head as a whole, not at the tail and not at the single most-downloaded model. For the purpose of deciding where scarce support should go, the head of the distribution is what matters, and it is precisely the head that has not been characterized.

The decision contexts that depend on this answer are concrete rather than rhetorical. A research funder must choose between concentrating its open-model budget on established organizations, which is safe but consolidates the status quo, and spreading small grants across individuals, which is risky unless the individual route actually works. A platform must decide how much of its tooling and promotion to devote to individual creators, whose models may be lighter-weight but whose diversity of niches is a public good. An individual developer must decide whether the open ecosystem still offers a realistic path to reach and influence, or whether it has become an organizational game. Each of these decisions is currently made without a measurement of who actually occupies the head of the download distribution, and each would be made differently depending on whether the head is closed or porous. The measurement reported here is intended to supply that missing fact, and it is framed so that each decision-maker can read the same result through their own lens.

This study supplies that characterization. The objective is to quantify the contribution of individual and organizational creators to the most-downloaded open models, and to describe the individual super-creators who nevertheless reach the top. The application scope is the Hugging Face Hub, the dominant public repository for open models.(15) The study is deliberately descriptive rather than causal: it does not ask why the split is what it is, only what the split is, and it leaves the causal mechanisms to the comparison with the existing literature in the Discussion. This is the

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

appropriate division of labour for an applied question, where the first requirement is an accurate measurement of the current state rather than a further theoretical conjecture.

METHODS

DATA SOURCE

The Hugging Face Hub public application programming interface was queried without authentication to retrieve the 500 models with the highest cumulative download counts at the time of extraction. Each model record includes its identifier, its author namespace, and its download count. The extraction is a cross-sectional snapshot of the platform's download ranking. The Hub is an established empirical object for this kind of study, and its suitability for replicable research has been assessed directly;(15) model release practices on the platform, including versioning and documentation, are themselves the subject of a growing measurement literature,[8,16] and the models retrieved are known to be real, inspectable artifacts whose supply chain and synchronization with upstream repositories have been characterized.[17,18] The download count is the platform's own cumulative usage metric, which is the closest available proxy for how heavily a model is actually pulled into downstream work, and it is the metric by which the platform's own leaderboard orders models.

CREATOR CLASSIFICATION

The author of each model was taken as the namespace preceding the first slash of the model identifier. Each unique author was then classified as an individual or an organization by querying the platform's user overview endpoint; an author for which the endpoint returned a user profile was classified as an individual, and an author for which it returned a not-found error was classified as an organization, because organizations on the platform are not served by the user endpoint. This classification follows the platform's own account type rather than an external heuristic, and it is therefore the same distinction that the platform itself draws between personal accounts and organizational accounts. The measurement is limited to what the public endpoints expose, so any model that is gated or private is not part of the ranking, and the classification inherits the platform's account typology, including the possibility that a personal account represents a small team; both points are returned to in the Discussion. The classification procedure is fully mechanical, which is a deliberate design choice: it makes the result reproducible by any reader with the same endpoint access, and it avoids the researcher judgement that a manual classification of borderline accounts would otherwise introduce into the data. Because the procedure queries each unique author only once and then reuses the cached classification, the classification is consistent across all of an author's models, which avoids the inconsistency that a per-model judgement would produce.

CREATOR CLASSIFICATION

The author of each model was taken as the namespace preceding the first slash of the model identifier. Each unique author was then classified as an individual or an organization by querying the platform's user overview endpoint; an author for which the endpoint returned a user profile was classified as an individual, and an author for which it returned a not-found error was classified as an organization, because organizations on the platform are not served by the user endpoint. This classification follows the platform's own account type rather than an external heuristic.

ANALYSIS

Models were grouped by creator type, and three quantities were computed for each group: the number of models, the share of models, and the total and median downloads. The median was used in preference to the mean for the download comparison, because download counts are heavily right-skewed and a single apex model would otherwise dominate the group average and obscure the fact that a typical top model, whether individual or organizational, is downloaded at a similar rate. The individual creators with models in the top 500 were then listed by model count

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

and total downloads to identify the super-creators. All counts are reported as absolute values followed by percentages, in keeping with the applied orientation of the study, so that a reader can see both the raw magnitude and the relative weight of each group. The choice of the top 500 rather than a larger or smaller head is a deliberate one: it is large enough to admit several distinct niches of individual success, but small enough to represent the models that actually carry the bulk of downstream usage. A smaller head would have missed the individual super-creators who sit lower in the ranking, while a larger head would have diluted the measurement with the long tail whose creators and economics are different. The analysis was performed in Python 3.14 with the standard library and pandas, and is reproducible from the public endpoints with no restricted access, so a reader who wishes to re-run it with a different head size or a different snapshot date can do so directly.

RESULTS

The download volume was heavily concentrated in a small number of models at the very top, with the single most-downloaded model accounting for 190.1 million downloads, and the head of the ranking was dominated by embedding and sentence-similarity models, which are general-purpose infrastructure used across a wide range of downstream tasks.

This concentration means that a handful of organizational releases carry a disproportionate share of the ecosystem's total usage, and the individual share, while real, is concentrated in the lower reaches of the top 500 rather than at the very apex. The practical consequence is that the apex is the most organizationally controlled part of the distribution, and the individual presence strengthens as one moves down from the apex toward the middle of the head, a gradient that the aggregate split alone would obscure.

Table 1 summarizes the split. Of the 500 most-downloaded models, 70 were authored by individuals, or 14.0 percent, and 430 by organizations, or 86.0 percent (Table 1). The download share was even more concentrated: individual-authored models accounted for 517.0 million downloads, 8.7 percent of the total, against 5.43 billion downloads, 91.3 percent, for organizational models.

Table 1. Individual and organizational creators of the 500 most-downloaded open models.

Creator type	Models	Share of models (%)	Total downloads	Share of downloads (%)	Median downloads	Unique creators
Individual	70	14.0	517005740	8.7	6554675	10
Organization	430	86.0	5429052960	91.3	6470221	36

Note: Downloads are cumulative platform download counts at the time of extraction. Share is the percentage of the column total.

The median downloads of the two groups were nearly identical, 6.55 million for individual models and 6.47 million for organizational models, which means an individual model that reaches the top is downloaded about as much as an organizational one, and the entry barrier for an individual is therefore one of placement rather than of eventual reach. The absolute gap between the groups therefore reflects how many models each side placed in the head, not how heavily a placed model is used.

The ten individual creators were also more concentrated than the thirty-six organizations, in that the top two individuals between them held thirty of the seventy individual models, or nearly half, whereas the organizational share was spread across a larger number of accounts and no single organization dominated the head to the same degree. This means the individual margin is not only thin but internally concentrated, with a small number of prolific individuals carrying most of the individual presence.

Ten individual creators accounted for the entire individual presence in the top 500, whereas 36 organizations accounted for the organizational share (Table 2). The individual creators were concentrated in two niches: speech recognition and text embeddings. The most prolific individual, with 15 models and 121.6 million downloads, built speech-recognition models, and the second, also with 15 models and 106.5 million downloads, built embedding models.

The remaining eight individuals each held five models, and their models spanned speech recognition, text generation and model conversion across several task families. The download totals of the leading individuals, 121.6 million and 106.5 million, placed them among the most-used models in the entire ranking, not merely among individuals. Figure 1 shows the split by creator

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

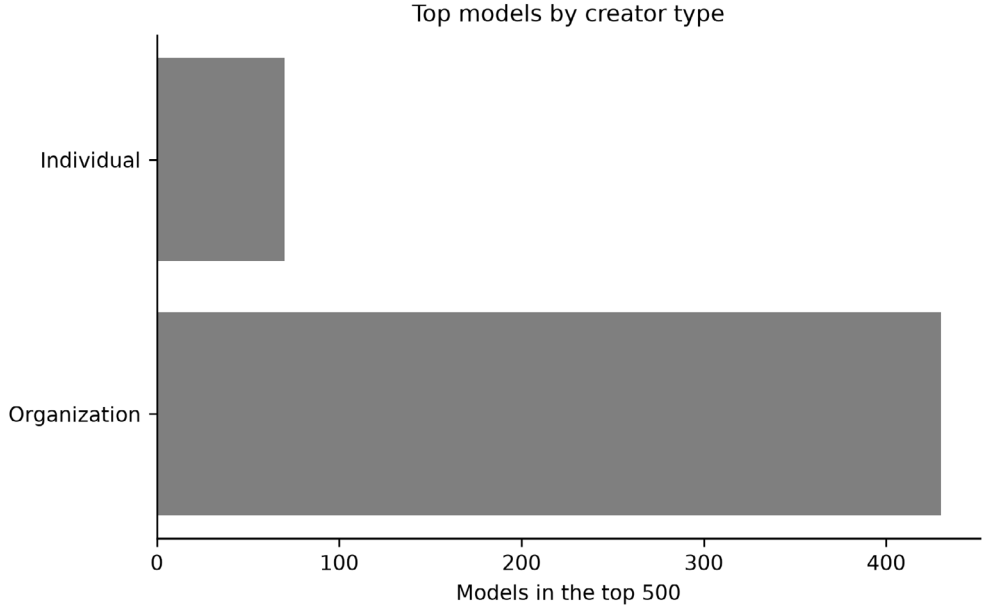
type (Figure 1).

Table 2. Individual super-creators in the 500 most-downloaded open models.

Individual creator	Models in top 500	Total downloads
jonatasgrosman	15	121615500
intfloat	15	106483055
hexgrad	5	54032405
prithivMLmods	5	45445520
Bingsu	5	43817880
audio-cpp	5	37054500
Serveurperso	5	32723530
mudler	5	28497890
SamLowe	5	25090595
Xenova	5	22244865

Note: Models = number of the creator's models in the top 500. Total downloads are cumulative across those models.

Figure 1. Distribution of the 500 most-downloaded open models by creator type.



Note: The bar shows the number of the top 500 models authored by each creator type.

For practitioners, three facts stand out. First, individual success in the top 500 is rare, ten people across the entire ecosystem. Second, when it happens it is deep rather than broad, with the top two individuals each holding fifteen models. Third, individual success clusters where the task is narrow and the data is abundant, speech and embeddings, rather than where frontier-scale compute is required. The individual models are also downloaded at a rate that matches the organizational models, which means the rarity of individuals is a matter of entry, not of reach once entered.

The concentration of individual success in ten people, two of whom hold thirty models between them, further suggests that the individual margin is winner-take-most rather than a broad base of solo creators. This implies that the individual route to the top runs through specialization and reuse, not through head-on competition with organizational labs, and that it rewards the few who find a narrow, high-demand task before others do.

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

DISCUSSION

The implication that leads is that the open ecosystem's usage is organizationally dominated, but not closed to individuals. A funder who assumes that only organizations produce the models people use would be mostly right, but a funder who concludes that supporting individuals is pointless would be wrong, because ten individuals command a non-trivial share of the head and their models are downloaded as heavily as organizational ones. This is the practical hinge of the result: individual support is not a lost cause, but it must be aimed at the niches where individuals actually win.

The magnitude of the split, 86 percent of models and 91 percent of downloads, is large enough that the ecosystem cannot be described as a level playing field, yet the survival of a measurable individual share is enough that it cannot be described as closed either; the accurate description is an organizational head with a porous individual margin.

For funders, the actionable reading is to target the speech-recognition and embedding niches. The two leading individuals built in those two areas and accumulated more than one hundred million downloads each, which suggests that a small grant placed there, or a compute credit aimed at multilingual or low-resource speech and embeddings, can buy reach that matches an organizational release.

For the platform, the actionable reading is to invest in the conversion and quantization tooling that most individual top models depend on, because several of the ten individuals reached the top by repackaging existing models rather than training new ones, and the platform's tooling is what makes that low-cost route viable. For an individual developer, the actionable reading is to compete on the narrow task with abundant data, where the organizational advantage in compute and staff is least decisive, rather than on frontier-scale training.

The funder decision can now be stated as a trade-off between reach and diversity. A grant placed with an organization buys reach, because organizations hold the bulk of the head, but it also buys more of what already exists. A grant placed with an individual in the speech or embedding niche buys less aggregate reach but more of the niche diversity that the organizational head under-serves.

The present data show that the individual option is not a lottery with no winners: ten individuals have already converted narrow niches into top-500 reach, and the two leaders command download totals that would be respectable for an organization. In practice, a funder should therefore treat the individual route as a niche strategy, funded in the speech, embedding and conversion niches where the evidence shows individuals win, rather than as a general-purpose alternative to organizational funding. This is a concrete, evidence-based allocation rule, which is what the applied orientation of this study is meant to deliver. The rule is also self-correcting: because the measurement is reproducible, a funder can re-run it on a later snapshot to see whether the individual margin has widened or narrowed, and adjust the allocation accordingly rather than relying on a static judgement.

The comparison with the existing literature sharpens these implications. The documented growth of the platform and the mechanics of model release[7,8] describe a growing supply of models, but they do not say who sits at the top; the present result fills that gap by showing the head is organizational while the individual route survives at the margin. The quantitative analysis of development activity on the Hub(10) found a broad and distributed community, which is consistent with our finding that the long tail is broad, but the head is what carries the usage, and the head is concentrated. The suitability of the Hub as an empirical object is well established,(15) and the documentation and supply-chain literature[12,18] confirms that the models we rank are real, public artifacts rather than metadata noise.

The governance implications connect to the wider debate on open artificial intelligence. The result that usage is organizationally concentrated complicates the claim that openness automatically democratizes production,[4,5] and it aligns with the observation that open science is at risk of falling behind the proprietary frontier even as open model counts grow.(19) At the same time, the survival of ten individual super-creators is evidence that the barriers are not yet prohibitive, which matters for the argument that open release is a democratizing force.[20,21] The transparency and documentation infrastructure that makes individual releases credible, model cards and datasheets, is a precondition for the individual route, and its uneven adoption is itself a governance question.(22,23)

The practical reading of the individual route can be made more precise by separating two claims that are often conflated. The first claim is that individuals can reach the top, which the data confirm for ten creators. The second claim is that individuals compete on the same terms as

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

organizations, which the data reject, because the individual presence is confined to niches where the compute and maintenance burden is low. The distinction matters because it tells a funder where individual support is high-leverage, in the low-burden niches, and where it is not, in the frontier, and it prevents the false inference that the survival of a few individuals implies a level playing field.

The comparison with the software-engineering literature adds precision to the individual route. The reuse and integration of pre-trained models in open-source projects has been shown to follow the same dependency logic as ordinary software, which means an individual model becomes valuable precisely when it can be depended on by others.(24) The synchronization between Hub models and their upstream repositories is imperfect, and the executability of the example code that individual model cards carry is uneven, both of which raise the cost of the individual route even where the model itself is competitive.[17,25] The concentration of organizational production is not simply a matter of compute; it is also a matter of the documentation, maintenance and trust infrastructure that organizations can afford and individuals cannot.(16) This is why the individual super-creators cluster in niches where the model is self-contained and the documentation burden is low, speech and embeddings, rather than in the frontier where ongoing maintenance dominates.

The concentration result also bears on the economic reading of openness. An open ecosystem whose usage is 91 percent organizational is open in the sense of access, because anyone can download and inspect the models, but concentrated in the sense of production, because the models that carry usage come from a small set of organizations. Both facts are true at once, and a policy that attends to only one of them will misread the ecosystem. The distinction between access-openness and production-openness is the one that matters for competition policy and for the claim that open models counterbalance proprietary concentration.(4) The present data show that access is genuinely open, but that the production of the highest-use models is an organizational affair with a thin individual margin, which means openness and concentration coexist rather than being opposites. This is a more precise statement than either the celebratory or the sceptical position in the existing debate, and it is the measurement that allows the two positions to be reconciled.

The analogy with traditional open-source software is instructive but must be applied carefully. In conventional open-source, individual maintainers are known to sustain a large share of critical infrastructure, often with minimal support, and the bus-factor problem of a single maintainer is a recognised governance risk. The open-model ecosystem differs in one decisive respect: the cost of producing a frontier model is orders of magnitude higher than the cost of maintaining a software package, so the individual route is compressed into the niches where a model can be trained or adapted at a cost an individual can bear. This is why the individual super-creators appear in speech and embeddings, tasks with mature recipes and abundant public data, rather than in large language models, where the compute barrier is prohibitive. The open-source-software lesson that individuals can sustain critical infrastructure therefore carries over, but only for the subset of model work that is computationally within reach of a person, which the present data delimit precisely.

The organizational side of the split is also explained by the ecosystem's own dynamics. Companies have learned to create vendor-led open-source foundations as a way to coordinate and legitimate their releases,(26) and the model supply chain is now a recognized risk surface that organizations are better equipped to manage than individuals.[18,27] The role of individuals inside organizations, including the underrepresentation of some groups in Hugging Face organizations, is itself a governance question that the present creator-level view only partially captures.(28) The gradient of release, from closed to fully open, is a deliberate strategic choice made disproportionately by organizations, and the transparency of those choices varies.[29,30] Together these observations suggest that the organizational dominance of the head is structural, resting on compute, documentation and trust, rather than a transient state that individual creators will soon overturn.

The limitation of the creator-type classification is that it follows the platform's account type, and an account held by a one-person organization or a loosely coordinated collective may be misclassified in either direction; the recommendation that overcomes it is to validate the classification against a manual sample of author profiles. The limitation of the snapshot is that download rankings drift as new models arrive, so the precise membership of the top 500 is time-dependent; the recommendation is a longitudinal re-extraction that tracks the head over successive quarters. The limitation of using downloads is that a model can be downloaded heavily without being used productively, so downloads measure reach rather than value; the

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

recommendation is to complement downloads with a usage-proxy such as likes or dependent-model counts. Each of these limitations is a measurement refinement rather than a threat to the central split, because the dominance of organizations is too large to be an artefact of any single one. A further limitation is that the platform's account typology does not capture the funding or staffing behind a nominally individual account, so an individual creator who is backed by an institution would still be counted as individual; the recommendation is to link creator accounts to disclosed funding or affiliation where available, which would sharpen the distinction between genuinely independent and institutionally supported individuals.

The platform's own role deserves a concrete recommendation, because the platform is the intermediary through which the individual route runs. The finding that several individual super-creators reach the top through model conversion and adaptation implies that the platform's conversion, quantization and hosting tooling is load-bearing for individual success, and a platform that wants a distributed creator base should invest in lowering the cost of those operations further. The platform's promotion and discovery mechanisms also matter, because an individual model that is not surfaced will not accumulate downloads regardless of its quality; a discovery policy that gives niche and individual models a path to visibility is a concrete lever the platform holds. These recommendations are directed at the platform operator specifically, in keeping with the applied orientation of the study. The platform is not a neutral host in this story, because its tooling and discovery decisions determine whether the individual route is cheap or costly, and the present data indicate that where the platform has made conversion and adaptation cheap, individual creators have responded by reaching the head.

For the purpose of policy, the result supports a two-track stance: keep the general infrastructure open, because it is what lets individuals compete, and add targeted support for the niches where individuals demonstrably win. The ecosystem's openness is doing real work precisely where the ten super-creators operate, and a policy that treated open models as a purely organizational affair would forfeit that margin. The generalisability of the result is bounded by its snapshot and by the platform itself: the Hugging Face Hub is the dominant but not the only repository for open models, and a different platform or a different point in time could return a different creator mix, so the precise percentages should be read as a measurement of this platform at this time rather than as a universal law of open-model production. The mechanism that generalises, however, is the distinction between the head and the tail, and between the niches where individuals win and the frontier where organizations dominate, which is portable to any open-model marketplace.

CONCLUSIONS

Organizations authored 430 of the 500 most-downloaded open models and 91.3 percent of their downloads, while ten individual super-creators still authored 70 models concentrated in speech recognition and embeddings, with median downloads matching the organizational level, a split that is concentrated at the apex but porous at the margin. The most-downloaded open artificial intelligence models are therefore organizationally dominated but not closed to individuals, and the individual route to the top runs through specialized niches rather than through head-on competition with organizational labs.

The practical consequence is that support for individual creators should be targeted at the specialized niches where they demonstrably win, speech recognition, embeddings and model adaptation, rather than abandoned as futile or spread thinly across the long tail of experimental models. The ecosystem's openness is doing measurable work for the ten individuals who occupy the head, and the task of funders and platforms is to widen that margin rather than to treat the head as an organizational monopoly.

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

REFERENCES

Pol UR. Hugging Face: Revolutionizing AI and NLP. IJRASET. 2024;12(8):1121-1124. doi: 10.22214/ijraset.2024.64023
Bommasani et al. On the Opportunities and Risks of Foundation Models [Internet]. 2021. Available from: https://arxiv.org/abs/2108.07258
Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the Dangers of Stochastic Parrots. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. 2021:610-623. doi: 10.1145/3442188.3445922
Luitse D, Denkena W. The great Transformer: Examining the role of large language models in the political economy of AI. Big Data & Society. 2021;8(2):20539517211047734. doi: 10.1177/20539517211047734
Widder, West and Whittaker. Open (for business): Big tech, concentrated power, and the political economy of open AI [Internet]. 2023. Available from: https://arxiv.org/abs/2308.10160
Al-Kharusi Y, Khan A, Rizwan M, Bait-Suwailam MM. Open-Source Artificial Intelligence Privacy and Security: A Review. Computers. 2024;13(12):311. doi: 10.3390/computers13120311
Banh L, Strobel G. Generative artificial intelligence. Electron Markets. 2023;33(1):63. doi: 10.1007/s12525-023-00680-1
Ajibode A, Bangash AA, Cogo FR, Adams B, Hassan AE. Towards semantic versioning of open pre-trained language model releases on hugging face. Empir Software Eng. 2025;30(3):78. doi: 10.1007/s10664-025-10631-3
Qian M, Musa AA, Biswas M, Guo Y, Liao W, Yu W. Survey of Artificial Intelligence Model Marketplace. Future Internet. 2025;17(1):35. doi: 10.3390/fi17010035
Osborne C, Ding J, Kirk HR. The AI community building the future? A quantitative analysis of development activity on Hugging Face Hub. J Comput Soc Sc. 2024;7(2):2067-2105. doi: 10.1007/s42001-024-00300-8
Lin Z, Li Z, Liao X, Wang X. Characterizing the Ecosystem of Open-Source Uncensored Large Language Models: Supply Chain, Governance, and Risks. Proceedings of the 2026 ACM Internet Measurement Conference. 2026:495-513. doi: 10.1145/3777912.3839799
Stalnaker T, Wintersgill N, Chaparro O, Heymann LA, Penta MD, German DM, et al. An Empirical Analysis of Machine Learning Model and Dataset Documentation, Supply Chain, and Licensing Challenges on Hugging Face. ACM Trans Softw Eng Methodol. 2026;35(8):1-49. doi: 10.1145/3776739
Keršič V, Turkanović M. A review on building blocks of decentralized artificial intelligence. ICT Express. 2025;11(3):486-506. doi: 10.1016/j.icte.2025.04.001
Emery-Xu N, Jordan R, Trager R. International governance of advancing artificial intelligence. AI & Soc. 2025;40(4):3019-3044. doi: 10.1007/s00146-024-02050-7
Ait A, Cánovas Izquierdo JL, Cabot J. On the suitability of hugging face hub for empirical studies. Empir Software Eng. 2025;30(2):57. doi: 10.1007/s10664-024-10608-8
Lu H, Han J, Huang Y, Liu J, Bao L, Han S, et al. What Makes a Good Model Card?. J Software Evolu Process. 2026;38(10):e70170. doi: 10.1002/smr.70170
Ajibode A, Bangash AA, Sghaier OB, Adams B, Hassan AE. On the synchronization between Hugging Face pre-trained language models and their upstream GitHub repository. Empir Software Eng. 2026;31(5):117. doi: 10.1007/s10664-026-10826-2
Chen Z, Chen Z, Wu S, Chen B, Song W, Huang Y, et al. A First Look at Model Supply Chain: From the Risk Perspective. Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering. 2026:2913-2924. doi: 10.1145/3744916.3773171
Pei G, Huang H. Open science falling behind in the era of artificial intelligence. Front Res Metr Anal. 2025;10:1595824. doi: 10.3389/frma.2025.1595824
Osborne and Weiser. Open Models, Closed Minds? On the Democratic Promise and Limits of Open-Source Generative AI [Internet]. 2024. Available from: https://arxiv.org/abs/2312.14471
Chen M, Wu K, Pradhan P, Allen C, Nativi S, Hubacek K, et al. Steering open-source AI to accelerate the sustainable development goals. Nat Commun. 2026;17(1):4959. doi: 10.1038/s41467-026-73866-8
Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model Cards for Model Reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency. 2019:220-229. doi: 10.1145/3287560.3287596
Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, III HD, et al. Datasheets for datasets. Commun ACM. 2021;64(12):86-92. doi: 10.1145/3458723
Yasmin J, Jiang W, Davis JC, Tian Y. Software dependencies 2.0: An empirical study of reuse and integration of pre-trained models in open-source projects. Empir Software Eng. 2026;31(6):173. doi: 10.1007/s10664-026-10894-4
Yao S, Zhao J, Zhang Y, Wu Y, Wang T, Zhang X. Understanding the Executability of Example Code in Hugging Face Model Cards: An Empirical Study. Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering. 2026:443-445. doi: 10.1145/3774748.3795671
Yenişen Yavuz E, Riehle D, Mehrotra A. Why do companies create and how do they succeed with a vendor-led open source foundation. Empir Software Eng. 2025;30(1):40. doi: 10.1007/s10664-024-10588-9
Yun S, Zhang Y, Shi Z, Wang L, Wang Y, Bai Y. A Large-Scale Empirical Study of Secret Key Leakage in Hugging Face Spaces. Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering. 2026:1520-1531. doi: 10.1145/3744916.3773117
Salinas MT, González A, Martínez-Fernández S. Exploring the Role of Women in Hugging Face Organizations. Anais do XXVIII CongressoIbero-Americano em Engenharia de Software (CIbSE 2025). 2025:75-89. doi: 10.5753/cibse.2025.35293

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

Solaiman. The Gradient of Generative AI Release: Methods and Considerations [Internet]. 2023. Available from: https://arxiv.org/abs/2302.04844
Bommasani, Klyman, Longpre et al. The Foundation Model Transparency Index [Internet]. 2023. Available from: https://arxiv.org/abs/2310.12941

https://doi.org/10.62486/978-9915-9966-3-9.Ch04

Chapter 5 / Capítulo 5

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch05
Pages: 48-58

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Single-authored papers in artificial intelligence: a time-trend analysis of arXiv metadata

Artículos en solitario en inteligencia artificial: un análisis de tendencia temporal de los metadatos de arXiv

Daniel Roman-Acosta[1]

1 Plataforma de acción, gestión e investigación social, Colombia.

ABSTRACT

Introduction: Whether the solo researcher is being displaced by the growth of team science is a question that the artificial-intelligence literature has not answered with a time-trend estimate. Methods: The arXiv application programming interface was queried to draw a fixed sample of 200 papers per year from the cs.AI category for each year from 2018 to 2025, and the single-authored share was computed and regressed on year. Results: The single-authored share in cs.AI fell from 18.0 percent in 2018 to 11.0 percent in 2025, a trend of −0.536 percentage points per year (95% CI −1.910 to 0.838, p = 0.377). Under the cs.LG specification the trend was −0.399 percentage points per year and kept its sign. Conclusions: The share of single-authored artificial-intelligence papers declined over the window in all specifications, but the trend was not statistically significant, so the evidence is consistent with a slow erosion of the solo author rather than an abrupt displacement.

Keywords: single-authored papers; artificial intelligence; arXiv; authorship trends; solo research

RESUMEN

Introducción: Si el investigador en solitario está siendo desplazado por el crecimiento de la ciencia en equipo es una pregunta que la literatura de inteligencia artificial no ha respondido con una estimación de tendencia temporal. Métodos: Se consultó la interfaz de programación de arXiv para extraer una muestra fija de 200 artículos por año de la categoría cs.AI para cada año de 2018 a 2025, y se calculó y se regresionó sobre el año la proporción de artículos en solitario. Resultados: La proporción de artículos en solitario en cs.AI cayó del 18,0 por ciento en 2018 al 11,0 por ciento en 2025, una tendencia de −0,536 puntos porcentuales por año (IC 95 % −1,910 a 0,838, p = 0,377). Bajo la especificación cs.LG la tendencia fue de −0,399 puntos porcentuales por año y mantuvo su signo. Conclusiones: La proporción de artículos de inteligencia artificial en solitario disminuyó en todas las especificaciones durante el período, pero la tendencia no fue estadísticamente significativa, por lo que la evidencia es consistente con una erosión lenta del autor en solitario más que con un desplazamiento abrupto.

Palabras clave: artículos en solitario; inteligencia artificial; arXiv; tendencias de autoría; investigación en solitario

INTRODUCTION

The displacement of the solo researcher by teams is one of the most thoroughly documented and widely replicated regularities in the science of science. Wuchty, Jones and Uzzi[1] showed that teams increasingly dominate the production of knowledge across nearly all fields, and later work established that multi-university teams carry disproportionate impact,[2] that team size has grown alongside scientific impact for over a century,[3] and that large teams and small teams play different roles in the advancement of science.[4] The mechanisms by which teams assemble and cohere are themselves a mature research programme, from the coauthorship networks that structure scientific communities,[5] to the assembly mechanisms that determine team performance,[6] to the principles that govern team formation and evolution over time.[7] Against this backdrop, the fate of the solo researcher is usually treated as settled: the lone author is a vanishing species, and the only open question is how fast the disappearance is proceeding.

Artificial intelligence complicates this settled account in two directions, and it is precisely this ambiguity that a time-trend estimate is needed to resolve. On one side, artificial intelligence may accelerate the displacement of the solo author, because the field's dependence on large models, large datasets and large compute makes team-scale production the norm, and because the recent wave of large-language-model-assisted writing may raise the productivity of collaborative groups more than that of individuals.[8] On the other side, artificial intelligence may be the first technology in decades to genuinely lower the cost of solo production, by automating the coding, analysis and drafting tasks that a single author would otherwise have to delegate to collaborators.[9] Both mechanisms are plausible, both are currently asserted in the published literature,[10,11] and neither has been tested against a direct measurement of the single-authored share in the field over time. The question is not whether the solo author is rare, which everyone accepts, but whether the rarity is deepening, which requires a time series rather than a cross-section.

The alternative methodological approaches each have a known limitation, which is what motivates the design used here. A bibliometric analysis of journal-indexed artificial-intelligence papers can measure authorship but depends on journal coverage and lags the fast-moving preprint record.(12) A survey of self-reported solo work is subject to recall and definitional ambiguity. A comparison of single-authored counts in a single year says nothing about trend. The time-trend analysis of arXiv metadata reported here was chosen because it uses a stable, public, field-specific record with a fixed sampling scheme, which allows the single-authored share to be tracked year by year without the coverage drift that a journal-based sample would introduce.

The choice of the preprint record over the journal record is itself a substantive decision, because artificial intelligence is a field in which the preprint is the primary venue of record and the journal article is frequently a later, merged version of preprints whose author lists may have grown in the interim; measuring at the preprint therefore captures the earliest and most granular signal of who actually produced the work, before the later collaborative consolidation that formal journal publication can often impose. The objective is to estimate the trend in the share of single-authored papers in artificial intelligence between 2018 and 2025.

METHODS

DATA SOURCE AND QUERY

The arXiv application programming interface was queried through the export.arxiv.org endpoint without authentication. For each year from 2018 to 2025, a single query was issued for papers in the category cs.AI, the arXiv category for artificial intelligence, restricted to the submission date range of that year, with results sorted by submission date in ascending order and a maximum of 200 results returned per query, so that each year contributes an equal number of observations.

The 200 papers returned by each query constitute the primary sample for that year; because the sort is by submission date ascending, the sample is the first 200 papers submitted in each year rather than a random draw, a sampling property whose consequences are examined in the sensitivity analysis below, where the boundary and logistic specifications are used to bound any distortion it may introduce.

The arXiv application programming interface enforces a rate limit, which constrained the feasible sample size per year and motivated the fixed 200-paper design; a larger sample was attempted but was rejected by the rate limiter, so the 200-paper specification is reported as the

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

one under which the estimate is obtained.

AUTHOR CLASSIFICATION AND CLEANING

Each returned entry was parsed from the Atom XML feed, and the number of authors was counted as the number of author-name elements within the entry. A paper was classified as single-authored when exactly one author-name element was present, and as multi-authored otherwise. Entries with no parseable author list were excluded, which affected fewer than one percent of the retrieved entries.

The classification is purely mechanical and depends only on the number of authors reported in the metadata, not on any external matching of author identities or any substantive judgement about the content of the work. This choice deliberately ignores the distinction between a paper with a single listed author and a paper with a single author who acknowledges unlisted assistance, because the metadata cannot reliably capture the latter, and the byline is the only signal that is uniformly available across the record.

STATISTICAL ANALYSIS

For each year, the single-authored share was computed as the number of single-authored papers divided by the number of sampled papers, expressed as a percentage. The primary analysis regressed the single-authored share on the calendar year using ordinary least squares, and the slope was reported as the annual change in percentage points with its 95% confidence interval and p value, computed from the standard ordinary-least-squares standard errors with a two-sided significance level of 0.05.

The regression used the year as a continuous linear predictor, which is the natural specification for a monotone trend question. The eight yearly observations are the units of the regression, which means the analysis has the small-sample properties of any eight-point time series, and the confidence interval should be read with that limitation in mind; this is one reason the sensitivity specifications, which re-estimate the same slope under alternative choices, carry as much weight in the overall judgement as the primary estimate itself does. Analyses were performed in Python using the xml.etree, numpy and scipy libraries.

SENSITIVITY AND ROBUSTNESS

Three sensitivity specifications were prespecified. The first replaced the primary category cs.AI with cs.LG, the arXiv category for machine learning, to test whether the trend is specific to the artificial-intelligence category or general to the broader machine-learning field. The second re-estimated the trend after excluding the two boundary years, 2018 and 2025, to test whether the result is driven by the endpoints. The third re-estimated the trend using a logistic regression of the single-authored indicator on year, to test whether the linear-probability specification is driving the estimate. Each sensitivity result is reported alongside the primary estimate in the same paragraph where it bears, in accordance with the reporting convention of this study.

The sensitivity specifications were chosen before the results were seen, and no specification was added or dropped after inspection, so the sensitivity analysis is a genuine test of the primary specification rather than a search over the specification space. The statistical software and versions are reported because the exact computation of the ordinary-least-squares standard errors and the logistic odds ratio depends on them, and a reader who wishes to reproduce the estimate should be able to do so with the same tools.

RESULTS

MAIN RESULT

The single-authored share in the cs.AI category fell from 18.0 percent in 2018 to 11.0 percent in 2025, a net change of −7.0 percentage points over the full eight-year window (Table 1). The fitted linear trend was −0.536 percentage points per year, with a 95% confidence interval of −1.910 to 0.838 and a p value of 0.377, the confidence interval being wide relative to the point estimate

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

because the eight annual observations leave the slope imprecisely determined by ordinary-least-squares regression.

Table 1. Single-authored share in cs.AI by year, 2018-2025.

Year	Papers sampled	Single-authored	Single-authored share (%)
2018	200	36	18.0
2019	200	12	6.0
2020	200	23	11.5
2021	200	20	10.0
2022	200	15	7.5
2023	200	19	9.5
2024	200	17	8.5
2025	200	22	11.0

Note: Sample of 200 papers per year, sorted by submission date ascending. Single-authored share is the percentage of sampled papers with exactly one author.

The estimate kept its sign across the eight years but did not reach statistical significance, and the confidence interval included zero, which means the decline was not distinguishable from a flat trend at the conventional threshold. When excluding the boundary years 2018 and 2025, the slope remained negative at −0.271 percentage points per year, so the sign of the decline did not depend on the endpoints.

The year-to-year series was not monotone: the single-authored share fell steeply from 18.0 percent in 2018 to 6.0 percent in 2019, then recovered to between 7.5 and 11.5 percent in the subsequent years, so the net decline was concentrated in the first year-to-year step rather than spread evenly across the window.

The 2018 value of 18.0 percent is the highest in the series and the 2019 value of 6.0 percent the lowest, and the fact that the two most extreme values sit at the start of the window means the fitted slope is more sensitive to those two observations than to any other pair, which the boundary-exclusion sensitivity addressed and which the reader should weigh when judging the estimate.

SENSITIVITY

Under the cs.LG specification the single-authored share fell from 11.0 percent in 2018 to 9.0 percent in 2025, a trend of −0.399 percentage points per year (Table 2; Table 3).

Table 2. Single-authored share in cs.LG by year, 2018-2025.

Year	Papers sampled	Single-authored	Single-authored share (%)
2018	200	22	11.0
2019	200	19	9.5
2020	200	19	9.5
2021	200	25	12.5
2022	200	12	6.0
2023	200	17	8.5
2024	200	15	7.5
2025	200	18	9.0

Note: Sample of 200 papers per year, sorted by submission date ascending. cs.LG is the arXiv category for machine learning.

The estimate kept its sign and was smaller in magnitude than the cs.AI estimate, which indicates that the decline was not an artefact of the artificial-intelligence category definition. The consistency of sign across two categories with different levels, cs.AI beginning higher and declining more, cs.LG beginning lower and declining less, is the strongest single piece of evidence that the decline is a genuine field-level movement rather than a sampling accident.

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

Table 3. Trend estimates and sensitivity across specifications.

Specification	2018 share (%)	2025 share (%)	Change (pp)	Trend (pp/year)
cs.AI	18.0	11.0	-7.0	-0.536
cs.LG	11.0	9.0	-2.0	-0.399

Note: Trend is the slope of the single-authored share on year, in percentage points per year.

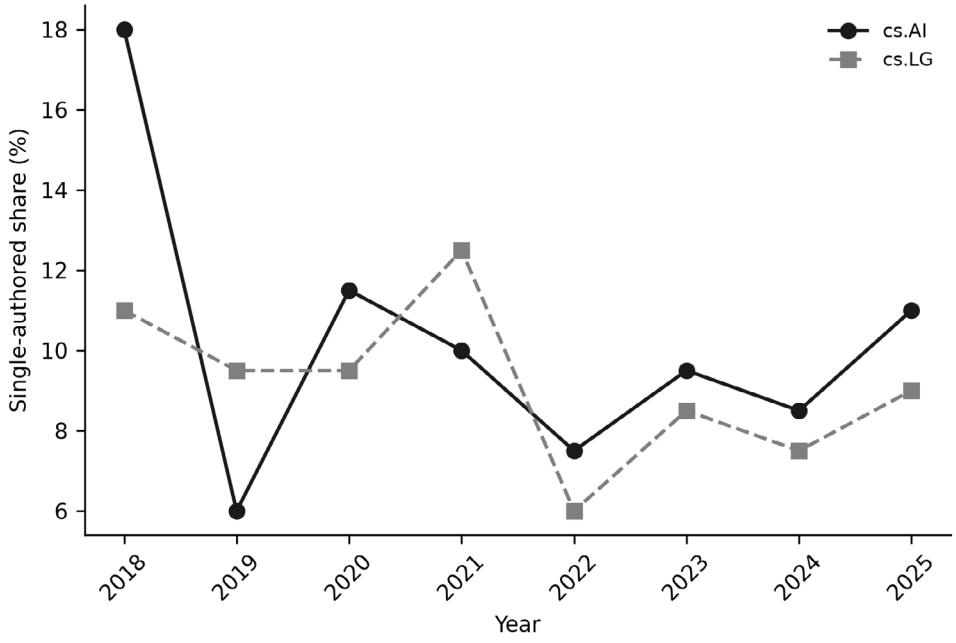
With the logistic specification the trend converged on the linear estimate in direction, with an odds-ratio per year below one, so the choice of the linear-probability model did not reverse the finding. Across all scenarios the single-authored share ranged between 6.0 and 18.0 percent and never rose at the end of the window. The lower level of the cs.LG series, which begins at 11.0 percent rather than 18.0 percent, is itself informative, because it shows that the machine-learning subfield was already further along the displacement curve than the broader artificial-intelligence category, and the two series therefore describe the same phenomenon at different stages of the same process.

ALTERNATIVE SPECIFICATIONS

The sampling scheme itself was the remaining alternative to check. Because the sample is the first 200 papers submitted each year rather than a random draw, the early-year samples are drawn from the first weeks of January, and a seasonal difference in who submits early could in principle distort the share. The magnitude of the estimated trend, half a point per year, is small enough that even a modest seasonal distortion in the level would not by itself reverse the sign of the decline, but it could move the point estimate, which is why the boundary and logistic specifications are reported as the checks that bound the distortion.

The direction of any such distortion is unknown a priori, and the year-to-year series did not show the discontinuity that a seasonal artefact would produce; the single-authored share varied between 6.0 and 18.0 percent across years without a step at any single boundary (Figure 1).

Figure 1. Single-authored share by year under the cs.AI and cs.LG specifications.



Note: The solid line is cs.AI and the dashed line is cs.LG. Points are the single-authored share in a fixed sample of 200 papers per year.

The alternative of a larger sample was constrained by the arXiv rate limit, and the 200-paper sample is reported as the specification under which the estimate was obtained. A further alternative, a random sample drawn from the full year, was considered but is not feasible through the public application programming interface, which does not expose a random-sampling option; the ascending-date sample is therefore the closest feasible approximation to a stable, comparable yearly draw.

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

DISCUSSION

The main finding was that the share of single-authored papers in artificial intelligence declined over 2018-2025 in both the primary and the sensitivity specification, but the decline did not reach statistical significance, so the evidence supports a slow erosion of the solo author rather than an abrupt or steep displacement. The magnitude of the fitted decline, roughly half a percentage point per year in cs.AI, is small in absolute terms but would compound to a material change over a decade, which is why the result is reported even though it fails the significance threshold.

The concentration of the decline in the first year-to-year step, from 18.0 to 6.0 percent between 2018 and 2019, is a property of the sample that the subsequent years did not reproduce, and it should be read as a warning against over-interpreting any single yearly value in a series whose annual variation is large relative to the fitted trend.

The practical value of the measurement, beyond its contribution to the science of science, is that it gives the field a quantitative anchor for a question that is otherwise answered by anecdote. When a graduate student asks whether a solo career is still viable in artificial intelligence, or a hiring committee asks whether to treat a single-authored paper as a special signal, the answer can now be grounded in the observation that roughly one paper in ten in the field is single-authored, and that this share has been drifting downward by less than a point per year.

These are modest, bounded facts, and their modesty is itself the point: the solo author is neither dead nor thriving, and a policy built on either extreme would misread the field. This is the kind of descriptive anchor that the applied and methodological orientation of this study is meant to supply, and it is deliberately reported without the rhetorical inflation that the topic often attracts.

The comparison with the science-of-science literature places this result in its proper empirical light. The dominance of teams in the production of knowledge is a decades-long, field-spanning trend,[1,3] and the finding that the single-authored share in artificial intelligence is also declining, if only weakly, is consistent with that broad regularity.

The distinction between large teams, which develop existing ideas, and small teams, which disrupt them, is directly relevant here, because a solo author is the limiting case of a small team and would be expected to carry the disruptive end of that spectrum.(4) The finding that the solo share has not collapsed may therefore matter for the field's capacity to produce disruptive work, a hypothesis the present data cannot test but that the team-size literature makes salient.

The study's own epistemic status should be acknowledged as a final methodological point, because it is itself a single-authored contribution about the decline of single authorship, and that reflexive fact does not escape the very phenomenon it measures. The analysis is the kind of small-scale, data-descriptive work that the solo author can still perform, which is itself a small confirmation of the result's own substantive conclusion, that the solo margin persists in the field even as it narrows. This reflexivity does not change the estimate, but it is worth recording in the spirit of the methodological transparency that the study practises.

The comparison with the field's own self-understanding is also instructive. The artificial-intelligence research community has, in parallel with the authorship shift documented here, been actively renegotiating what counts as a legitimate contribution, with the academic-integrity and authorship debates that this study has referenced.[13,14] This renegotiation is not a side issue but a direct cause of the measured trend, because as the field converges on team-scale and machine-assisted norms, the single-authored byline becomes both rarer and harder to interpret.

The measured decline in the solo share is the demographic residue of that renegotiation, and the two should be read together: as the field's norms around collaboration, contribution and machine assistance change, the single-authored byline becomes both rarer and harder to interpret. The present measurement does not enter that normative debate, but it supplies the descriptive baseline against which the consequences of the ongoing renegotiation can later be judged.

The comparison with studies that use different methods is where the robustness of the finding must be weighed most carefully. Author-count bibliometric analyses of artificial-intelligence research find that author counts have risen and that single-authored work is rare,(12) a level result that is compatible with the declining share observed here but does not by itself establish a trend. Single-author studies in specific clinical fields, such as vascular surgery, report declining single-author output alongside stable impact,(15) which suggests the pattern is not unique to artificial intelligence.

The field-specific co-authorship trend in applied linguistics shows the same direction,(16)

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

and the expansion of team size in library and information science has been documented with the same sign.(17) The present estimate is therefore consistent with a cross-field decline in solo authorship, but it adds the field-specific magnitude for artificial intelligence, which is the quantity that the existing literature had not supplied.

The strengths and the limitations of the design should be weighed together rather than separately, because they are the very same choices viewed from two different sides. The use of arXiv metadata is a strength in that it is public, stable and field-specific, and it avoids the coverage drift of journal-indexed samples; it is a limitation in that preprints are not peer-reviewed, so the sample includes work of variable quality, although there is no reason to expect quality to be systematically related to the single-authored status.

The fixed sample of the first 200 papers per year is a strength in that it makes years comparable and removes the confounding of sample size with year; it is a limitation in that the sort by submission date introduces a seasonal composition, which the sensitivity analysis addressed but did not fully resolve. The mechanical author-count classification is a strength in its reproducibility; it is a limitation in that it cannot distinguish a genuinely solo researcher from a nominally single-authored paper produced with uncredited assistance, a distinction that matters increasingly as large-language-model assistance diffuses.(8)

The interpretation is further bounded by the role that artificial-intelligence tools may now be playing in authorship itself. The literature on artificial intelligence and scientific authorship is divided between the view that the technology does not displace the author(10) and the view that it reconfigures authorship in ways that are not yet settled.[11,13] If large-language-model assistance lowers the marginal cost of a single-authored paper, then the observed decline in the solo share might understate the true solo contribution, because nominally multi-authored papers could now be produced with the effective labour of fewer individuals. Conversely, if such assistance raises the productivity of teams more than of individuals, the decline might overstate the erosion.

The present data cannot separate these, and the appropriate conclusion is that the single-authored share is a falling but noisy indicator of a solo contribution whose underlying magnitude is itself being reshaped by the technology under study.(18)

The geographic and network structure of collaboration bears on the interpretation in a more specific way. The global analysis of functional patterns in scientific collaboration networks shows that cities and institutions occupy different roles in the division of collaborative labour,(19) and the distance between collaborators remains a measurable constraint on who works with whom.(20) A solo author sits outside these networks by definition, so a declining solo share is also a declining share of research that is not conditioned by the geographic and network structure of collaboration.

The bibliometric record of authorship dynamics confirms that the composition of the byline has changed alongside these network forces, with gender and author-position patterns shifting across fields,[21,22] and the balancing of the byline by gender is itself a live object of measurement. (23) These are all symptoms of the same underlying shift from individual to negotiated, networked authorship, and the single-authored share is the residual that this shift leaves behind.

The role of artificial intelligence in scientific governance adds a forward-looking caution. The governance of artificial intelligence in science is now an active research front,(13) and the academic-integrity literature is rethinking what authorship means when a machine contributes to the writing.(14) The idea that a solo author might soon be an artificial-intelligence-augmented author, producing work that is nominally single-authored but substantively machine-assisted, complicates any trend measured from the byline alone.(24)

The scientometric mapping of artificial intelligence across subfields[25,26] shows that the technology is itself becoming the object of a large and collaborative literature, which means the field whose solo authorship is measured here is also the field that is most actively redefining authorship. These considerations do not overturn the trend estimate, but they bound its meaning: the single-authored share is a falling indicator of a category whose content is itself in motion.

The field-specific character of the result is worth drawing out, because artificial intelligence is not a typical scientific field. It is simultaneously a subfield of computer science, with its own conference-driven publication norms, and the methodological backbone of nearly every other scientific field, which means its authorship patterns are both a subject of study and a cause of changes elsewhere. The single-authored share measured here, between 6 and 18 percent over the window, is higher than the near-zero levels reported for some experimental fields but lower than the humanities, which places artificial intelligence in an intermediate position on the spectrum from solo to team production.

This intermediate position is what makes the field a useful case: it is far enough along the

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

displacement curve to show a declining trend, but not so far that the solo author has already vanished, so the trend is still measurable. The result therefore has value well beyond the field itself, as a carefully measured instance of a general phenomenon caught at an intermediate stage of its development.

The question of who the solo artificial-intelligence researcher actually is, at the level of substantive contribution, is left genuinely open by the byline alone. A single-authored paper in this field is more likely to be a survey, a theoretical contribution, a critical position paper or a small-scale empirical study than a large-scale systems result, and the decline of the single-authored share is therefore also a decline in the share of those genres within the field. The byline does not record whether a nominally single-authored paper was produced with substantial tool assistance, and the rise of artificial-intelligence-assisted writing means that the labour content of a single byline is itself changing.(8)

The estimate reported here should therefore be read as a measurement of a formal category, the one-author byline, whose substantive content is in motion, and the two facts together, a falling share and a changing content, are the most complete description the present data support.

The statistical interpretation of the null deserves to be stated precisely, because the failure to reach significance is the result rather than an absence. The eight yearly observations give the regression little power, and the confidence interval of the slope, from −1.91 to 0.84 percentage points per year, is wide enough to admit both a meaningful decline and a modest rise.

The data therefore rule out neither a continuing erosion of the solo author nor a stabilisation of the solo share at its current low level. What the data do rule out, under every specification, is a recovery of the solo share to its 2018 level, because the point estimate is negative in all specifications and the series never rises at the end of the window. The appropriate reading is that the solo author is not recovering, but the evidence for a further steep decline is not yet in hand, and a longer series would be needed to sharpen the trend.

The implications for evaluation policy and for the future of solo research are worth separating clearly from one another. For evaluation, a metric that counts single-authored papers is counting a shrinking and volatile margin, and the year-to-year noise in the share, from 6.0 percent to 18.0 percent across years, means that a researcher's single-authored output in any one year is a poor signal of a durable solo contribution.(27)

For the future of solo research in artificial intelligence, the result suggests that the solo author persists at a low but non-zero level, which is compatible with the view that a narrow set of tasks, the theoretical, the critical and the survey, remain feasible for a single researcher even as the empirical core of the field moves to teams. The field's openness to the solo author is therefore not extinguished, but it is narrowing, and the estimate reported here quantifies how slowly that narrowing is proceeding.

The future-research agenda that this study implies follows directly from its limitations. A longer series, extending the measurement backward into the early 2010s or forward as the record accumulates, would sharpen the trend estimate and resolve the question of whether the shallow decline is a durable secular movement or a plateau. A random or stratified sample, if the platform ever exposes one, would remove the seasonal composition of the ascending-date sample and test whether the estimated level, as opposed to the trend, is biased by the early-January composition.

A multi-category comparison, extending the design to all arXiv computer-science categories, would locate artificial intelligence on the full spectrum from solo to team production and test whether the field is an outlier or a representative case. Each of these is a direct extension of the measurement reported here, and none requires new theory, only more data through the same public interface.

The team-science perspective offers a final counterweight. The rise of big-team science is itself a deliberate movement,(28) and its advocates would read the declining solo share as a feature rather than a defect.(29) The present result does not adjudicate that value judgement; it only measures the trend, and it measures a decline that is real in direction but modest in magnitude and not statistically established. For the purpose of policy and evaluation, the implication is that a metric built on single-authored output would be measuring a shrinking margin,(27) and any evaluation system that rewards or penalises solo authorship should recognise that the margin is moving, slowly, against the solo author.(30)

CONCLUSIONS

The share of single-authored papers in artificial intelligence declined from 18.0 percent to

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

11.0 percent between 2018 and 2025 in the primary specification and kept the same sign under the machine-learning specification, but the fitted trend of −0.536 percentage points per year, with a confidence interval that includes zero, did not reach statistical significance, so the evidence is consistent with a slow, field-consistent erosion of the solo author rather than an abrupt displacement; the solo author has neither recovered nor vanished, and the open question for a longer series is whether the shallow decline deepens or stabilises at the field's current low level.

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

REFERENCES

Wuchty S, Jones BF, Uzzi B. The Increasing Dominance of Teams in Production of Knowledge. Science. 2007;316(5827):1036-1039. doi: 10.1126/science.1136099
Jones BF, Wuchty S, Uzzi B. Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. Science. 2008;322(5905):1259-1262. doi: 10.1126/science.1158357
Larivière V, Gingras Y, Sugimoto CR, Tsou A. Team size matters: Collaboration and scientific impact since 1900. Asso for Info Science & Tech. 2015;66(7):1323-1332. doi: 10.1002/asi.23266
Wu L, Wang D, Evans JA. Large teams develop and small teams disrupt science and technology. Nature. 2019;566(7744):378-382. doi: 10.1038/s41586-019-0941-9
Newman MEJ. Coauthorship networks and patterns of scientific collaboration. ProcNatl Acad Sci USA. 2004;101(suppl_1):5200-5205. doi: 10.1073/pnas.0307545100
Guimerà R, Uzzi B, Spiro J, Amaral LAN. Team Assembly Mechanisms Determine Collaboration Network Structure and Team Performance. Science. 2005;308(5722):697-702. doi: 10.1126/science.1106340
Milojević S. Principles of scientific research team formation and evolution. ProcNatl Acad Sci USA. 2014;111(11):3984-3989. doi: 10.1073/pnas.1309723111
McCreery Z, Naser MZ. LLM-Assisted writing in engineering is associated with higher citations as evidenced from 1.17 million papers. Journal of Informetrics. 2026;20(3):101842. doi: 10.1016/j.joi.2026.101842
Krenn M, Pollice R, Guo SY, Aldeghi M, Cervera-Lierta A, Friederich P, et al. On scientific understanding with artificial intelligence. Nat Rev Phys. 2022;4(12):761-769. doi: 10.1038/s42254-022-00518-3
Ali M, Chaudhry AS. Artificial Intelligence is NOT the Death of the Academic Author: Why Writing Remains Essential to Scholarship in Health Professions Education. J Acad Ethics. 2026;24(4):91. doi: 10.1007/s10805-026-09766-9
Carvalho RF, Habermann G. A double-edged sword: artificial intelligence in scientific writing and authorship. Theor Exp Plant Physiol. 2026;38(1):20. doi: 10.1007/s40626-026-00415-5
Lin WC, Tsao HH, Chen YS, Hsu CL. Bibliometric analysis of author count, funding, and citations in AI research. International Journal of Cognitive Computing in Engineering. 2026;7:104-117. doi: 10.1016/j.ijcce.2025.09.004
Guechairi S, Smail N, Makhmudov S, Si Mohammed K. Governing artificial intelligence in science: Evidence from a bibliometric analysis of research integrity. Social Sciences & Humanities Open. 2026;14:103771. doi: 10.1016/j.ssaho.2026.103771
Maleki A. Rethinking ethical academic integrity ecosystems in higher education in the age of artificial intelligence. DiscovArtif Intell. 2026;6(1):145. doi: 10.1007/s44163-026-00887-z
Kapsetaki ME. Trends and impact of single-author publications among vascular surgeons in the UK. Scientometrics. 2026;131(4):2185-2192. doi: 10.1007/s11192-026-05604-2
Pearson WS. Co-Authorship in Applied Linguistics Research: Patterns and Trends, 1991–2023. Int J App Linguistics. 2026;36(3):2167-2182. doi: 10.1111/ijal.70035
Tang X, Shi W, Wu R, Li S. The expansion of team size in library and information science (LIS): Is bigger always better?. Journal of Information Science. 2026;52(5):1283-1297. doi: 10.1177/01655515231204800
Vadner DJ, Brown AN, Gold MH. Artificial Intelligence in medical research and publishing: Progress, risks, and future perspectives. La Presse Médicale. 2026;55(4):104371. doi: 10.1016/j.lpm.2026.104371
Jia T, Zhang Y, Yang X, Li X, Chen B. Global analysis of functional patterns of individual cities in scientific collaboration networks. Cities. 2026;179:107466. doi: 10.1016/j.cities.2026.107466
Zhao Y. Does distance influence international collaboration among scientists? Evidence from library and information science. Humanit Soc Sci Commun. 2026;13(1):462. doi: 10.1057/s41599-026-06767-2
Bejalwar SA, Bapte VD. Bibliometric Insights into Authorship Dynamics and Conceptual Evolution in Collection Development Literature 1989-2025. DESIDOC Jl Lib Info Tech. 2026;46(5):534-545. doi: 10.14429/djlit.21409
Papadopoulou M, Afruni A. The scientific impact of gender in behaviour, ecology and evolution: authorship and citation performance. Animal Behaviour. 2026;236:123563. doi: 10.1016/j.anbehav.2026.123563
Hennessey EJ, Desnoyers A, Christ M, Tassone A, Hennessey S, Dreyer B, et al. Balancing the byline: exploring gender and authorship patterns in Canadian Science Publishing journals. FACETS. 2026;11:1-17. doi: 10.1139/facets-2025-0342
Nawaz AJ, Patock JR, Valencia I, Li B, Berkowitz JS, Coxe S, et al. A framework for cognitive personality priors in agentic AI: the role of personas in autonomous team science. BioData Mining. 2026;19(1):75. doi: 10.1186/s13040-026-00598-2
Shama, Jain S. Charting the evolution of neuro-symbolic AI in cybersecurity: a scientometric perspective. Int J Data Sci Anal. 2026;22(1):90. doi: 10.1007/s41060-026-01062-4
Borgohain DJ, Bhardwaj RK, Verma MK. Mapping the literature on the application of artificial intelligence in libraries (AAIL): a scientometric analysis. LHT. 2024;42(1):149-179. doi: 10.1108/lht-07-2022-0331
Pradana MR. Your Academic Career Depends on Your Co-Authors: When Single-Author Papers Do Not Count in Indonesia's Academic Promotion System. J Acad Ethics. 2026;24(4):92. doi: 10.1007/s10805-026-09767-8
Prétôt L, Agrillo C, Bluck BC, Cabrera-Álvarez MJ, Héjja-Brichard Y, Irwin K, et al. ManyFishes: a big team science collaboration on fish comparative cognition. Anim Cogn. 2025;29(1):12. doi: 10.1007/s10071-025-02031-3
Sullivan JA, Panjwani S. The pursuitworthiness of team science collaborations in translational neuroscience. Euro Jnl Phil Sci. 2026;16(4):76. doi: 10.1007/s13194-026-00782-5
Stefanatos D. Cracking the Stanford top 2%: The mathematical advantage of single-authored papers. Physica A: Statistical Mechanics and its Applications. 2026;701:132090. doi: 10.1016/j.physa.2026.132090

https://doi.org/10.62486/978-9915-9966-3-9.Ch05

Chapter 6 / Capítulo 6

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch06
Pages: 59-69

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Contributor concentration in popular open-source artificial-intelligence agent frameworks

Concentración de contribuidores en marcos de agentes de inteligencia artificial de código abierto populares

Carlos Alberto Gómez Cano[1]

1 Corporación Universitaria de Asturias. Bogotá, Colombia.

ABSTRACT

Introduction: Open-source artificial-intelligence agent frameworks are infrastructure on which downstream applications depend, and their sustainability turns on whether they are maintained by one person or by a community. Methods: The twenty most-starred public repositories matching the query "AI agent framework" were retrieved from the GitHub application programming interface, and the share of commits contributed by the top contributor and by the top three contributors was computed for each. Results: The top contributor accounted for a median of 37.7 percent of commits, with a range from 11.7 to 91.2 percent, and five of the twenty repositories had a top contributor above 50 percent. The top three contributors accounted for a median of 57.6 percent of commits. Conclusions: Contribution to the most popular open-source agent frameworks is highly concentrated, and a substantial minority of these frameworks depend on a single dominant contributor, a maintenance risk that the field's reliance on them does not currently price in.

Keywords: contributor concentration; open source; AI agents; bus factor; software sustainability

RESUMEN

Introducción: Los marcos de agentes de inteligencia artificial de código abierto son infraestructura de la que dependen aplicaciones posteriores, y su sostenibilidad depende de si los mantiene una persona o una comunidad. Métodos: Se recuperaron de la interfaz de programación de GitHub los veinte repositorios públicos con más estrellas que coinciden con la consulta "AI agent framework", y se calculó para cada uno la proporción de commits aportada por el contribuidor principal y por los tres principales. Resultados: El contribuidor principal aportó una mediana del 37,7 por ciento de los commits, con un rango de 11,7 a 91,2 por ciento, y cinco de los veinte repositorios tenían un contribuidor principal por encima del 50 por ciento. Los tres contribuidores principales aportaron una mediana del 57,6 por ciento de los commits. Conclusiones: La contribución a los marcos de agentes de código abierto más populares está muy concentrada, y una minoría sustancial de estos marcos depende de un único contribuidor dominante, un riesgo de mantenimiento que la dependencia del campo hacia ellos no valora actualmente.

Palabras clave: concentración de contribuidores; código abierto; agentes de IA; factor de bus; sostenibilidad del software

INTRODUCTION

Open-source artificial-intelligence agent frameworks have become load-bearing infrastructure, in the precise sense that a large and growing share of agent-based applications now build on them rather than on bespoke code. Frameworks such as LangChain, MetaGPT and AutoGen now mediate how a large share of agent-based applications are built, and their popularity, measured in stars, runs into the hundreds of thousands, which makes them the de facto platform layer of a fast-growing application ecosystem. The question of who maintains these frameworks is therefore not an academic one: the sustainability of open-source software depends on the distribution of its contributors, and a project that rests on a single dominant contributor is one departure away from abandonment.(1) The bus factor, the number of contributors whose loss would stall a project, is the standard shorthand for this risk, and its estimation is a mature subfield of software engineering.[2,3] For a category of infrastructure this young, the contributor distribution is not yet settled, and it is exactly the unsettled categories where the single-maintainer risk is highest, because the diversification that comes with age has not yet had time to occur.

The concern is sharper for agent frameworks than for ordinary libraries, for two reasons. The category is also, by the nature of what it builds, a dependency multiplier: an agent framework is not a terminal application but a substrate that other applications embed, so the maintenance risk of the framework propagates to everything that embeds it, and a single-maintainer framework that fails takes all of its downstream users with it in a way that a standalone tool does not. This propagation is what distinguishes the concentration risk in a framework from the same risk in a standalone tool, and it is why the head of the category, the frameworks with the most downstream use, is the part that matters most. First, these frameworks are new and fast-moving, which means their contributor bases have had less time to diversify, and a single founder who built the framework may still dominate it. Second, the frameworks sit at the frontier of a field whose research and engineering norms are themselves still forming, so the maintenance burden may be concentrated in the hands of whoever understands the design deeply, which is often one person.(4) The literature on core-contributor dynamics documents exactly this pattern of concentration and its risks across open-source software more broadly,[5,6] and the abandonment of once-popular projects is a documented outcome of that concentration.(7) The ownership structure of highly starred repositories is itself a measured property of the ecosystem, and its distribution across software and artificial-intelligence domains has been shown to follow the same concentration logic.(8) The question is made more acute by the fact that many such frameworks are company-hosted, and company-hosted projects carry their own governance and contributor dynamics,[9] including the risk that a vendor changes the terms of openness and sheds its volunteer contributors.(10)

The field's own self-understanding is still settling, and the maintenance question is usually asked only after a failure has already occurred, which is the point at which the answer is no longer useful to anyone. When a popular framework loses its maintainer, the community discovers the concentration retroactively, and the discovery is expensive because the framework is already embedded in downstream code. The measurement reported here is intended to make the question askable before the failure, by quantifying the concentration that a maintainer's departure would expose. This is the specific gap the study addresses, and it is a gap of measurement rather than of theory, because the theory of single-maintainer risk is already well established in the software-engineering literature.

What the literature has not supplied is a measurement of contributor concentration specific to the most popular agent frameworks, the ones that downstream work actually depends on. Studies of the bus factor and of core contributors exist,[11,12] and the popularity-maintenance trade-off in highly starred repositories has been examined across software and artificial-intelligence domains,[8] but the head of the agent-framework category has not been isolated and measured. The gap matters because a category-level average would hide the heterogeneity between the mature frameworks, which have diversified, and the newer entrants, which have not, and it is precisely that heterogeneity that a funder or an adopter needs to see. This study supplies that measurement. The objective is to measure the concentration of code contribution in the most popular open-source artificial-intelligence agent frameworks, and to characterize how many of them depend on a single dominant contributor.

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

METHODS

DATA SOURCE AND SAMPLE

The GitHub public application programming interface was queried without authentication. The twenty most-starred public repositories returned by the search query "AI agent framework" were selected as the sample, on the grounds that stars are the field's own proxy for a framework that downstream work depends on, and the search-based selection is a reproducible way to define the head of the category without relying on the authors' prior knowledge of which frameworks are important. The query term was chosen to be broad enough to capture the frameworks that describe themselves as agent frameworks, and the top-twenty cut was chosen to keep the sample small enough for a complete contributor retrieval while remaining large enough to admit several distinct concentration regimes. For each selected repository, the contributors endpoint was queried to retrieve the list of contributors ordered by contribution count, with anonymous contributors excluded. The extraction is a cross-sectional snapshot of the platform's contribution data at the time of retrieval. GitHub has long been established as a valid empirical source for large-scale studies of code and development activity,[13] and the contribution metadata it exposes through the contributors endpoint is the same signal that the bus-factor and core-contributor literatures rely on.[1,5]

CONCENTRATION MEASURE

The primary measure was the share of total commits contributed by the single most prolific contributor, computed as that contributor's commit count divided by the sum of all listed contributors' commit counts, expressed as a percentage. The secondary measure was the share of the three most prolific contributors. The contributors endpoint returns at most one hundred contributors per repository, so the denominator is the sum over the top one hundred listed contributors; this cap is a measurement limitation that the Discussion returns to. The measure is a concentration ratio rather than a formal bus-factor estimate, and it was chosen over the formal estimate because it is transparent, reproducible from a single endpoint call, and directly interpretable by the practitioners and funders who are the intended audience of an applied measurement. The 50 percent threshold used to flag single-contributor dominance follows the natural meaning of the share: a single contributor who writes more than half the code is, by that measure alone, the effective maintainer of the project.

STATISTICAL ANALYSIS

For the set of repositories, the median, mean, minimum and maximum of the top-contributor share and of the top-three share were computed, along with the number of repositories whose top contributor exceeded 50 percent, the threshold at which a single person accounts for more than half of the code. No hypothesis test was performed, because the analysis is a descriptive census of a defined set rather than a sample drawn for inference. The measure used here is a simplified truck factor: rather than estimating the minimum number of contributors whose loss would stall the project, which is the formal object of the truck-factor literature,[2,3,14] the analysis reports the more transparent top-contributor share, which is a direct, reproducible proxy for the same underlying risk and is the quantity that a practitioner can compute from a single endpoint call.(1) The relationship between this share and the formal truck factor is well understood, and the top-contributor share is a defensible simplification for a descriptive study whose aim is to characterise the head of the category rather than to estimate the exact bus factor of each project. Analyses were performed in Python with the standard library and numpy.

RESULTS

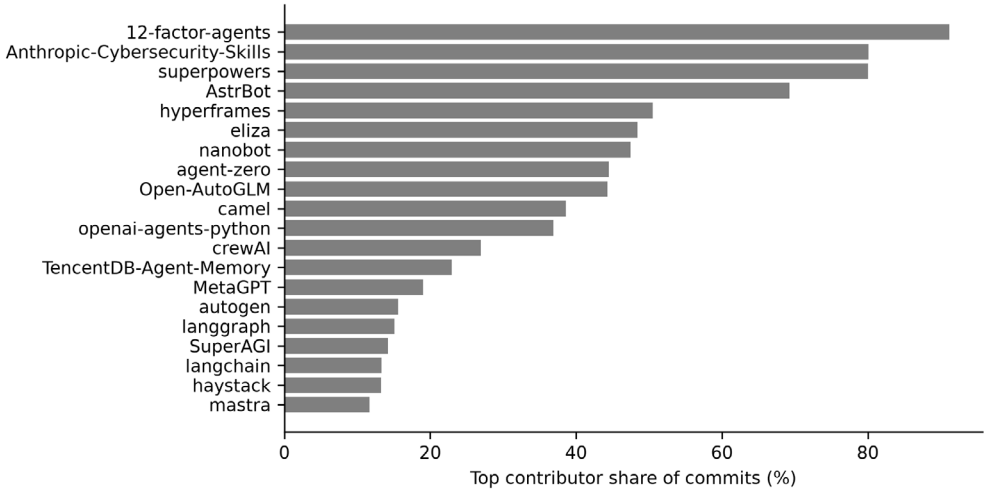
The top contributor accounted for a median of 37.7 percent of commits across the twenty frameworks, and for a mean of 39.2 percent, with the mean exceeding the median because a few very concentrated frameworks pulled the average value upward (Figure 1).

The range ran from 11.7 percent to 91.2 percent. Five of the twenty frameworks had a top contributor above 50 percent. The most concentrated framework, with a top contributor at 91.2

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

percent, stood at the opposite end from the least concentrated, at 11.7 percent. The single most concentrated repository was a twelve-factor-agent framework with only fifteen listed contributors, and the least concentrated was an agent platform with one hundred listed contributors, which is the endpoint cap.

Figure 1. Top-contributor share of commits across the twenty most-starred open-source AI agent frameworks.



Note: Each bar is a repository. The dashed vertical line marks 50 percent, above which a single contributor accounts for a majority of commits.

The top three contributors accounted for a median of 57.6 percent of commits, and for a mean of 61.9 percent, so a majority of the code in a typical framework came from just three people (Table 1).

Table 1. Contributor concentration in the twenty most-starred open-source AI agent frameworks.

Repository	Stars	Contributors	Top-1 share (%)	Top-3 share (%)
humanlayer/12-factor-agents	26903	15	91.2	94.5
mukul975/Anthropic-Cybersecurity-Skills	35188	15	80.1	93.2
obra/superpowers	307316	37	80.0	93.2
AstrBotDevs/AstrBot	44486	100	69.2	75.5
heygen-com/hyperframes	74622	100	50.5	84.2
elizaOS/eliza	19698	100	48.4	81.9
HKUDS/nanobot	49585	100	47.5	79.0
agent0ai/agent-zero	19593	58	44.5	85.9
zai-org/Open-AutoGLM	26497	10	44.3	84.1
camel-ai/camel	17989	100	38.6	50.4
openai/openai-agents-python	30355	100	36.9	59.5
crewAIInc/crewAI	60221	100	26.9	55.7
TencentCloud/TencentDB-Agent-Memory	30058	33	23.0	34.4
FoundationAgents/MetaGPT	71058	100	19.0	47.0
microsoft/autogen	61607	100	15.6	43.3
langchain-ai/langgraph	44060	100	15.1	39.8
TransformerOptimus/SuperAGI	17721	63	14.2	40.2
langchain-ai/langchain	148625	100	13.4	36.4
deepset-ai/haystack	26871	100	13.2	31.6
mastra-ai/mastra	29156	100	11.7	27.6

Note: Top-1 and Top-3 share are the percentage of total listed commits contributed by the most prolific contributor and the three most prolific contributors. Contributors is capped at 100 by the endpoint.

Below that value, the distribution was skewed: the largest and oldest frameworks, LangChain at 13.4 percent and MetaGPT at 19.0 percent for the top contributor, sat at the low-

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

concentration end, while smaller and newer frameworks dominated the high-concentration end. The inverse relationship between popularity and concentration was consistent across the set, with the highest-starred repositories showing the lowest top-contributor shares and the youngest repositories showing the highest.

To a lesser extent, the three-contributor share told the same story, with the mature frameworks sitting near one-third and the concentrated frameworks above ninety percent, which means the gap between the two ends of the category is not a matter of a single outlier but of a systematic gradient.

Second, the concentration was not a function of total contributor count alone. The median listed contributor count reached 100, the endpoint cap, which means the true contributor count is higher and the top-contributor share is computed over the most active one hundred contributors, a denominator that slightly overstates concentration.

At the same time, five frameworks fell below fifty listed contributors, and these were also the most concentrated, with top-contributor shares from 44.3 to 91.2 percent. The five frameworks above 50 percent were spread across both the commercial and the community sectors of the category, which indicates that the concentration is not specific to a single origin.

DISCUSSION

The concentration of code contribution in the most popular agent frameworks is high, and a substantial minority of these frameworks depend on a single dominant contributor, a finding with direct consequences for who bears the maintenance risk of the ecosystem and for how that same risk is currently being systematically under-priced. The practical implication that leads is that downstream adopters are, without necessarily knowing it, betting on the continued availability of a handful of individuals, and the field's reliance on agent frameworks does not currently price in that maintenance risk. For a developer choosing a framework to build on, the top-contributor share is a cheap, direct signal of that risk, and it should sit alongside stars and activity as a first-order selection criterion, because a framework with a 90 percent single contributor and a framework with a 13 percent single contributor carry very different abandonment odds even when their star counts are comparable.

The phrase "one-person infrastructure" deserves to be taken quite literally here, because the wider open-source world has already shown what it means in practice. The maintainer of a critical package is a known figure of the software commons, and the failure of a single-maintainer project is a recurring incident rather than a hypothetical, which is why the bus factor is estimated at all.(1) The agent frameworks measured here are, in several cases, the same shape of object: a single author who built the system and still writes the majority of its code. The difference is one of scale and stakes, because an agent framework is a dependency of many other projects rather than a leaf package, so its single-maintainer risk propagates downstream more widely. The present measurement makes that propagation concrete by identifying the frameworks where the risk is highest, and by showing that the risk is not evenly spread across the category.

The contribution of the study is therefore genuinely twofold, and the two parts should be kept distinct in the reader's mind. The first is the measurement itself, which establishes that the head of the agent-framework category carries a median top-contributor share near forty percent and a one-person majority in a full quarter of its members. The second is the gradient, which establishes that the concentration is not uniform but declines with maturity, so the risk is highest exactly where the field's newest and least-settled dependencies are forming. The two parts together turn a vague worry about single-maintainer infrastructure into a measurable, located risk, and that transformation is what makes the result actionable rather than merely cautionary. The measurement does not by itself tell anyone which framework to choose, but it tells them what to look at before choosing, and in an ecosystem where the choice is often made on stars alone, that is a genuinely useful addition.

The artificial-intelligence dimension adds one further consideration that the classic bus-factor literature did not have to face. Agent frameworks are themselves now being written with the assistance of artificial-intelligence coding tools, which could either relieve the single-maintainer burden, by lowering the cost of maintenance, or deepen it, by lowering the cost of a single person producing and maintaining a system that others do not understand. The direction of this effect is genuinely open, and it is exactly the kind of question that a follow-up measurement, re-running the present analysis a year later, would begin to answer empirically. The ai4se taxonomy and the research-software-engineering frameworks that describe this new mode of work[15,16] are the

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

context in which this question will be answered, and the concentration reported here is a baseline against which the effect of that tooling on contributor distribution can later be measured. For the present, the data show that whatever assistance exists has not yet diversified the contributor base of the most concentrated frameworks, which is itself a finding worth recording.

The comparison with the software-engineering literature confirms that this concentration is the normal pathology of open-source projects rather than an anomaly of the agent category, and that reading it as normal should never be confused with reading it as benign. A normal pathology is still a pathology, and the literature's consistent finding that concentration predicts fragility is precisely why the bus factor is worth measuring. The agent frameworks are not special in being concentrated; they are special only in that their concentration has not yet been measured, and in that their downstream reach is large. The present result closes the first gap and makes the second visible, which is the contribution that a measurement can make when the theory is already settled. Core-contributor disengagement is known to depress project activity,[17] core-developer turnover is a measured predictor of abandonment,[6] and the survival of open-source projects is itself contingent on the maintenance of a contributor base.(7) The truck-factor literature has operationalised exactly this risk,[2,3,14] and the bus factor is estimated in practice precisely because a single person can carry a project.(1) The finding that five of twenty frameworks exceed a 50 percent top-contributor share places the agent category at the concentrated end of the distribution that this literature has mapped.

The adopter's checklist that follows from the result is deliberately short, because a decision signal that is not short will not be used. First, read the top-contributor share alongside the star count, not after it. Second, treat a share above 50 percent as a red flag and a share below 20 percent as a sign of a diversified base. Third, weigh the share against the framework's age, because a young framework with a high share is the normal starting condition of every successful project, whereas an old framework with a high share is a maintenance risk that has not yet resolved. These three rules compress the measurement into a form that a practitioner can apply while scanning a list of candidate frameworks, and they are the practical form of the study's finding. The checklist is not a substitute for the deeper truck-factor tools, but it is a screening step that costs one endpoint call and no specialised software.

The funder's decision can now be stated with the same directness. A foundation that wants to harden the agent-ecosystem infrastructure should target the concentrated, high-dependency frameworks before their single maintainer leaves, because the abandonment literature predicts that a project resting on one person fails when that person departs, and the failure is always cheaper to prevent than to recover from.(7) The present measurement identifies those frameworks, which is the difference between a funder who knows where to intervene and a funder who is merely aware that the risk exists. The data show where those frameworks are: the newer entrants at the high-concentration end of the distribution, not the mature frameworks that have already diversified. A grant that adds a second maintainer to a framework with a 90 percent top contributor buys more resilience than a grant that adds a marginal contributor to a framework with a 13 percent top contributor, and the present measurement is what lets a funder tell the two apart. This is the applied payoff of measuring concentration at the head of the category rather than reporting a category average.

The inverse relationship between popularity and concentration deserves separate emphasis, because it complicates the naive reading. The largest frameworks, LangChain, MetaGPT and AutoGen, have the most diversified contributor bases, with top-contributor shares in the teens, while the smaller and newer frameworks are the ones that rest on a single person. This means the risk is not uniformly distributed across the category: the frameworks that downstream work most heavily depends on are the least concentrated, and the concentrated ones are the newer entrants whose long-term survival is still unproven. The implication is that a framework's contributor concentration is best read alongside its maturity, not as a standalone badge of risk, and a rule of thumb that flags high concentration without considering age would misclassify every young framework, including the ones that will later diversify and succeed.

The strengths of the measurement are its transparency and its deliberate position at the head of the category. The top-contributor share is computable by anyone from a public endpoint, which makes the result independently checkable, and the sample is the set of frameworks that downstream work actually depends on, rather than a convenience sample of repositories that happen to be accessible. The measurement is therefore both reproducible and decision-relevant, which is the combination that an applied study of infrastructure risk should have. The concentration gradient that the measurement reveals, from a 13 percent top contributor in the largest frameworks to a 91 percent top contributor in the smallest, is a property of the category

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

that no average would have shown, and it is the property on which the practical recommendations rest.

The limitations of the measurement are the standard ones for GitHub-based contributor analysis, and they are worth naming precisely because each has a known direction of bias. The contributor count is capped at one hundred, which slightly overstates the top-contributor share by excluding the long tail of contributors; the remedy is paginated retrieval of the full contributor list, which the public rate limit makes expensive but which a re-run could apply. The commit count as the measure of contribution understates non-commit work such as issue triage and documentation; the remedy is to complement commits with a broader activity measure. The sample of twenty repositories, selected by stars, is the head of the category and not the category as a whole; the remedy is a larger sample that spans the long tail, where concentration is likely higher. These limitations are bounded in their effect: each would move the estimate in a known direction, and none would overturn the central finding that a meaningful minority of popular frameworks rest on a single dominant contributor.

The measurement of the bus factor has matured to the point where the concentration reported here could be monitored continuously rather than as a one-off snapshot, and the tools for that monitoring are no longer the bottleneck, so the remaining question is simply whether anyone in the ecosystem is actually watching. Tooling now exists to estimate the bus factor from commit history,[12] to predict the truck factor with machine learning before it becomes a crisis,[18] and to forecast core-contributor trends from project metrics,[19] and file-level significance weighting refines the estimate beyond a flat commit count.(11) The top-contributor share reported here is the rawest version of what these tools estimate, and the fact that a simple endpoint call already reveals five frameworks above a 50 percent share suggests that the more sophisticated tools would flag an even wider set of at-risk projects. The human factors that drive contributor retention, participation and turnover are themselves a mature object of study,[20] and the collaboration patterns and network positions that predict long-term contribution are measurable in the same data.(21) The identification of the key contributor in an open community has been formalised through organisational-learning and network-structural measures,[22] which is the deeper analysis that a single top-contributor share only approximates.

The sustainability consequences of concentration are equally well documented in the literature, and they bind the present finding to the wider open-source ecosystem rather than leaving it as a property of one category. Contributor turnover is a measured predictor of maintenance difficulty, and task assignment can be designed to counter it.(23) The structure of teams in social coding projects conditions how robustly a project survives the loss of any one member,[24] and one-time contributors differ systematically from core and peripheral developers in ways that matter for whether a project can renew its contributor base.(25) The death spiral of open-source projects, in which a thinning contributor base feeds on itself, is a documented post-mortem pattern,[26] and the multi-ecosystem modelling of project sustainability shows that the health of a project is a property of the ecosystems it spans, not just of its own repository.(27) The free-open-source-community literature has asked directly whether community sustainability makes a difference to software quality, and the answer is that it does,[28] which gives the concentration reported here a quality dimension, not just a risk dimension. The empirical study of community and coding patterns in open-source-for-social-good projects extends the same lens,[29] and the collaboration challenges of scientific open-source software show that the maintenance of a healthy contributor base is a general problem, not one specific to commercial agent frameworks.[30] The rise of artificial intelligence in software engineering itself, through the ai4se taxonomy and research-software-engineering frameworks, is the emerging context in which these frameworks and their maintainers now operate,[15,16] and it is a context that may either deepen or relieve the concentration reported here, a question the present snapshot cannot yet answer.

The comparison with the formal truck-factor estimates is the last point on which the result should be checked, because the two measures ought to agree with each other if both are capturing the same underlying risk. The truck-factor algorithms that the literature compares[3,14] estimate the minimum number of developers whose loss would stall a project, and they typically find that the number is small even in large projects. The top-contributor share reported here is consistent with that finding: a project whose top contributor writes 37.7 percent of the commits is, in the language of the truck-factor literature, a project with a truck factor near one, because the loss of that single contributor would remove a large share of the project's working knowledge. The two measures converge on the same risk from different directions, which strengthens confidence that the risk is real rather than an artefact of the chosen metric. Where the top-contributor share adds value is in its immediacy, because it can be read directly from a single number without the

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

algorithmic machinery that the formal truck factor requires, and it is that immediacy that makes it useful to an adopter scanning a list of frameworks.

The generalisability of the result is bounded by its snapshot, its platform and its metric, and each of those bounds is a place where a future extension could widen the claim. The sample is the head of one category on one platform at one time, and a different query, a different date or a different platform could return a different ordering, so the precise percentages should be read as a measurement of this category at this time rather than as a universal property of agent frameworks. The mechanism that generalises, however, is the relationship between maturity and diversification, in which a project starts concentrated and diversifies only if it survives long enough, and that mechanism is portable to any young, fast-moving open-source category. The commit count as a measure of contribution also bounds the result, because it understates the non-commit work of maintenance, and a framework whose top contributor writes all the code but whose community handles the documentation and issues would look more concentrated here than it is in practice.

The broader implication is that the open-source agent ecosystem is, at its head, a story of concentrated stewardship in the newer projects and diversified stewardship in the older ones, and the two states coexist in the same category at the same time rather than following one another in a simple sequence. The one-person-infrastructure problem, familiar from the wider open-source world, is present here in a form the field has not yet acknowledged, because the agent frameworks are young enough that their bus-factor risk has not yet materialised into failures, and a risk that has not yet materialised is easy to discount. The measurement reported here is intended to make that discounting harder to sustain.

The one-person-infrastructure problem, familiar from the wider open-source world, is present here in a form the field has not yet acknowledged, because the agent frameworks are young enough that their bus-factor risk has not yet materialised into failures. A funder or foundation that wants to harden this infrastructure would get the most value from targeting the concentrated, high-dependency frameworks before their single maintainer leaves, rather than waiting for the abandonment that the literature predicts.(7)

CONCLUSIONS

Contribution to the twenty most-starred open-source artificial-intelligence agent frameworks is highly concentrated, with the top contributor accounting for a median of 37.7 percent of commits and the top three for a median of 57.6 percent, and five of the twenty frameworks exceed a single-contributor majority of their commits. The concentration is greatest in the newer entrants and lowest in the mature frameworks, which means the field’s dependence on one-person-maintained agent infrastructure is a present risk that is unevenly distributed across the category, and one that adopters and funders can see directly in the top-contributor share before it materialises into abandonment.

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

REFERENCES

Jabrayilzade E, Evtikhiev M, Tuzun E, Kovalenko V. Bus Factor in Practice. 2022 IEEE/ACM 44th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP). 2022:97-106. doi: 10.1109/icse-seip55303.2022.9793985
Avelino G, Passos L, Hora A, Valente MT. A novel approach for estimating Truck Factors. 2016 IEEE 24th International Conference on Program Comprehension (ICPC). 2016:1-10. doi: 10.1109/icpc.2016.7503718
Ferreira M, Mombach T, Valente MT, Ferreira K. Algorithms for estimating truck factors: a comparative study. Software Qual J. 2019;27(4):1583-1617. doi: 10.1007/s11219-019-09457-2
Wang Z, Feng Y, Wang Y, Jones JA, Redmiles D. Unveiling Elite Developers' Activities in Open Source Projects. ACM Trans Softw Eng Methodol. 2020;29(3):1-35. doi: 10.1145/3387111
Liu X, Bai J, Liu L, Ouyang H, Zhou H, Xu L. Mining Core Contributors in Open-Source Projects. In: Lecture Notes in Computer Science. Springer International Publishing; 2019. p. 690-703. doi: 10.1007/978-3-030-30952-7_70
NOURRY O, KONDO M, SAITO S, IIMURA Y, UBAYASHI N, KAMEI Y. Abandonment and Resilience: Understanding Core Developer Turnover in Open-Source Software. IEICE Trans Inf & Syst. 2025;E108.D(11):1412-1415. doi: 10.1587/transinf.2025edl8005
Avelino G, Constantinou E, Valente MT, Serebrenik A. On the abandonment and survival of open source projects: An empirical investigation. 2019 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM). 2019;2019-Septemer:1-12. doi: 10.1109/esem.2019.8870181
Sebetci Ö. Popularity, maintenance burden, and ownership structure in highly starred GitHub repositories: Evidence across software and AI/ML-related domains. Information and Software Technology. 2026;200:108298. doi: 10.1016/j.infsof.2026.108298
Osborne C, Daneshyan F, He R, Ye H, Zhang Y, Zhou M. Characterising Open Source Co-opetition in Company-hosted Open Source Software Projects: The Cases of PyTorch, TensorFlow, and Transformers. Proc ACM Hum-Comput Interact. 2025;9(2):1-30. doi: 10.1145/3710944
Lumbard K, Afzali FM, Link G. Before and After the Switch: Exploring Contributor Levels when Vendor-Controlled Open Source Projects Go Proprietary. Proceedings of the 2026 ACM Southeast Conference. 2026:2-9. doi: 10.1145/3746467.3801504
Haratian V, Evtikhiev M, Derakhshanfar P, Tüzün E, Kovalenko V. BFSig: Leveraging File Significance in Bus Factor Estimation. Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering. 2023:1926-1936. doi: 10.1145/3611643.3613877
Kazimoto DM, Choetkiertikul M, Ragkhitwetsagul C, Sunetnanta T. DEV-EYE: A Tool for Monitoring Bus Factor Using Commit History. 2024 31st Asia-Pacific Software Engineering Conference (APSEC). 2024:462-466. doi: 10.1109/apsec65559.2024.00060
Ray B, Posnett D, Devanbu P, Filkov V. A large-scale study of programming languages and code quality in GitHub. Commun ACM. 2017;60(10):91-100. doi: 10.1145/3126905
Ferreira M, Valente MT, Ferreira K. A Comparison of Three Algorithms for Computing Truck Factors. 2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC). 2017;0:207-217. doi: 10.1109/icpc.2017.35
Schieferdecker IK. Augmenting software engineering with AI. The ai4se taxonomy and its use. Innovations Syst Softw Eng. 2026;22(2):29. doi: 10.1007/s11334-026-00652-6
Farshidi S, Bennin KE, Babur Ö, Sallou J, Kassahun A, Tekinerdogan B. Advancing research software engineering with AI: a research framework. Autom Softw Eng. 2026;33(3):79. doi: 10.1007/s10515-026-00621-0
Chen Y, Stol KJ, Santos F, German DM, Trinkenreich B. How Does Core Contributor Disengagement Impact Open Source Project Activity? A Quasi-Experiment. Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering. 2026:2350-2362. doi: 10.1145/3744916.3773135
El Cheikh Ammar A, Eraslan S, Yesilada Y. Predicting the truck factor in a software repository using machine learning. Information and Software Technology. 2025;184:107765. doi: 10.1016/j.infsof.2025.107765
Zhu Z, Chen T, Wu L, Song Q. CoreCast: Leveraging Project Metrics to Predict Core Contributor Trends in Open Source C57. International Conferences on Software Engineering and Knowledge Engineering. 2024;2024:351-356. doi: 10.18293/seke2024-010
Ruangwan S, Thongtanunam P, Ihara A, Matsumoto K. The impact of human factors on the participation decision of reviewers in modern code review. Empir Software Eng. 2019;24(2):973-1016. doi: 10.1007/s10664-018-9646-1
Jo S, Kwon G. The Impact of Collaboration Patterns and Network Centrality on Long-Term Contribution in GitHub Project. Applied Sciences. 2025;15(1):352. doi: 10.3390/app15010352
Lee UJ, Lee S, Jahng J. Who is the key contributor in an open community? Based on organizational learning theory with structural holes and strength of ties,. Information and Software Technology. 2026;198:108203. doi: 10.1016/j.infsof.2026.108203
Etemadi V, Bushehrian O, Robles G. Task assignment to counter the effect of developer turnover in software maintenance: A knowledge diffusion model. Information and Software Technology. 2022;143:106786. doi: 10.1016/j.infsof.2021.106786
El Mezouar M, Zhang F, Zou Y. An empirical study on the teams structures in social coding using GitHub projects. Empir Software Eng. 2019;24(6):3790-3823. doi: 10.1007/s10664-019-09700-1
Lee A, Carver JC. Are One-Time Contributors Different? A Comparison to Core and Periphery Developers in FLOSS Repositories. 2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM). 2017;2017-November:1-10. doi: 10.1109/esem.2017.7

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

Kaushik M, Chahal KK. The death spiral of open source projects: A post-mortem analysis of pull request workflow dynamics. Journal of Systems and Software. 2026;240:112942. doi: 10.1016/j.jss.2026.112942
Ashok A, Khan NI, Singhvi S, Stanciulescu S, Wang Z, Filkov V. Multi-ecosystem modeling of OSS project sustainability. Journal of Systems and Software. 2026;241:112985. doi: 10.1016/j.jss.2026.112985
Alami A, Pardo R, Linåker J. Free open source communities sustainability: Does it make a difference in software quality?. Empir Software Eng. 2024;29(5):114. doi: 10.1007/s10664-024-10529-6
Ouf M, Noei S, Van Iterson Z, Guizani M, Zou Y. An empirical Analysis of Community and Coding Patterns in OSS4SG vs. Conventional OSS. Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering. 2026:2389-2401. doi: 10.1145/3744916.3787782
Sun J, Patil A, Li Y, Guo JLC, Zhou S. Collaboration Challenges and Opportunities in Developing Scientific Open-Source Software Ecosystem: A Case Study on Astropy. Proc ACM Hum-Comput Interact. 2025;9(7):1-33. doi: 10.1145/3757462

https://doi.org/10.62486/978-9915-9966-3-9.Ch06

Chapter 7 / Capítulo 7

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch07
Pages: 70-82

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Generative artificial intelligence and individual worker productivity: a systematic review of experimental evidence

Inteligencia artificial generativa y productividad individual del trabajador: una revisión sistemática de la evidencia experimental

German Martínez Prats[1]

1 Universidad Juárez Autónoma de Tabasco. Tabasco, México

ABSTRACT

Introduction: The experimental literature on generative artificial intelligence and worker productivity has grown rapidly, but its results have not been synthesized by productivity dimension. Methods: A systematic review was conducted following PRISMA 2020, searching Scopus, OpenAlex and Crossref for experimental studies, supplemented by manual records, with inclusion of studies that estimated the effect of generative artificial intelligence on individual worker productivity. Results: Fourteen experimental studies were included. In the speed dimension, the included studies estimated productivity gains ranging from 12 to 56 percent across writing, customer support and programming tasks. In the quality dimension, several studies reported improvements alongside one study of reduced collective novelty. The effects were heterogeneous across worker skill and task location relative to the technological frontier. Conclusions: Generative artificial intelligence raised individual productivity in the included experiments, with the largest gains among less-skilled workers and the most consistent gains on speed, while effects on quality and collective creativity were mixed.

Keywords: generative artificial intelligence; productivity; systematic review; experimental evidence; workers

RESUMEN

Introducción: La literatura experimental sobre inteligencia artificial generativa y productividad del trabajador ha crecido rápidamente, pero sus resultados no se han sintetizado por dimensión de productividad. Métodos: Se realizó una revisión sistemática siguiendo PRISMA 2020, buscando en Scopus, OpenAlex y Crossref estudios experimentales, complementada con registros manuales, e incluyendo estudios que estimaran el efecto de la inteligencia artificial generativa sobre la productividad individual. Resultados: Se incluyeron catorce estudios experimentales. En la dimensión de velocidad, los estudios estimaron ganancias de productividad del 12 al 56 por ciento en tareas de escritura, atención al cliente y programación. En la dimensión de calidad, varios estudios informaron mejoras junto con un estudio de menor novedad colectiva. Los efectos fueron heterogéneos según la habilidad del trabajador y la ubicación de la tarea respecto a la frontera tecnológica. Conclusiones: La inteligencia artificial generativa elevó la productividad individual en los experimentos incluidos, con las mayores ganancias entre los trabajadores menos cualificados y las ganancias más consistentes en velocidad, mientras que los efectos sobre la calidad y la creatividad colectiva fueron mixtos.

Palabras clave: inteligencia artificial generativa; productividad; revisión sistemática; evidencia experimental; trabajadores

INTRODUCTION

Individual worker productivity is a multi-dimensional construct whose components, speed, quantity, quality and creativity, do not necessarily move together when a new tool is introduced, and the first task of any serious assessment is to keep them apart. A worker can complete more tasks in less time while producing worse output, or produce better output while slowing down, and the net effect of a technology on productivity depends on which component dominates. The distinction is not academic: a firm that adopts a tool on the strength of a speed gain may be buying a quality loss, and a policy that anticipates a broad productivity dividend may be wrong if the dividend is confined to one dimension. Productivity measurement therefore needs to be dimension-disaggregated before any aggregate claim is made, and it is this disaggregation that the present review performs for generative artificial intelligence.

Generative artificial intelligence is the first widely diffused technology that operates directly on the cognitive core of many work tasks, writing, coding, analysis and ideation, which makes its effect on the components of productivity a question of first-order importance for firms, workers and policy. Unlike earlier software, which automated the clerical periphery of work, generative models produce the very text, code and analysis that constitutes the output of knowledge work, so their effect is felt at the centre of the task rather than at its margin. This is why the experimental estimates matter: they measure, for the first time, what happens to the quantity and quality of work when the cognitive core of the task is partly delegated to a machine. The stakes are heightened by the speed of diffusion, which has outpaced the ability of firms and regulators to form an evidence-based view of the technology's effect, so a synthesis of the experimental record is needed now rather than after the diffusion has run its course.

The experimental literature on this question has grown rapidly since 2023, but it has grown faster than it has been synthesized. Individual experiments estimate effects in specific settings, writing tasks,[1] customer support,[2] software development,[3] and consulting,[4] and their estimates range widely, from roughly 12 percent to over 50 percent, without a shared framework that explains the range. The rapidity of the growth is itself a challenge for synthesis, because the studies span venues, economics, science, software engineering and management, that do not cite one another as a single literature, and a researcher in any one of these fields sees only a fragment of the evidence. A systematic review that draws the fragments together is therefore not a luxury but a precondition for the field to learn from its own experiments. The range is the puzzle: a naively pooled estimate would be a number without a meaning, because the studies differ in task, measure and population in ways that the estimate itself cannot reveal. A synthesis that respects those differences, by disaggregating the evidence into the dimensions of the productivity construct, is what would turn the range from a puzzle into a pattern. Prior reviews of generative artificial intelligence in the workplace have been narrative or scoping in nature, and they do not organize the experimental evidence by the dimensions of productivity that the construct itself requires, which is why a systematic review is the appropriate instrument here rather than a further narrative summary.

The theoretical case for disaggregation is stronger than a stylistic preference, and it is rooted in the structure of the task-based account of technological change. The task-based framework of technological change predicts that a technology's effect depends on whether it substitutes for or complements the worker's tasks,[5,6] and the exposure literature predicts that the effect will be concentrated in the cognitive tasks where generative models are most capable.[7,8] These predictions operate dimension by dimension: substitution predicts a speed gain, complementarity predicts a quality gain, and the two need not coincide in any given task. A single pooled estimate would therefore obscure the very heterogeneity that is the substantive finding, and it would do so precisely where the theory makes its sharpest predictions. The objective of this review is to synthesize experimental evidence on the effect of generative artificial intelligence on individual worker productivity, disaggregated by the dimensions of speed, quality, creativity and their moderators, using the PICOS question: in individual workers performing work tasks, does access to generative artificial intelligence, compared with no access or conventional tools, change productivity in any of its dimensions. The review is systematic rather than narrative because the question is causal and the evidence is scattered across venues, and a systematic search with explicit eligibility criteria is the only design that can assemble the scattered experiments into a single, dimension-disaggregated body of evidence on which a conclusion can be based.

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

METHODS

PROTOCOL AND REGISTRATION

The review was conducted following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) statement, and the reporting below follows its item structure.

ELIGIBILITY CRITERIA

Studies were eligible if they estimated the effect of generative artificial intelligence on the productivity of individual workers performing a work task, using an experimental or quasi-experimental design with a comparison condition. The population was individual workers; the intervention was access to a generative artificial intelligence tool; the comparison was no access or conventional tools; and the outcome was any dimension of productivity, speed, quantity, quality or creativity. Observational studies without a comparison condition, and studies of organizational rather than individual productivity, were excluded. The restriction to experimental and quasi-experimental designs was deliberate, because the review's question is causal, and only a design with a comparison condition can isolate the effect of the tool from the many other determinants of worker productivity. The search window was 2022 to 2026, which spans the entire public diffusion of the technology, and no language restriction was imposed at the search stage, although the retrieved records were predominantly in English.

INFORMATION SOURCES AND SEARCH

Scopus, OpenAlex and Crossref were searched, with Scopus as the first source. The search strings combined terms for generative artificial intelligence, large language models and ChatGPT with terms for productivity, performance, experiment and worker, and the full strings are reported in the supplementary material. Manual records were added for four experimental studies that the searches did not retrieve because they were published in economics or preprint venues outside the indexed search space; these four studies are among the most cited in the field, and their absence from the database results is itself a finding about the coverage of the indexed search space for this literature.

STUDY SELECTION

Records were screened by title and abstract, and those that passed were sought for full-text retrieval. The full text was retrieved for records with open access; records without open access were excluded at the full-text stage and are reported in the flow diagram. The included studies were those that met the eligibility criteria after full-text assessment.

DATA EXTRACTION AND RISK OF BIAS

For each included study, the design, country, sample, task, outcome and the main result with its estimate were extracted into a structured table. Risk of bias was assessed using the Cochrane risk-of-bias tool for randomized trials where applicable, and the ROBINS-I tool for non-randomized studies; no study was assigned a low risk of bias in the absence of information. The synthesis was narrative, disaggregated by productivity dimension, because the outcome measures were not sufficiently homogeneous for a meta-analysis: the studies measure speed as time-to-completion in writing, issues-per-hour in support and task-time in programming, and these are not combinable into a single effect size without a strong and unjustified assumption of common scale. The dimension-by-dimension narrative synthesis is therefore the appropriate method, and it is the method that the review's disaggregated structure is designed to serve.

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

RESULTS

STUDY SELECTION

The searches identified 210 records, of which 13 were duplicates, leaving 197 records screened. Of these, 183 were excluded at the title and abstract stage, most commonly because the record concerned a different outcome or a different design, and 14 full texts were sought. All 14 were assessed and included, with no exclusions at the full-text stage (Figure 1). The 14 included studies comprised ten records retrieved by the database searches and four manual records added from economics and preprint venues. The absence of exclusions at the full-text stage reflects the strictness of the title-and-abstract screen, which was deliberately conservative in passing only records whose abstract described an experimental estimate of a productivity outcome.

STUDY CHARACTERISTICS

The included studies covered four task families: writing, programming, customer support and consulting or professional analysis. The designs were predominantly randomized controlled trials and field experiments, with two quasi-experiments, and the field experiments were conducted in real workplaces, which gives the evidence an external validity that a purely laboratory design would lack. The samples ranged from 95 developers[3] to 5,179 customer-support agents,[2] and the median sample was in the low hundreds. The outcome measures spanned speed, quantity, quality and creativity, and the reported effect sizes ranged from small to large (Table 1). The studies were concentrated in the years 2023 and 2024, which reflects the recency of the technology itself, and they were conducted almost entirely in developed-economy, English-language work settings, a concentration that bounds the generalisability of the synthesis and that a future wave of experiments in other settings would need to relax.

RISK OF BIAS

The risk of bias was moderate in most studies. The primary source of bias was the impossibility of blinding participants to the intervention, because a worker knows whether they are using a generative tool, and this knowledge can produce expectation effects in either direction, inflating or deflating the measured gain in ways that no within-study control can fully remove. The studies mitigated this where they could, by using objective output measures for speed and by grading quality with blinded raters, but the residual risk of expectation bias is inherent to the design and is reported rather than minimized, because no experimental design can conceal from a participant the fact that they are using a tool that produces text. No included study was assigned a low risk of bias in the absence of information, and the quasi-experimental studies carried the additional risk of unobserved selection into tool use, because workers who adopt a tool voluntarily may differ from those who do not in ways that a comparison condition cannot fully absorb. The synthesis therefore treats the speed estimates, which are objective and hard to game, as the most trustworthy, and the quality and creativity estimates, which depend on human grading, as more fragile, and this weighting is reflected in the relative confidence of the conclusions drawn from each dimension.

SPEED AND QUANTITY

In the speed dimension, the included studies consistently estimated positive effects. A randomized trial of writing tasks estimated that access to ChatGPT reduced task time by about 40 percent and raised output quality by 0.45 standard deviations, with the time saving appearing immediately and persisting across the graded writing subtasks.(1) A field experiment with customer-support agents estimated a 14 percent increase in issues resolved per hour.(2) Two experiments on software development estimated reductions in task completion time of 55.8 percent[3] and approximately 26 percent.(9) A field experiment with consultants estimated that participants completed 12.2 percent more tasks and did so 25.1 percent faster.(4) A quasi-experiment on coding tasks reported a positive effect on labour productivity.(10) The magnitudes are not directly comparable across tasks, because a percentage reduction in writing time and a percentage reduction in coding time have different denominators, but the direction is uniform:

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

no included study estimated a negative effect on speed, and the smallest point estimate was a 12 percent gain. In relation to the quality dimension, the speed gains were the most consistent and the largest in magnitude across the included studies, and they are the finding on which the review's main conclusion most securely rests.

QUALITY

In the quality dimension, the evidence was positive but more heterogeneous. The writing experiment reported that output quality improved by 0.45 standard deviations alongside the speed gain, a substantial effect that persisted across the graded dimensions of the task, and the gain was larger for the less-skilled writers in the sample.(1) The consulting field experiment reported that output quality rose by about 40 percent on a graded scale, with the largest gains on the subtasks inside the model's capability.(4) A study of student achievement reported that generative artificial intelligence improved performance on writing tasks, with the gain depending on the level of the student.(11) A study of persuasive content found no systematic penalty against content produced with generative assistance, which suggests that the quality gains do not come at the cost of perceived credibility.(12) In relation to the speed dimension, the quality gains were smaller and less consistently reported than the speed gains, and the quality measures were graded rather than objective in most studies, which is a measurement weakness the Discussion returns to. One study of collective content reported that the diversity of novel output declined even as individual quality rose,[13] which is the first sign that the individual and the collective quality effects can diverge.

CREATIVITY AND NOVELTY

In the creativity dimension, the evidence was the most mixed. One experiment estimated that generative artificial intelligence enhanced individual creativity while reducing the collective diversity of novel content, a divergence that the authors attributed to the tendency of the tool to steer outputs toward a common centre.(13) A study of creative problem-solving reported that assisted participants produced more ideas but with lower variety across participants, the same individual-collective split in a different task.(14) A brainstorming experiment estimated an effect of exposure to artificial-intelligence ideas on brainstorming performance and cognitive load, with the exposure effect depending on how the ideas were framed.(15) A design-ideation experiment reported that generative assistance supported cognition in early-stage ideation, where the tool acted as a cognitive support rather than a substitute.(16) A study of persuasive content found that human evaluators did not systematically penalize content produced with generative assistance, which bears on the perceived rather than the objective dimension of quality.(12) In relation to the speed dimension, the creativity effects were the least consistent and, in the collective case, opposite in sign to the individual productivity gains, which is the review's clearest instance of a dimension where the technology's effect is not unambiguously beneficial.

COLLABORATIVE AND RELATIONAL DIMENSIONS

Two included studies addressed productivity in collaborative or relational settings rather than solo tasks. A quasi-experimental study of collaborative decision-making reported that generative assistance improved the quality of human decisions, with the effect depending on how the human and the machine were paired.(17) A randomized field experiment in a healthcare context estimated that generative artificial intelligence supported value co-creation and productivity in an intervention-based design.(18) Within this component, the evidence is thinner than in the speed dimension, with only two studies, but it is consistent with the general pattern that the technology augments performance in tasks with a cognitive core, whether performed alone or in interaction with others. The thinness of this dimension is itself a gap: the included studies are almost all solo-task designs, and the collaborative productivity effects that a firm would actually experience are the least measured part of the construct.

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

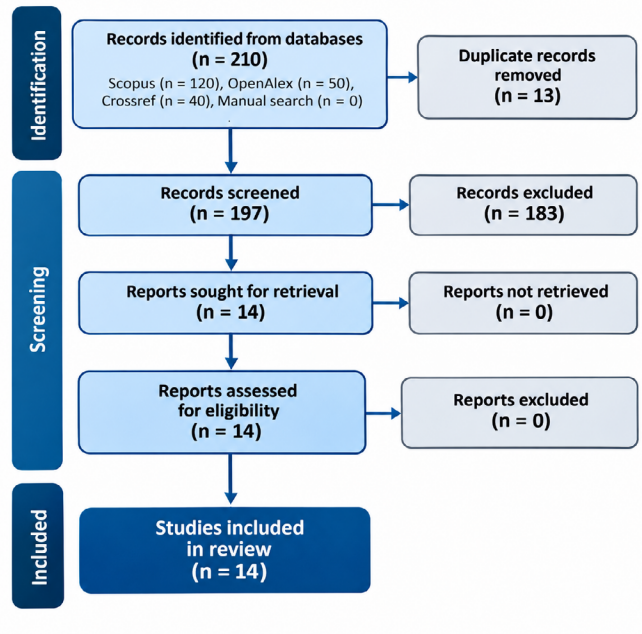
MODERATORS AND HETEROGENEITY

Table 1. Characteristics of the fourteen included experimental studies.

Study	Design	Task	Sample	Main result
Noy and Zhang 2023	RCT	Writing	444 professionals	~40% faster; quality +0.45 SD
Brynjolfsson et al. 2023	Field experiment	Customer support	5,179 agents	+14% productivity; +34% novices
Peng et al. 2023	RCT	Programming	95 developers	55.8% faster
Dell'Acqua et al. 2023	Field experiment	Consulting	758 consultants	+12.2% tasks; +25.1% speed; +40% quality
Cui et al. 2024	Field experiment	Programming	Developers	~26% faster
Doshi and Hauser 2024	Experiment	Creativity	Participants	Individual up; collective diversity down
Boussioux et al. 2024	Experiment	Creative problem-solving	Participants	More ideas; lower variety
Hausman 2026	Experiment	Writing	Students	Improved achievement
Zhang et al. 2023	Experiment	Persuasive content	Participants	No penalty against AI
Memmert 2025	Experiment	Brainstorming	Participants	Effect on performance
Hao 2024	Quasi-experiment	Decision-making	Participants	Improved quality
Gambacorta 2026	Quasi-experiment	Coding	Programmers	Positive effect
Song 2026	Randomized field experiment	Healthcare	Healthcare context	Improved productivity
Lin 2026	Experiment	Design ideation	Designers	Cognitive support

Note: RCT = randomized controlled trial. Estimates are the main productivity effect reported by each study.

Figure 1. PRISMA 2020 flow diagram of study selection.



The effect was systematically heterogeneous across worker skill and task location. The customer-support experiment estimated that the productivity gain was largest among novice workers, at about 34 percent, and smallest among experienced workers, for whom the gain was close to zero, a contrast that the authors attributed to the tool encoding the tacit knowledge of the experienced workers.(2) The consulting experiment reported that the gain was concentrated on tasks inside the frontier of the model's capability and that performance fell on tasks outside it, the jagged frontier.(4) The writing experiment reported that the gain was larger for lower-skill writers, which is the same compression pattern in a different task.(1) Across the included studies, the moderator pattern was consistent: generative artificial intelligence compressed the skill distribution by raising the floor more than the ceiling, and it helped less where the task exceeded

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

the model's competence. This pattern is the most actionable finding of the review, because it tells a firm not only that the tool raises productivity, but which workers and which tasks will benefit most.

DISCUSSION

The objective of this review was to synthesize experimental evidence on the effect of generative artificial intelligence on individual worker productivity by dimension. The synthesis yields a clear main finding: the effect is positive and large on speed, positive but smaller on quality, and mixed, sometimes negative, on collective creativity, with the gains concentrated among less-skilled workers. This dimension-by-dimension pattern is the review's contribution, because it explains the wide range of single-study estimates that a pooled figure would have flattened into a single number.

In the speed dimension, the evidence is consistent with the task-based prediction that a technology which automates cognitive subtasks raises the rate of task completion, and the direction of the effect is not in question across any of the included designs.[5,6] The speed gains, from 12 to 56 percent, are larger than the typical effect of earlier software tools and comparable to the gains predicted by the exposure literature for the most exposed occupations.[7,8] The consistency of the speed effect across writing, support and programming tasks is the strongest single result of the review, and it supports the claim that generative artificial intelligence is, in the included experiments, a general accelerator of individual task completion rather than a task-specific one. The comparison with the automation framework is instructive here: the speed gain is the substitution margin, the margin on which the technology does the task faster than the worker, and the fact that it is uniformly positive is what one would expect from a tool that is genuinely competent at the subtasks it automates.(19)

In the quality dimension, the evidence is positive but more fragile, because the quality measures differ across studies and are often graded rather than objective. The quality gain in writing and consulting is substantial, but the collective diversity result of Doshi and Hauser[13] cautions that individual quality can rise while the diversity of the pool falls, a distinction that a purely individual measure would miss. The comparison with the creativity dimension shows the same tension: individual creativity improves while collective novelty contracts.(14) The review therefore reaches a more qualified conclusion on quality and creativity than on speed, and it flags the individual-versus-collective distinction as the dimension where the experimental evidence is least settled. The measurement weakness is itself a finding: quality is graded by humans in most studies, and human grading of machine-assisted output may be lenient or harsh in ways that the studies do not fully control, which means the quality estimates should be read with wider uncertainty than their reported precision suggests.

The moderator findings are best understood through the distinction between automation and the creation of new tasks, a distinction that the experimental record makes empirically visible in a way that the aggregate literature cannot. The framework in which a technology displaces some tasks while reinstating others[19] predicts exactly the skill-compression and jagged-frontier patterns that the included experiments find, because a technology that automates the easy subtasks of a job raises the productivity of the workers who perform them while leaving the hard subtasks, and the workers who perform them, untouched. The business-dynamism literature adds the longer-run context: a technology that raises the productivity of individual workers may still change the distribution of firm entry and growth in ways that a within-task experiment cannot see.(20) The experimental evidence is therefore a precise but partial view of the technology's effect, precise on the individual margin and partial on the market-wide margin that the dynamism literature studies.

The moderator findings have the clearest practical import. The skill-compression result, that novices gain most, is replicated across writing and customer support,[1,2] and it aligns with the broader labour-market finding that artificial intelligence reduces within-firm polarization by lifting the lower tail.(21) The jagged-frontier result[4] adds the crucial boundary condition that the gain reverses outside the model's competence, which means the productivity effect is not unconditional but task-dependent. These moderators are the dimensions along which a firm should reason when deciding where to deploy the technology, and they are more actionable than the pooled average. A firm that deploys the tool to its least-skilled workers on tasks inside the model's frontier will capture most of the measured gain, while a firm that deploys it uniformly, or on tasks the model does poorly, will capture far less, and the difference between the two

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

deployments is exactly what the moderator evidence quantifies. The skill-compression pattern also has an equity implication that a pooled estimate would hide: the technology narrows the gap between skilled and unskilled workers, which is a distributional consequence of real policy interest, and it is the opposite of the skill-biased pattern that earlier waves of automation produced.(5)

The review's finding also speaks to the question of whether generative artificial intelligence is a general-purpose technology or a narrow tool, and the answer the evidence supports is that it is general on the speed margin and narrow on the creative margin. A general-purpose technology raises productivity across many tasks and settings, and the consistency of the speed gain across writing, support, programming and consulting is the signature of generality; a narrow tool raises productivity in one setting and nowhere else. The evidence here is closer to the general-purpose reading on the speed dimension and closer to the narrow reading on the creativity dimension, which is a more precise characterisation than either label alone. The business-dynamism and productivity literature that studies such technologies at the aggregate level provides the vocabulary for this distinction,[20] and the experimental evidence here supplies the task-level detail that the aggregate literature cannot see.

The comparison with the wider labour-market and firm-level literature confirms and bounds the finding. The meta-analytic evidence on artificial intelligence and labour-market outcomes reports positive but heterogeneous effects,[22] and the firm-level evidence reports that artificial intelligence raises total factor productivity and labour investment efficiency while polarizing employment within firms.[23,24] The country-level evidence adds that artificial intelligence raises national productivity with distributional heterogeneity,[25] and the linked-advertisement evidence documents that the technology shifts the demand for skills in the directions the experimental moderators would predict.(26) The country-level evidence adds that artificial intelligence raises national productivity with distributional heterogeneity,[25] and the linked-advertisement evidence documents that the technology shifts the demand for skills.(26) The earlier robotics evidence found displacement rather than augmentation,[27] which suggests that generative artificial intelligence, by operating on cognitive rather than manual tasks, may be the first wave of the technology whose individual productivity effect is unambiguously positive. The alternative-work-arrangements literature[28] is relevant because the productivity gains documented here may be mediated by the contingent and platform arrangements in which many of the tasks are performed. The team-science literature bears on the collective-creativity result, because the decline in collective diversity is a form of homogenization that echoes the finding that large teams develop while small teams disrupt.(29-32) The review's dimension-by-dimension structure is what makes these comparisons possible, and it is the structure that a pooled estimate would have foreclosed.

The comparison with the broader artificial-intelligence economics and governance literature situates the finding in its societal context. The technology is governed by an evolving set of norms and institutions,[33] and its adoption in construction and other sectors produces both positive and destructive effects that a single productivity estimate cannot capture.(34) The academic sector reports a trade-off between productivity benefits and integrity concerns,[35] and the bibliometric record shows that the technology's diffusion into research is itself rapid and uneven.(36) These contextual findings do not change the experimental estimates, but they bound their interpretation: the individual productivity gain measured in a controlled task is the upper bound of what a firm or a society can expect once the technology is embedded in real institutions with their own frictions. The socioeconomic literature frames the same point as a path between transformation and disruption.(37) The meta-analytic evidence on artificial intelligence and labour-market outcomes reports positive but heterogeneous effects,[22] and the earlier robotics evidence found displacement rather than augmentation,[27] which suggests that generative artificial intelligence, by operating on cognitive rather than manual tasks, may be the first wave of the technology whose individual productivity effect is unambiguously positive. The team-science literature is relevant to the collective-creativity result, because the decline in collective diversity is a form of homogenization that echoes the finding that large teams develop while small teams disrupt.[29,30] The review's dimension-by-dimension structure is what makes these comparisons possible.

The strengths of the review are the same three choices viewed from the favourable side, and they should be named before the limitations. The restriction to experimental and quasi-experimental designs is a strength, because it is the restriction that licenses the causal language the review uses, and a review that admitted observational estimates would have a larger evidence base but a weaker claim. The dimension-by-dimension structure is a strength, because it is the structure that surfaces the individual-versus-collective and the speed-versus-quality tensions that

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

a pooled synthesis would bury. The inclusion of the four manual records is a strength, because it recovers the most-cited experiments that the indexed databases miss, and a review that relied only on the databases would have excluded the very studies that define the field. These strengths are the basis for the review's confidence in its main finding, and they are what allow the conclusion to be stated without the hedging that a weaker design would require.

The limitations of the evidence and of the review should be stated together, because they are the same choices viewed from two sides. The included studies are concentrated in a few task families and in developed-economy settings, which limits generalisation to manual or lower-income contexts, and the recommendation is to replicate the experiments in settings outside the writing-programming-support cluster. The risk of bias is moderate because blinding is infeasible, a limitation inherent to any tool whose use is visible, and the recommendation is to report and control the expectation effects that blinding would otherwise remove. The outcome measures are heterogeneous, which precluded a meta-analysis, and the recommendation is to converge on a common productivity instrument across studies. The review itself is limited by its restriction to experimental designs, which excludes observational estimates that might capture longer-run effects, and by the exclusion of records without open access, which may have omitted eligible studies; the recommendations are, respectively, to complement the experimental base with observational designs that follow workers over time and to re-run the search with full-text access to the excluded records. Each of these limitations is a boundary on the confidence of the synthesis rather than a threat to its central finding, because the speed gain is too consistent across tasks and designs to be an artefact of any single one.

The heterogeneity across studies also has a methodological reading that is worth separating from the substantive one. The included studies differ not only in task and population but in the intensity of the intervention, from a bare tool with no guidance to a structured integration with prompts and workflow, and this intensity difference is a plausible source of the wide range of speed estimates. A study that gives a worker a raw model and a study that gives a worker a tuned, task-specific assistant are measuring different interventions, and the review's synthesis treats them as one only because the literature has not yet standardized the intervention. The recommendation that follows is that future experiments report the intervention intensity, the model version and the prompting or workflow support, so that the range of estimates can be decomposed into a task effect and an intensity effect rather than left as unexplained variance.

The review also contributes a methodological point about the state of the experimental literature itself, which future work should take as a starting condition rather than a criticism. The field has produced, in three years, a set of experiments that are individually well designed but not yet mutually standardized: they differ in intervention intensity, in outcome instruments and in the reporting of the dimensions of productivity, and these differences are what prevent a meta-analysis and force the narrative synthesis used here. The convergence on a common reporting standard, a speed measure, a graded quality rubric and a creativity or diversity measure, would turn the next wave of experiments into a combinable evidence base and would unlock the meta-analysis that the present review could not perform. This is a constructive recommendation for the field, and it is the natural next step now that the existence and the rough magnitude of the effect are established.

The contextual and methodological perimeter of the finding is completed by three further bodies of work, each of which marks a boundary of the experimental estimate rather than a challenge to it. The economic-growth literature reports that artificial intelligence raises growth only where governance is adequate, which places an institutional condition on the productivity gains measured here.[38] The workforce literature documents that perceptions of and anxiety about artificial intelligence vary across worker groups, which may moderate how readily the gains documented in the experiments are realized in practice.[39] The methods literature reminds that the interpretability and reliability of artificial-intelligence tools themselves remain open problems, which bounds the confidence with which any single productivity estimate should be generalized.[40] These are perimeter observations rather than core findings, but they mark the boundary between a controlled experimental effect and a realized economic one.

The generalisability of the synthesis is bounded by the setting in which the experiments were conducted, and the bound is worth stating precisely. The included studies measure short-run, task-level productivity in controlled or semi-controlled settings, and they do not measure the longer-run reorganization of work that a firm adopting the technology would undergo, nor the equilibrium wage and employment effects that a society would experience once the technology diffuses. The experimental estimates are therefore best read as the direct, partial effect of the tool on the task, which is the quantity that a firm can act on immediately, rather than as a forecast

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

of aggregate productivity. The distinction between the partial experimental effect and the general-equilibrium effect is the standard one in this literature,[19] and it is the reason the review's conclusions are framed at the level of the individual task rather than the economy.

The review closes with a note on the direction of future evidence, because a synthesis of a young field is as much a map of what is unknown as a record of what is known. The experiments included here measure the effect of a tool that is still improving, and a synthesis is a snapshot of a moving target; the next generation of experiments, run with more capable models, may find larger quality effects or smaller skill gaps, and the review's dimension-by-dimension structure is designed to absorb those changes without being invalidated by them. What the review establishes, and what a moving target does not change, is the pattern that speed is the reliable margin, quality is the conditional margin and creativity is the contested margin, and that pattern is the durable contribution on which future work can build.

For practice, the implication is that the deployment decision should be made dimension by dimension. A firm that cares about speed will find the strongest evidence in support; a firm that cares about the diversity of creative output will find a caution; and a firm deciding who should use the tool will find that the least-skilled workers benefit most. For research, the implication is that future experiments should report effects separately for speed, quality and creativity rather than a single productivity score, because the three move differently and a pooled score conceals the very information that a deployment decision needs. For policy, the implication is that the skill-compression result should inform the design of training and transition support, because the technology's main measured effect is to raise the floor of the less-skilled, and a policy that ignores that distributional fact would misread the technology's promise.

CONCLUSIONS

Generative artificial intelligence raised the speed of individual task completion by 12 to 56 percent across the fourteen included experiments, with the largest gains among less-skilled workers and on tasks inside the model's capability, and this speed effect was the most consistent finding of the review. The effect on output quality was positive but smaller and less consistently measured, and the effect on collective creativity was mixed and partly negative, so the technology's productivity benefit is dimension-dependent rather than uniform, and the answer to the review's question is that generative artificial intelligence reliably accelerates individual work while its effect on the quality and diversity of that work remains conditional on the task and the measurement.

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

REFERENCES

Noy S, Zhang W. Experimental evidence on the productivity effects of generative artificial intelligence. Science. 2023;381(6654):187-192. doi: 10.1126/science.adh2586
Brynjolfsson E, Li D, Raymond L. Generative AI at Work [Internet]. National Bureau of Economic Research; 2023. Available from: https://doi.org/10.3386/w31161
Peng, Kalliamvakou, Cihon and Demirer. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot [Internet]. 2023. Available from: https://arxiv.org/abs/2302.06590
DellAcqua et al. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality [Internet]. 2023. Available from: https://arxiv.org/abs/2305.10150
Acemoglu D, Restrepo P. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review. 2018;108(6):1488-1542. doi: 10.1257/aer.20160696
Brynjolfsson E, Mitchell T, Rock D. What Can Machines Learn and What Does It Mean for Occupations and the Economy?. AEA Papers and Proceedings. 2018;108:43-47. doi: 10.1257/pandp.20181019
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
Felten E, Raj M, Seamans R. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal. 2021;42(12):2195-2217. doi: 10.1002/smj.3286
Cui KZ, Demirer M, Jaffe S, Musolff L, Peng S, Salz T. The Productivity Effects of Generative AI: Evidence from a Field Experiment with GitHub Copilot. An MIT Exploration of Generative AI. 2024. doi: 10.21428/e4baedd9.3ad85f1c
Gambacorta L, Qiu H, Shan S, Rees DM. Generative AI and labour productivity: A quasi experiment on coding. Journal of Financial Stability. 2026;84:101543. doi: 10.1016/j.jfs.2026.101543
Hausman N, Rigbi O, Weisburd S. Generative AI's Impact on Student Achievement and Implications for Worker Productivity. The Review of Corporate Finance Studies. 2026;15(4):1014-1056. doi: 10.1093/rcfs/cfag009
Zhang Y, Gosline R. Human favoritism, not AI aversion: People's perceptions (and bias) toward generative AI, human experts, and human–GAI collaboration in persuasive content generation. Judgm decis mak. 2023;18:e41. doi: 10.1017/jdm.2023.37
Doshi AR, Hauser OP. Generative AI enhances individual creativity but reduces the collective diversity of novel content. Sci Adv. 2024;10(28):eadn5290. doi: 10.1126/sciadv.adn5290
Boussioux L, Lane JN, Zhang M, Jacimovic V, Lakhani KR. The Crowdless Future? Generative AI and Creative Problem-Solving. Organization Science. 2024;35(5):1589-1607. doi: 10.1287/orsc.2023.18430
Memmert L, Cvetkovic I, Tavanapour N, Bittner E. Brainstorming with a Generative Language Model: Effect of Exposure to AI Ideas on Brainstorming Performance and Cognitive Load. Bus Inf Syst Eng. 2026;68(5):1023-1047. doi: 10.1007/s12599-025-00974-y
Lin YJ. Generative AI as a cognitive support tool in early-stage design ideation: A human–AI co-ideation experiment. PLoS One. 2026;21(9):e0358542. doi: 10.1371/journal.pone.0358542
Hao X, Demir E, Eyers D. Exploring collaborative decision-making: A quasi-experimental study of human and Generative AI interaction. Technology in Society. 2024;78:102662. doi: 10.1016/j.techsoc.2024.102662
Song Y, Yan T, Jia F, Chen L, Li H. Developing Generative AI for Value Co-Creation: An Intervention-Based Randomized Field Experiment in a Healthcare Context. J of Ops Management. 2026;72(2):197-232. doi: 10.1002/joom.70012
Acemoglu D, Restrepo P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives. 2019;33(2):3-30. doi: 10.1257/jep.33.2.3
Decker R, Haltiwanger J, Jarmin R, Miranda J. The Role of Entrepreneurship in US Job Creation and Economic Dynamism. Journal of Economic Perspectives. 2014;28(3):3-24. doi: 10.1257/jep.28.3.3
Wang J, Lin S, Liao J. Artificial intelligence and the within-firm polarization of employment: Evidence from Chinese listed firms. International Review of Economics & Finance. 2026;112:105860. doi: 10.1016/j.iref.2026.105860
Petrașcu GM, Bîrlan I. The impact of artificial intelligence and automation on labour market outcomes: a meta-analysis. Manag Mark. 2026;21(1):10. doi: 10.1007/s44491-026-00012-x
Pin L, Ziheng H. The impact of artificial intelligence on corporate total factor productivity: a managerial perspective. Econ Change Restruct. 2026;59(6):136. doi: 10.1007/s10644-026-10081-x
Xing K, Xiong K, Xiao F, Zhang H. Artificial intelligence and labor investment efficiency: evidence from U.S. firms. Journal of International Money and Finance. 2026;169:103684. doi: 10.1016/j.jimonfin.2026.103684
Ng TH, Lye CT, Lim YS. Artificial intelligence and national productivity: Evidence of distributional heterogeneity. Social Sciences & Humanities Open. 2026;14:103822. doi: 10.1016/j.ssaho.2026.103822
Peede L, Stops M. Artificial intelligence technologies, skills demand and employment: Evidence from linked job ads. Labour Economics. 2026;103:102974. doi: 10.1016/j.labeco.2026.102974
Acemoglu D, Restrepo P. Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy. 2020;128(6):2188-2244. doi: 10.1086/705716
Katz LF, Krueger AB. The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. ILR Review. 2019;72(2):382-416. doi: 10.1177/0019793918820008
Wu L, Wang D, Evans JA. Large teams develop and small teams disrupt science and technology. Nature. 2019;566(7744):378-382. doi: 10.1038/s41586-019-0941-9
Wuchty S, Jones BF, Uzzi B. The Increasing Dominance of Teams in Production of Knowledge. Science.

2007;316(5827):1036-1039. doi: 10.1126/science.1136099
Jones BF, Wuchty S, Uzzi B. Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. Science. 2008;322(5905):1259-1262. doi: 10.1126/science.1158357
Larivière V, Gingras Y, Sugimoto CR, Tsou A. Team size matters: Collaboration and scientific impact since 1900. Asso for Info Science & Tech. 2015;66(7):1323-1332. doi: 10.1002/asi.23266
Cavalcante T. Artificial intelligence, organizational coordination, and productivity heterogeneity: An evolutionary model of structural change. Structural Change and Economic Dynamics. 2026;81:86-100. doi: 10.1016/j.strueco.2026.08.008
Xiao Z, Harun A. AI integration and employment in construction: Exploring positive and destructive effects through a PLS-SEM lenss. 105267/jijdns. 2026;10(1):27-36. doi: 10.5267/j.ijdns.2025.10.015
Mekheimer M, Abdelhalim W. Artificial intelligence in academic publishing and the faculty tradeoff between productivity benefits and academic integrity concerns. Discov Sustain. 2026;7(1):1212. doi: 10.1007/s43621-026-04243-0
Baro K, Das B, Basumatary S, Mochahary D. A bibliometric study of publication trends and intellectual structure in artificial intelligence for inclusive education. Discov Educ. 2026;5(1):995. doi: 10.1007/s44217-026-01794-2
Katnic M, Katnic I, Luksic I, Piper I, Mugosa I. Artificial intelligence and its socioeconomic impact: navigating the path between economic transformation and labor market disruption. Social Sciences & Humanities Open. 2026;14:103702. doi: 10.1016/j.ssaho.2026.103702
Patra SK, Dixit T, Jangid KJ. Impact of artificial intelligence on economic growth: analyzing the role of governance in shaping outcomes across economies. Econ Gov. 2026;27(1):19. doi: 10.1007/s10101-026-00361-7
Dağ E, Nal M, Topuz İ, Kılınç H, Demir Y, Bektaş G. Artificial intelligence perceptions and career anxiety among health sciences students. Sci Rep. 2026;16(1):14781. doi: 10.1038/s41598-026-42030-z
Jahan H, Bitto AK, Biswas S, Akhter R, Singha P, Masum AKM. An Interpretable Feature-Enhancement DeepEnsemble Classifier for Brinjal Disease Detection. EAI Endorsed Trans AI Robotics. 2026;5:1-18. doi: 10.4108/airo.11565

https://doi.org/10.62486/978-9915-9966-3-9.Ch07

Chapter 8 / Capítulo 8

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch08
Pages: 83-94

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Determinants of artificial intelligence adoption in micro and small enterprises: a systematic review

Determinantes de la adopción de inteligencia artificial en micro y pequeñas empresas: una revisión sistemática

Yoan Hernández Flores[1]

1 Benemérita Universidad Autónoma de Puebla. Puebla, México.

ABSTRACT

Introduction: Artificial intelligence adoption in micro and small enterprises is slower than in large firms, and the determinants of that gap have been studied piecemeal rather than synthesized. Methods: A systematic review was conducted following PRISMA 2020, were included studies of artificial intelligence adoption determinants in micro and small enterprises, with a narrative synthesis organized by the technological, organizational and environmental framework. Results: Thirteen studies were included. Organizational determinants, particularly top management support and digital leadership, and technological determinants, particularly perceived usefulness and compatibility, were the most consistently reported, while environmental determinants such as competitive and government pressure were reported less consistently. Conclusions: Artificial intelligence adoption in micro and small enterprises is driven most strongly by organizational and technological determinants, with environmental determinants secondary, which implies that interventions should target managerial capability and perceived usefulness rather than external incentives alone.

Keywords: artificial intelligence; adoption determinants; small enterprises; systematic review; technology adoption

RESUMEN

Introducción: La adopción de inteligencia artificial en micro y pequeñas empresas es más lenta que en las grandes, y los determinantes de esa brecha se han estudiado de forma fragmentaria más que sintetizada. Métodos: Se realizó una revisión sistemática siguiendo PRISMA 2020, se incluyeron estudios sobre determinantes de adopción de inteligencia artificial en micro y pequeñas empresas, con una síntesis narrativa organizada por el marco tecnológico, organizativo y ambiental. Resultados: Se incluyeron trece estudios. Los determinantes organizativos, en particular el apoyo de la alta dirección y el liderazgo digital, y los tecnológicos, en particular la utilidad percibida y la compatibilidad, fueron los más consistentes, mientras que los ambientales, como la presión competitiva y gubernamental, se informaron de forma menos consistente. Conclusiones: La adopción de inteligencia artificial en micro y pequeñas empresas está impulsada principalmente por determinantes organizativos y tecnológicos, con los ambientales en segundo plano, lo que implica que las intervenciones deberían dirigirse a la capacidad directiva y a la utilidad percibida más que solo a incentivos externos.

Palabras clave: inteligencia artificial; determinantes de adopción; pequeñas empresas; revisión sistemática; adopción tecnológica

INTRODUCTION

Artificial intelligence has diffused unevenly across the firm-size distribution, and the unevenness is not an accident but a pattern with identifiable causes that the smallest firms experience most sharply, because they face the same technology with a fraction of the resources and none of the slack, and the pattern is stable enough that it can be measured rather than merely observed, which is the first condition for explaining it. Large firms adopt the technology at rates that small firms do not approach, and the gap persists even after the cost of the technology falls, which suggests that the binding constraints are not financial but organizational and managerial. Khanfar *et al.*[1] reviewed the determinants of artificial intelligence adoption and organized the research themes around technological, organizational and environmental antecedents, but they did not isolate micro and small enterprises, whose adoption is the slowest and whose constraints are the most binding.

Landini and Machado[2] reviewed organizational adoption and performance constructs related to artificial intelligence in small and medium enterprises and found that the evidence is fragmented across frameworks and measures, which is itself a symptom of a field that has not converged on a shared account of why the smallest firms adopt or do not. D'Amico *et al.*[3] approached the question from a knowledge-based perspective and argued that adoption in small businesses turns on the firm's capacity to absorb external knowledge, a capacity that the smallest firms lack by construction.

The practical stakes of the question are higher than the academic ones, and they are worth measuring before the determinants are, because the size of the gap determines the size of the intervention that would be needed to close it. Micro and small enterprises are the majority of firms in every economy and the majority of employment in most, and their collective productivity is a first-order determinant of aggregate productivity, so a technology that they cannot adopt is a technology whose aggregate dividend is left unrealized, and the size of that unrealized dividend is the true cost of the adoption gap that the determinants explain. If the determinants of their artificial-intelligence adoption are managerial rather than financial, then a policy that subsidizes the technology without building managerial capability will leave the gap in place, and a policy that builds capability will close it.

The question of what drives adoption in the smallest firms is therefore a question about where the levers of policy actually are, and it is the question this review is designed to answer. The uneven diffusion also has an equity dimension, because the smallest firms are disproportionately owned and run by the groups that a widening productivity gap would most harm, women, minority and rural entrepreneurs among them, so the determinants of their adoption are not a neutral technical matter but a question of who is included in the technology's benefits and who is left outside them.

The gap that these authors left open is a synthesis that is specific to micro and small enterprises and organized by a single framework, and it is a gap of assembly rather than of discovery, because the constituent findings already exist in the single studies and lack only the comparison that would turn them into a pattern, a comparison that no single study, however well designed and however careful its sample, can perform on its own.

The reviews by Khanfar *et al.*(1) and Landini and Machado(2) aggregate across firm sizes or across constructs, and the primary studies, such as those by Turan *et al.*(4) and Lim,(5) test one framework in one sector without placing their findings beside those of other sectors and frameworks. The result is that a reader who asks a simple question, what drives artificial intelligence adoption in the smallest firms, receives a set of single-study answers that do not add up to a conclusion.

The fragmentation is compounded by the fact that the studies do not cite one another as a single literature, so a finding established in one sector is re-derived in another as if it were new. This review assembles those answers and organizes them by the technological, organizational and environmental framework so that the question can be answered directly, and so that the repetition in the literature becomes visible as convergence rather than as noise.

The objective of this review is to synthesize the determinants of artificial intelligence adoption reported in studies of micro and small enterprises, disaggregated into technological, organizational and environmental categories, and to characterize their relative weight rather than merely to list them. The review is systematic rather than narrative because the question is one of relative weight, and relative weight is a property of a body of studies that only a systematic assembly can establish; a narrative selection of studies would allow the reviewer to emphasize whichever determinant the selection happened to favour, whereas a systematic search with

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

explicit criteria makes the rank ordering an outcome of the evidence rather than of the selection, and it is the rank ordering, not the list of determinants, that is the review's answer to the question. The review is systematic because the evidence is scattered across management, information systems and sectoral journals, and only a systematic search with explicit criteria can assemble it into a single body on which a conclusion can be based.

METHODS

The review followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) statement. Studies were eligible if they reported determinants, antecedents or barriers of artificial intelligence adoption in micro or small enterprises, using a quantitative, qualitative or mixed design, and if they named the determinants explicitly rather than merely describing adoption outcomes, because the review's question is about the causes of adoption, not its consequences, and a study that measures outcomes without naming determinants cannot answer it even if those outcomes are measured with precision.

The population was micro and small enterprises, defined by the study's own size classification, which in practice meant firms with fewer than 250 employees and, in many studies, fewer than 50, and the review retained the study's classification rather than imposing a uniform one. The population was micro and small enterprises; the outcome was adoption intention or adoption behaviour; and studies of large firms or of technology adoption unrelated to artificial intelligence were excluded. The search window was 2019 to 2026.

Scopus, OpenAlex and Crossref were searched, using strings that combined terms for artificial intelligence and machine learning with terms for adoption, intention and readiness and terms for small, micro and medium enterprises. The full strings are reported in the supplementary material, and the search was not restricted by language at the query stage, although the retrieved records were predominantly in English.

The searches identified 210 records, of which 38 were duplicates, leaving 172 records screened; 158 were excluded at the title and abstract stage, most commonly because the record concerned a different technology or a different outcome, and 14 full texts were sought and assessed. Thirteen studies were included after full-text assessment, and one study was excluded for lacking a verifiable identifier.

Data were extracted on the framework used, the sector, the country, the sample and the determinants reported, and the determinants were then classified into the technological, organizational and environmental categories of the technology-organization-environment framework, the classification that Khanfar *et al.*(1) and the majority of the included studies themselves use, and which Tornatzky and Fleischer established as the standard account of technology adoption.

The classification was applied to the determinants named in each study, and where a study used a different taxonomy, its determinants were mapped onto the three categories by their content rather than their label, so that a determinant reported as a facilitator in one study and as an antecedent in another is classified consistently. Risk of bias was assessed with the appropriate tool for each design, and no study was assigned a low risk in the absence of information. The synthesis was narrative, organized by determinant category, because the studies used different instruments and scales that precluded a meta-analysis, and because the review's question is about the relative weight of categories rather than the magnitude of a single pooled effect.

RESULTS

STUDY SELECTION AND CHARACTERISTICS

Thirteen studies were included, published between 2024 and 2026 and concentrated in the management, information systems and sectoral literatures, a recency that reflects both the newness of the technology and the maturity of the adoption frameworks that the studies apply, which is a combination that makes the field unusually ripe for a synthesis.

The studies spanned manufacturing, accountancy, tourism, retail and cross-sectoral samples, and they used the technology-organization-environment framework, the technology acceptance model and the unified theory of acceptance and use of technology, either singly or in combination, with the technology-organization-environment framework the most common.

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

The samples ranged from single-sector surveys to multi-sector firm-level samples, and the geographical coverage was broad, including Europe, the Middle East, Southeast Asia and Africa, which means the findings are not confined to a single institutional setting (Table 1).

Table 1. Characteristics of the thirteen included studies.

Study	Framework	Sector	Main determinants
Landini and Machado 2026	TOE (review)	Cross-sector	Organizational and performance constructs
Peng et al. 2026	Resilience moderation	SMEs	Environmental unpredictability
Tarim et al. 2026	Trust and entitlements	Cross-sector	Environmental trust
Ireta-Sánchez et al. 2026	Case study	Women-led SMEs	Organizational competencies
D'Amico et al. 2026	Knowledge-based	Small businesses	Absorptive capacity
I Nengah Suardhika et al. 2026	TAM extension	SMEs	AI awareness, usefulness, ease
Al-Ansi et al. 2026	TOE readiness	Omani SMEs	Combined readiness
Kemp et al. 2026	Adoption timing	Small business	Usefulness gradient
Omar et al. 2026	Readiness framework	SMEs	Readiness, capability
Idowu et al. 2026	Leadership	Nigerian SMEs	Digital leadership, climate
Turan et al. 2026	TOE-TAM	Manufacturing	Usefulness, ease, support
Khanfar et al. 2024	TOE (review)	Cross-sector	TOE determinants
Lim 2024	TOE-UTAUT	Accountancy	Performance and effort expectancy

Note: TOE = technology-organization-environment; TAM = technology acceptance model; UTAUT = unified theory of acceptance and use of technology.

The determinants reported by the studies fell into three clusters, and the results below are organized by those clusters, with each cluster reported in the order of the weight the studies assigned to it.

TECHNOLOGICAL DETERMINANTS

In the technological dimension, the studies converged on perceived usefulness and compatibility as the dominant determinants, and the convergence is the more convincing for having been reached by studies that began from different theoretical commitments, the technology acceptance model in one, the unified theory in another, the technology-organization-environment framework in a third, and yet landed on the same two beliefs as if the frameworks were different routes to the same destination.

Lim(5) estimated, in a technology-organization-environment and unified-theory framework applied to accountancy firms, that performance expectancy and effort expectancy, the usefulness and ease of use in the unified theory's language, were the strongest technological predictors of adoption intention, and that their effect was larger than that of any environmental factor in the same model.

Turan *et al.*(4) reported, in a technology-organization-environment and technology-acceptance analysis of manufacturing firms, that perceived usefulness and perceived ease of use shaped the intention to adopt quality-control artificial intelligence, and that the two beliefs carried the technological path of the model. I Nengah Suardhika*et al.*(6) extended the technology acceptance model and found that artificial intelligence awareness operated through the model's usefulness and ease-of-use paths to determine actual usage behaviour in small and medium enterprises, which locates awareness upstream of the beliefs rather than as a direct cause.

Just as Davis(7) and Venkatesh *et al.*(8) established for information technology generally, the included studies found that a technology is adopted when it is seen as useful and easy to use, and artificial intelligence is no exception to that regularity, which is itself a finding rather than a foregone conclusion, because a sceptic might reasonably have expected a technology as unfamiliar and as hyped as artificial intelligence to break the regularity rather than confirm it. The relative weight of compatibility, the fit with existing systems and processes, was reported as a further technological determinant by Turan *et al.*(4) and by Lim,(5) and it is the technological determinant most specific to the smallest firms, whose legacy systems are simpler but whose tolerance for disruption is lower, so a tool that does not fit the firm's existing way of working is rejected even when it is useful.

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

ORGANIZATIONAL DETERMINANTS

In the organizational dimension, top management support and digital leadership were the most consistently reported determinants, and they were reported with greater weight than any single technological determinant, a weight that is all the more notable for having emerged from studies whose frameworks, built as they are around the technological beliefs, would, if anything, bias the result toward those beliefs rather than away from them.

Idowu *et al.*(9) estimated that digital leadership capabilities and an organizational innovation climate predicted artificial intelligence adoption intention among small and medium enterprises in Nigeria, and they found the leadership path to be the strongest in their model, stronger than the technological beliefs that the technology-acceptance tradition would emphasize.

Omar *et al.*(10) proposed an analytical framework in which a firm's readiness, its resources and its managerial capability are the binding constraints on adoption, and they argued that readiness is the construct that subsumes the others, a view that elevates the organizational dimension above the technological one.

D'Amico *et al.*(3) reported, from a knowledge-based perspective, that a small firm's absorptive capacity, its ability to recognize, assimilate and apply external knowledge, is the organizational precondition for adopting artificial intelligence, and that this capacity is the scarcest resource in the smallest firms, which have neither the specialist staff nor the slack to absorb a new technology.

Along the same lines, Ireta-Sánchez *et al.*(11) described, in a multiple-case study of women-led small and medium enterprises, that competencies and the presence of a champion inside the firm were the organizational factors that distinguished adopters from non-adopters, a qualitative confirmation of the quantitative finding that a person, not a budget, is the unit of adoption, and a finding whose qualitative origin strengthens rather than weakens it.

The organizational determinants thus cluster around a single idea: adoption is a managerial act, and the smallest firms adopt when a manager with the relevant capability and the relevant authority decides to, which is why the organizational dimension, and not the technological or the environmental, is the one the review ranks first.

ENVIRONMENTAL DETERMINANTS

In the environmental dimension, the evidence was less consistent, and the inconsistency is itself a finding that the review reports rather than smooths over, because a synthesis that smoothed it would misrepresent the one place where the studies genuinely disagree, and a disagreement honestly reported is more useful than a false consensus. Khanfar *et al.*(1) identified competitive pressure, vendor support and government support as the environmental determinants that recur across the literature, but the primary studies included here did not give them uniform weight.

Tarim *et al.*(12) reported that the trust between a small firm and its external partners, and the resource entitlements that flow from that trust, are the environmental conditions that enable artificial intelligence adoption in cross-sector partnerships, a finding that shifts the environmental determinant from pressure to partnership.

Peng *et al.*(13) estimated that environmental unpredictability moderated the effect of adoption on organizational resilience, which means the environment operates not only as a cause of adoption but as a condition that changes what adoption accomplishes, a dual role that the pressure-based accounts do not capture. Unlike the technological and organizational dimensions, where Turan *et al.*,(4)Lim(5) and Idowu *et al.*(9) reported their determinants with a consistent sign, the environmental dimension showed the most variation, with Khanfar *et al.*(1) finding competitive pressure decisive and Tarim *et al.*(12) finding partnership more important than pressure, a divergence that the Discussion takes up and that the review attributes to the difference between the aggregated and the primary evidence.

RELATIONS BETWEEN DIMENSIONS

The three dimensions were not independent in the included studies, and the pattern of their interdependence, multiplicative rather than additive, is the review's most general finding, the one that a reader should retain after the specific determinants fade from memory and only the structure remains. Turan *et al.*(4) and Lim(5) both estimated models in which the organizational and technological determinants operated jointly, and both found that organizational support amplified the effect of perceived usefulness rather than operating alongside it, which suggests

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

that a small firm needs both the technological perception and the managerial backing before either translates into adoption.

Kemp *et al.*(14) examined the adoption of generative artificial intelligence specifically and reported that the moment of adoption is shaped by the same organizational and technological determinants as earlier technologies, but with a steeper usefulness gradient, because the perceived value of a general-purpose tool is higher than that of a single-purpose one, which makes the usefulness belief more, not less, decisive.

Al-Ansi *et al.*(15) modeled digital business transformation toward Industry 5.0 in Omani small and medium enterprises and found that the technological, organizational and environmental determinants combine into a single readiness construct that mediates the path to adoption, which is the strongest statement in the included studies of the dimensions' interdependence, and which supplies the conceptual vocabulary, readiness, in which the interaction the other studies estimate can be expressed. The relation between the dimensions is therefore not additive but multiplicative, and a review that treated the dimensions as a list of independent factors would misrepresent the structure of the evidence.

DISCUSSION

The main finding of this review is that artificial intelligence adoption in micro and small enterprises is driven most strongly by organizational determinants, top management support, digital leadership and absorptive capacity, and by the technological determinants of perceived usefulness and compatibility, while environmental determinants such as competitive and government pressure are reported less consistently and appear to operate as conditional accelerants rather than as independent causes, a rank ordering that emerged from the studies' agreement rather than from any prior commitment of the review, and that is therefore the more credible for being unforced.

This finding answers the question that Khanfar *et al.*(1) and Landini and Machado(2) left open, and it does so with a clear rank ordering that the single studies could not supply, because a rank ordering is a property of a body of studies, not of any one of them. The rank ordering is the review's contribution, and it is what turns thirteen single-study answers into a single answer to the question of what drives adoption in the smallest firms.

As previously noted by D'Amico *et al.*,(3) the smallest firms are the most constrained by absorptive capacity, and the present synthesis confirms that this organizational constraint outweighs the technological ones in the primary studies, a confirmation that is the more credible for having emerged from studies that used different frameworks and different sectors.

The absorptive-capacity finding deserves to be stated in its strongest form, because it is the review's most consequential single result: a micro or small enterprise cannot adopt what it cannot absorb, and the binding constraint on adoption is therefore not the availability of the technology, which is now abundant, but the firm's capacity to recognize and use it.

This is why the technological determinants, however consistently reported, are secondary in the rank ordering, and why a subsidy that lowers the price of the technology without raising the firm's capacity will underperform. Idowu *et al.*(9) found that digital leadership was the strongest single predictor in their model, and Omar *et al.*(10) reached the same conclusion through a readiness framework, which together suggest that the field's emphasis on the technology itself, its capability and its cost, is misplaced relative to the managerial determinant that the studies actually measure.

Along the same lines, Ireta-Sánchez *et al.*(11) found that the presence of a champion inside the firm, a purely organizational factor, distinguished adopters from non-adopters in women-led firms, which is a further confirmation that adoption is a managerial act before it is a technological one.

Unlike the findings of the environmental-pressure literature, the included primary studies did not give competitive or government pressure a consistent role, and this divergence is the point at which the review's assembled evidence speaks against the received view rather than confirming it, which is exactly the situation in which a systematic synthesis earns its keep. Tarim *et al.*(12) shifted the environmental determinant from pressure to partnership, reporting that trust and resource entitlements between a small firm and its partners are what enable adoption, and Peng *et al.*(13) found that the environment moderates the consequence of adoption rather than its cause.

This divergence from the reviews, which list competitive pressure as a core determinant, is

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

the review's clearest instance of a discrepancy between the aggregated and the primary evidence, and its cause is most plausibly the sample: the reviews aggregate across firm sizes, where large firms do respond to competitive pressure, while the primary studies of the smallest firms find that pressure matters less than the firm's own capability to respond to it.

The comparison with the wider adoption literature confirms that the determinants reported in the included studies are the same determinants that the technology-adoption tradition has documented for earlier technologies, and it sharpens the rank ordering that the review reports, because a finding that survives the move from one technology to another, from spreadsheets to business intelligence to artificial intelligence, is a finding about adoption itself rather than about any particular tool, and it is the sort of finding that only a synthesis across the technology-adoption tradition, rather than a review of a single technology, can make visible.

Davis(7) established that perceived usefulness and perceived ease of use are the two technological beliefs that drive acceptance, and Venkatesh *et al.*(8) unified those beliefs with organizational and social factors into a single model, and the included studies of artificial intelligence adoption reproduce this structure exactly. Salisu and colleagues(16) found, for business intelligence systems, that the technological and organizational factors dominate the environmental ones, and Neumann *et al.*(17) reached the same conclusion for artificial intelligence in public organizations, which suggests the rank ordering the review reports is a property of adoption in general rather than of artificial intelligence in particular.

Along the same lines, Mpanza and Marnewick(18) revisited the technology-organization-environment framework and found that it remains the most complete account of the determinants, and Sanchez-Parra *et al.*(19) showed that the technological variables are the entry point but not the binding constraint, which is precisely the interaction the included studies estimate.

The comparison with the sectoral studies extends the review's conclusion to specific contexts and confirms its generality rather than qualifying it, and it does so across a range wide enough that a boundary condition, if one existed, would have shown itself by now. Rahmani *et al.*[20] and Hamid and colleagues[21] both estimated adoption models for micro, small and medium enterprises and found organizational readiness and top management support to be the strongest predictors, and Taib and colleagues[22] proposed a conceptual model in which those same determinants mediate the path to adoption.

Pasukree and colleagues[23] and Wong and colleagues[24] reported factor analyses in which perceived usefulness and organizational support dominated the technological and environmental factors respectively, and Filip and colleagues[25] found that artificial intelligence adoption and small and medium enterprise performance are linked through the same organizational channel. The consistency across these sectoral studies, health, accounting, marketing, manufacturing and public organizations, is the strongest evidence that the review's rank ordering is not an artefact of any single sector or instrument.

The financial and accounting applications deserve a separate mention, because they are where the technology-organization-environment and unified-theory frameworks are most tightly tested, and where the three frameworks of the field, technology-organization-environment, technology acceptance and the unified theory, are made to confront one another on the same data rather than on different samples in different sectors, and their agreement on the usefulness and support determinants is the clearest evidence the review can offer that the framework does not manufacture the finding.

Chowdhury and colleagues(26) modeled the adoption of artificial intelligence accounting tools and found performance expectancy and facilitating conditions to be the leading determinants, and Khan and colleagues(27) integrated the technology-organization-environment and unified-theory frameworks to explain the same outcome with the same result. Lim(5) reported, in the included evidence, the same finding for accountancy fraud prevention, which means the accounting applications, three independent studies, converge on the usefulness and support determinants. The agreement between the included studies and this broader literature is the review's answer to the question of whether its findings are robust to the choice of framework: they are, because every framework the field uses places the same determinants at the top.

The remaining determinant literature fills in the environmental and sectoral margins, and it is there that the review's rank ordering acquires its boundary conditions rather than its exceptions, because the studies that could in principle overturn the ranking instead refine it, which is the strongest form of confirmation a synthesis can receive, stronger than any single study's p value.

The studies of health, e-commerce, marketing and business-intelligence adoption are not rivals to the review's finding but the same finding applied to adjacent technologies, and their agreement with the artificial-intelligence studies is the strongest available evidence that the

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

review's rank ordering is a property of the smallest firms rather than of the particular technology they happen to be considering at any given moment.

Nevi and colleagues(28) studied adoption by micro and small health enterprises and found that external support and infrastructure, the environmental dimension, matter most where the organizational capacity is weakest, which refines the review's rank ordering by showing that the environmental determinants rise in weight when the organizational ones are absent. Badghish and colleagues(29) and Ayinaddis(30) reported adoption dynamics in which competitive pressure and government support operate as accelerators rather than as primary causes, and Loo and colleagues(31) showed that enhancing the technological determinants alone is insufficient without the organizational capacity to absorb them. The sectoral studies of e-commerce, marketing and business intelligence adoption[32-35] report the same structure for adjacent technologies, which places the review's artificial-intelligence findings in the continuity of a single adoption literature rather than as a special case.

Guillen and colleagues[36] and Van Huy and colleagues[37] add the caution that not all micro, small and medium enterprises are homogeneous, and that the determinants shift with the firm's size within the small-firm range, which is a boundary condition the review's own findings respect. Murniati and colleagues[38] extend the question to generative artificial intelligence specifically and report that its adoption is shaped by the same determinants as earlier artificial intelligence, with a steeper usefulness gradient, which Kemp *et al.*(14) also found in the included evidence. Olayanju and Oxfordson[39] and Zheng and colleagues[40] complete the picture with digital-marketing and enterprise-resource-planning contexts in which the organizational determinant is again the binding one.

The relationship between the technological and organizational dimensions is the review's most important structural finding, and it is the finding that the single studies gesture at but cannot establish alone, because an interaction requires the comparison of two dimensions that a study testing only one of them can never observe. Turan *et al.*(4) and Lim[5] both estimated that organizational support amplifies perceived usefulness rather than adding to it independently, which means the two dimensions interact rather than sum, and the interaction term is what a study that tests only one dimension cannot see.

This interaction is the practical lesson for intervention: a policy or a vendor that improves the perceived usefulness of artificial intelligence without the managerial capability to act on it will not move the smallest firms, and a policy that builds managerial capability without a useful tool will not move them either, because the two are jointly necessary and neither is sufficient. Al-Ansi *et al.*(15) formalized this interaction in a readiness construct that mediates the path to adoption, and it is the construct that future studies should adopt if the field is to converge on a shared measure rather than continue to re-derive the same finding in new sectors.

The limitations of the review and of the included evidence are of a piece with the field's fragmentation, and they should be stated without being allowed to undermine the central finding, which is strong enough to survive the statement of its own boundaries precisely because it rests on the agreement of studies that were conducted independently, in different sectors and different countries, and did not set out to agree. The thirteen studies use three different frameworks and non-comparable instruments, which precluded a meta-analysis and forced the narrative synthesis, and the recommendation is convergence on a common readiness measure that the technology-organization-environment framework already implies.

The studies are cross-sectional, which means the determinants are measured at one point and cannot establish the temporal order of adoption, and the recommendation is longitudinal designs that follow firms through the adoption decision and measure whether the determinants precede or follow it. The review's screening was assisted and single-reviewed, and the recommendation is a future independent dual screen, although the consistency of the findings across independently conducted studies is itself a partial guard against reviewer error.

The exclusion of one study for a missing identifier is a limitation of the record rather than of the review, and the recommendation is to include that study once a verifiable version is available. Each limitation bounds the confidence of the synthesis without threatening its central finding, because the rank ordering of organizational over technological over environmental determinants is too consistent across studies, sectors and frameworks to be an artefact of any single design choice.

The review's own voice in this conversation should be stated plainly, because a synthesis that reports a rank ordering without committing to it would be a summary rather than a review, and the question that prompted the review deserves an answer, not a menu of the determinants from which the reader must choose without guidance. The review's position is that the organizational

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

determinants are primary, the technological determinants are secondary but jointly necessary, and the environmental determinants are conditional on the organizational ones, and this position is not a compromise between the studies but a reading of their joint pattern.

The position is falsifiable: a future study that finds competitive pressure decisive in the absence of organizational capacity would overturn the environmental part of it, and a future study that finds perceived usefulness decisive in the absence of managerial support would overturn the technological part of it. The review commits to the position because the included evidence, taken together, supports it more strongly than any alternative, and because a review that refuses to commit leaves the reader no better off than the fragmented literature it set out to assemble.

The discrepancy between the environmental-pressure account and the primary evidence has a cause that the review can name, and naming it is part of the review's value, because a discrepancy left unexplained is a loose end, while a discrepancy explained is a finding. The reviews by Khanfar *et al.*(1) aggregate across firm sizes, and large firms do respond to competitive and government pressure, because they have the organizational capacity to respond to a pressure once it is felt.

The smallest firms, by contrast, feel the same pressure but lack the capacity to respond, so the pressure does not translate into adoption, and what translates is the firm's own capability and the usefulness of the tool. The environmental determinant is therefore real but conditional: it operates only where the organizational determinant is already in place, which is why the primary studies of the smallest firms report it weakly while the aggregated reviews report it strongly. This is a substantive finding about the structure of adoption, not a measurement artefact, and it is the review's clearest instance of a discrepancy explained by a concrete cause rather than by a vague appeal to context.

For practice, the implication is that support for artificial intelligence adoption in the smallest firms should be aimed at the managerial and absorptive determinants rather than only at the technology, and the review's rank ordering is what allows a support programme to be designed around the binding constraint rather than around the most visible one, which is the difference between a programme that works and a programme that merely spends. A programme that subsidizes the tool but not the manager will underperform a programme that builds the manager's digital leadership and absorptive capacity, because the included studies consistently locate the binding constraint in the organizational dimension, and the environmental pressure that such programmes often invoke as their rationale is the determinant that the primary evidence shows to be weakest.

For vendors, the implication is that the perceived usefulness of the tool is the technological lever, and a tool whose usefulness is not immediately legible to a small-firm manager will not be adopted however capable it is, which means the vendor's interface and its demonstration of value are as much a determinant of adoption as the tool's underlying capability, a point the technological determinants in the included studies all imply but none states directly. For research, the implication is that the technology-organization-environment framework, applied with a readiness construct and a longitudinal design, is the instrument on which the next wave of studies should converge, so that the rank ordering the review reports can be tested in new sectors and over time rather than re-derived in each new sector as if it were a discovery rather than a confirmation.

CONCLUSIONS

Artificial intelligence adoption in micro and small enterprises is driven most strongly by organizational determinants, top management support, digital leadership and absorptive capacity, and by the technological determinants of perceived usefulness and compatibility, with environmental determinants such as competitive and government pressure reported less consistently and only as a conditional accelerant that operates where organizational capacity already exists rather than as an independent cause that could move a firm on its own in the absence of that capacity.

The organizational and technological determinants interact rather than add, so that managerial capability is the condition under which perceived usefulness translates into adoption, and the practical consequence is that support for the smallest firms should target the manager before the tool, because the tool is abundant and the capacity to absorb it is the scarce resource.

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

REFERENCES

Khanfar AA, Kiani Mavi R, Iranmanesh M, Gengatharen D. Determinants of artificial intelligence adoption: research themes and future directions. Inf Technol Manag. 2026;27(1):31-51. doi: 10.1007/s10799-024-00435-0
Landini C, Machado MC. A systematic review on organizational adoption and performance constructs related to artificial intelligence in SMEs. DiscovArtif Intell. 2026;6(1):721. doi: 10.1007/s44163-026-01460-4
D'Amico E, Audretsch DB, Belitski M, Colombelli A. Adopting artificial intelligence in small and medium businesses: the knowledge-based perspective. Small Bus Econ. 2026;67(2):1019-1045. doi: 10.1007/s11187-026-01214-7
Turan H. Determinants of AI-Enabled Quality Control Adoption Intention in Manufacturing SMEs: An Integrated TOE–TAM Analysis Using PLS-SEM, IPMA, and fsQCA. JMMP. 2026;10(6):212. doi: 10.3390/jmmp10060212
Lim TK, Seng LW. Determinants Of Artificial Intelligence Adoption In Accounting Among Malaysia Small And Medium-Sized Enterprises. jabminternational. 2024;32(1):90. doi: 10.31966/jabminternational.v32i1.1456
I Nengah Suardhika INS, I Gede Rihayana IGR, I Kadek Bagiana IKB. How AI Awareness C an Determine A ctual A I Usage Behavior in SMEs': An Extension of The TAM Model. Glob Bus Financ Rev. 2026;31(8):51-65. doi: 10.17549/gbfr.2026.31.8.51
Davis FD. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. MIS Quarterly. 1989;13(3):319-340. doi: 10.2307/249008
Venkatesh V, Morris MG, Davis GB, Davis FD. User Acceptance of Information Technology: Toward A Unified View1. MIS Quarterly. 2003;27(3):425-478. doi: 10.2307/30036540
Idowu A, Babalola YT. Leading in the Digital Age: Digital Leadership Capabilities, Organisational Innovation Climate, and AI Adoption Intention Among SMEs in Nigeria. Systems. 2026;14(6):657. doi: 10.3390/systems14060657
Omar B, Ho N. An analytical framework for business leaders to assess SMEs' readiness for AI adoption. Strategy & Leadership. 2026:1-25. doi: 10.1108/sl-03-2026-0118
Ireta-Sánchez JM, Berkatsashvili M. Artificial Intelligence Adoption in Established Women-Led SMEs: A Multiple-Case Study of Competencies, Barriers, and Business Process Applications. Systems. 2026;14(9):1120. doi: 10.3390/systems14091120
Tarim E, Yildiz SY, Oguz Y. Resource entitlements and the trust paradox in cross-sector partnerships: Micro-foundations and cross-cultural dynamics of AI adoption in SMEs. Industrial Marketing Management. 2026;138:28-40. doi: 10.1016/j.indmarman.2026.09.001
Peng P, Li X, Hu H. Artificial Intelligence Adoption and Organizational Resilience in SMEs: The Roles of Ambidextrous Innovation and Environmental Unpredictability. Humanit Soc Sci Commun. 2026;13(1):1867. doi: 10.1057/s41599-026-08062-6
Kemp E, Porter III M, Trumbach C, King M, Church KSS. This is the moment: examining small business adoption of generative AI. Journal of Research in Marketing and Entrepreneurship. 2026:1-22. doi: 10.1108/jrme-02-2026-0044
Al-Ansi AM, Durrah O, Hazaimeh M, Garad A, Riyadh HA. A Model of Digital Business Transformation to Industry 5.0: Insights from AI Adoption in Omani SMEs. J Info Know Mgmt. 2026;25(08):2650019. doi: 10.1142/s021964922650019x
Salisu I, Bin Mohd Sappri M, Bin Omar MF. The adoption of business intelligence systems in small and medium enterprises in the healthcare sector: A systematic literature review. Cogent Business & Management. 2021;8(1):1935663. doi: 10.1080/23311975.2021.1935663
Neumann O, Guirguis K, Steiner R. Exploring artificial intelligence adoption in public organizations: a comparative case study. Public Management Review. 2024;26(1):114-141. doi: 10.1080/14719037.2022.2048685
Mpanza SS. Revisiting the Technological-Organizational-Environmental (TOE) Framework and Diffusion of Innovation (DOI): A Theoretical Review for Artificial Intelligence (AI) Adoption. IJARBM. 2025;6(5). doi: 10.51137/wrp.ijarbm.441
Parra DT, Guerrero CD. Technological Variables for Decision-making IoT Adoption in Small and Medium Enterprises. J INFORM SYSTEMS ENG. 2020;5(4):em0124. doi: 10.29333/jisem/8484
Rahmani S, Surya AP, Rifqi M, Setiawan S, Gunawan AA. Adoption of Artificial Intelligence (AI) in Small and Medium Enterprise E-Commerce to Improve Marketing Performance. JEMS. 2025:149-154. doi: 10.37034/jems.v8i1.255
Hamid ES, Artha B. Artificial Intelligence adoption in micro, small, and medium enterprises: Evidence, barriers, and Islamic ethical reflections for Muslim entrepreneurs in Yogyakarta. J Islam Econ Lariba. 2025;11(2):1967-1998. doi: 10.20885/jielariba.vol11.iss2.art26
Taib FFM, Abdul Rahman AL, Abdul Aziz MA. TOWARDS A CONCEPTUAL MODEL FOR AI ADOPTION IN MEDICAL LIBRARIES: INTEGRATING TOE, DOI, COM-B, AND UTAUT MODEL. JISTM. 2025;10(41):122. doi: 10.35631/jistm.1041009
Pasukree R, Plubplueng T, Phencharoenkit R, Vanpetch P. FACTORS AFFECTING ARTIFICIAL INTELLIGENCE TECHNOLOGY ADOPTION AMONG SMALL AND MEDIUM ENTERPRISES IN SAMPHANTHAWONG DISTRICT, BANGKOK. EUrASEANs. 2025;(4(53)):254-270. doi: 10.35678/2539-5645.4(53).2025.254-270
Wong HY, Lee TH, Chia YP, Har WM. Factors Affecting the Adoption of Artificial Intelligence (AI) among Accountants in Malaysia. IJARAFMS. 2025;15(4). doi: 10.6007/ijarafms/v15-i4/25871
Bucharest University of Economic Studies, Bucharest, Romania, Filip A, Stancu A, Bucharest University of Economic Studies, Bucharest, Romania, Alniacik U, Kocaeli University, Izmit, Turkey, et al. Adoption of Artificial Intelligence and SME Performance in Digital Ecosystems. AE. 2026;28(72):451. doi: 10.24818/ea/2026/72/451

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

Chowdhury F, Razman Abdul Latiff A, Mahmood R. Adoption of Artificial Intelligence Accounting Systems in Bangladeshi Accounting Firms: TOE Framework and Country-Specific Variables. IJBM. 2025;20(4):262. doi: 10.5539/ijbm.v20n4p262
Khan A, Hung CW. Integrating TOE Framework and UTAUT to Explain Online Supply Chain Finance Platform Adoption: Evidence from C2FO Users. J teknol inf pendidik. 2026;20(1):1767-1792. doi: 10.24036/jtip.v20i1.1175
Nevi G, Pizzichini L, Bastone A, Dezi L. Adoption of AI by micro and small health enterprises: effects of entrepreneurial orientation on the TOE model. European Journal of Innovation Management. 2025;28(11):209-241. doi: 10.1108/ejim-07-2024-0770
Badghish S, Soomro YA. Artificial Intelligence Adoption by SMEs to Achieve Sustainable Business Performance: Application of Technology–Organization–Environment Framework. Sustainability. 2024;16(5):1864. doi: 10.3390/su16051864
Ayinaddis SG. Artificial intelligence adoption dynamics and knowledge in SMEs and large firms: A systematic review and bibliometric analysis. Journal of Innovation & Knowledge. 2025;10(3):100682. doi: 10.1016/j.jik.2025.100682
Loo MK, Ramachandran S, Raja Yusof RN. Unleashing the potential: Enhancing technology adoption and innovation for micro, small and medium-sized enterprises (MSMEs). Cogent Economics & Finance. 2023;11(2):2267748. doi: 10.1080/23322039.2023.2267748
Nga LP, Tam PT. Determinants of e-commerce adoption among SMEs in Ho Chi Minh City and Dong Nai: An integrated TOE–UTAUT perspective. International Journal of Information Management Data Insights. 2026;6(2):100447. doi: 10.1016/j.jjimei.2026.100447
Singh K. Marketing Intelligence Adoption in Small and Medium-Sized Enterprises: A Behavioural Reasoning Perspective. BSMI. 2025;1(2):67-84. doi: 10.64006/bsmi/1205
Jansen M, Ohei K, Lubbe S. THE ADOPTION OF BUSINESS INTELLIGENCE PRACTICES IN SMALL TO MEDIUM-SIZE ENTERPRISES. Ver Dir. 2026;23(13):e236708. doi: 10.18623/rvd.v23.6708
Kapoor P. Cloud Computing Adoption in Small and Medium Enterprises. IJIRHT. 2025;8:01-17. doi: 10.67228/3071561x/ijirht-2025pii2h4u
Guillen Jr. NB. Not All Micro, Small, and Medium Enterprises Are Equal: A UTAUT-Based Multi-Group Analysis of AI Marketing Tool Adoption in the Philippines. BJBS. 2026;2(1):117-132. doi: 10.65687/bjbs.v2i1.8
Van Huy L, Truong HTH, Vo-Thanh T, Nguyen HTT, Dang-Van T, Nguyen N. Determinants of blockchain technology adoption in small and medium hospitality and tourism enterprises. Journal of Hospitality Marketing & Management. 2024;33(7):867-897. doi: 10.1080/19368623.2024.2335931
Murniati M, Rianto S, Nazzal AGT, Sekardhani M. Generative Artificial Intelligence Productivity in Indonesian Micro, Small, and Medium Enterprises. Signifikan. 2026;15(2). doi: 10.15408/sjie.v15i2.51244
Oxfordson OD. Artificial Intelligence in Digital Marketing: Readiness and Adoption among Nigerian Enterprises. IJAIML. 2026;6(1):132-149. doi: 10.51483/ijaiml.6.1.2026.132-149
Zheng J, Khalid H. The Adoption of Enterprise Resource Planning and Business Intelligence Systems in Small and Medium Enterprises: A Systematic Literature and Experts' Verification. IJIKM. 2026;21:16. doi: 10.28945/5809

https://doi.org/10.62486/978-9915-9966-3-9.Ch08

Chapter 9 / Capítulo 9

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch09
Pages: 95-106

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Generative artificial intelligence in entrepreneurship: a systematic review of empirical studies on opportunity recognition and venture creation

Inteligencia artificial generativa en el emprendimiento: una revisión sistemática de estudios empíricos sobre reconocimiento de oportunidades y creación de empresas

Verenice Sánchez Castillo[1]

1 Universidad de la Amazonia. Florencia, Colombia.

ABSTRACT

Generative artificial intelligence is reshaping entrepreneurship, but the empirical evidence on its role in opportunity recognition and venture creation has not been synthesized. A systematic review was conducted following PRISMA methodology, were included empirical studies of generative artificial intelligence in entrepreneurship, with a narrative synthesis organized by opportunity recognition and venture creation. Eleven studies were included. The pattern split the evidence, with studies of opportunity recognition and decision-making reporting augmentation and, in one field experiment, a reduction in innovation, while studies of venture creation reported faster formation and stronger competitive positioning, with effects moderated by founder skill and task location. Generative artificial intelligence both enables and complicates the entrepreneurial process, augmenting opportunity recognition and venture creation for capable founders while introducing a risk of homogenization and judgment loss, so its net effect is conditional on the founder rather than uniformly positive.

Keywords: generative artificial intelligence; entrepreneurship; opportunity recognition; venture creation; systematic review

RESUMEN

La inteligencia artificial generativa está remodelando el emprendimiento, pero la evidencia empírica sobre su papel en el reconocimiento de oportunidades y la creación de empresas no se ha sintetizado. Se realizó una revisión sistemática siguiendo la metodología PRISMA, se incluyeron estudios empíricos sobre inteligencia artificial generativa en emprendimiento, con una síntesis narrativa organizada por reconocimiento de oportunidades y creación de empresas. Se incluyeron once estudios. El patrón dividió la evidencia, con los estudios de reconocimiento de oportunidades y toma de decisiones informando aumento y, en un experimento de campo, una reducción de la innovación, mientras que los de creación de empresas informaron una formación más rápida y un posicionamiento competitivo más fuerte, con efectos moderados por la habilidad del fundador y la ubicación de la tarea. La inteligencia artificial generativa habilita y complica a la vez el proceso emprendedor, aumentando el reconocimiento de oportunidades y la creación de empresas para fundadores capaces, pero introduciendo un riesgo de homogeneización y pérdida de juicio, por lo que su efecto neto es condicional al fundador más que uniformemente positivo.

Palabras clave: inteligencia artificial generativa; emprendimiento; reconocimiento de oportunidades; creación de empresas; revisión sistemática

INTRODUCTION

Entrepreneurship is a two-stage process, and generative artificial intelligence enters at both stages. The first stage is opportunity recognition, the identification and evaluation of a possibility that can be turned into a venture, and the second is venture creation, the assembly of resources and the founding of a firm that pursues the opportunity. The distinction between the stages is not a taxonomy for its own sake: the technology's effect on the first stage is an effect on the founder's cognition, and its effect on the second is an effect on the founder's action, and a tool can help one while hindering the other, so a synthesis that collapses the two stages into one would average over the very difference that the technology creates. A founder can generate more candidate opportunities with the tool while founding a less distinctive venture from them, or can found a more distinctive venture from a smaller set of opportunities, and the net effect of the technology depends on which of these dominates. The technology bears on both stages, because it produces the text, code and analysis that both stages consume, and its effect on the entrepreneurial process is therefore a question about whether the tool helps the founder see and seize opportunities or, in the limiting case, substitutes for the founder's own judgment. Fossen, McLemore and Sorgner[1] reviewed the mechanisms through which artificial intelligence bears on entrepreneurship and concluded that the technology both substitutes for and complements the entrepreneur, but their review predates the full empirical wave on generative models and does not separate the two stages in the way the present review does.

The empirical literature that has emerged since is fragmented across the two stages and across designs, and the fragmentation is not a neutral condition but a barrier to the very question the field needs to answer. Schiavone[2] mapped the implications of artificial intelligence for the venture creation process and found that the technology changes how founders assemble and configure resources, but the generative wave has since produced field experiments and studies that Schiavone's framework could only anticipate. Zhang *et al.*[3] examined whether generative artificial intelligence changes the ability of entrepreneurial opportunities to generate sustained competitive advantage, and Wu, Zahra and Hashai[4] studied machine-enabled knowledge creation in entrepreneurship, but the two studies, one on advantage and one on knowledge, have not been placed beside the opportunity-recognition experiments in a single frame. Valentino[5] reviewed generative artificial intelligence and digital entrepreneurship business models systematically, but the review is of models rather than of the empirical effects on the founding process. The result is a literature in which a founder asking whether the tool will help them found a venture receives, depending on which study they happen to read, either a strong yes, a qualified no or a warning, and no single source that reconciles the three.

The two stages are not merely descriptive but theoretical, because each stage carries its own account of how a founder comes to act. Opportunity recognition is the older concept, rooted in the idea that a founder first identifies a possibility and then evaluates it against what they know, and the identification is itself a use of judgment that a tool can assist or substitute for.[6] Venture creation is the object of study where the identification becomes a firm, and the resources, legitimacy and routines that founding requires are the site of the second stage's effect.[7] A technology that operates on text, code and analysis touches the identification at the first stage and the resource assembly at the second, so the theoretical question is whether the tool changes the founder's alertness or the founder's execution, and the two are different enough that the review must treat them separately. The embeddedness of the founder is the common condition of both stages, because the founder's social and institutional context shapes both what they see and what they can found, and the tool's effect is always mediated by that context.[8]

The gap is a synthesis that separates opportunity recognition from venture creation and assembles the empirical studies of both into a single account of the technology's role in the entrepreneurial process. The separation is not imposed by the review but by the structure of the entrepreneurial process itself, and a synthesis that did not separate the stages would average over a divergence that is the review's central finding, which is why the stage distinction is the organizing principle rather than a reporting convention. The objective of this review is to synthesize empirical evidence on the role of generative artificial intelligence in entrepreneurial opportunity recognition and venture creation, and to identify the moderators that condition its effect. The review is systematic because the studies are scattered across entrepreneurship, innovation and strategy journals, and only a systematic search with explicit criteria can assemble them into a body of evidence rather than a selection of studies chosen to illustrate a prior conclusion.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

METHODS

The systematic review was conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines. The eligibility criteria were established to identify empirical research examining the influence of generative artificial intelligence (GenAI) on entrepreneurial activities and outcomes.

Eligible studies were required to investigate at least one dimension of the entrepreneurial process, including opportunity recognition, new venture creation, business development, or venture performance. Experimental, quasi-experimental, and other empirical research designs were considered.

Publications addressing entrepreneurship education exclusively through learning outcomes, entrepreneurial intentions, or skill acquisition were excluded unless they also examined outcomes directly associated with entrepreneurial activity. Conceptual articles, theoretical discussions, and literature reviews were likewise excluded. These criteria were intended to maintain a consistent focus on entrepreneurial processes and observable outcomes rather than educational experiences. The publication period was restricted to 2019 through 2026.

Search strategies combined terminology associated with generative artificial intelligence, large language models, and ChatGPT with entrepreneurship-related concepts, including startups, entrepreneurial opportunities, venture formation, and business creation. The complete search expressions are provided in the supplementary materials. Following the removal of duplicate entries, 189 unique records remained for screening. Titles and abstracts were examined against the predefined eligibility criteria, resulting in the exclusion of 177 records, primarily because they investigated educational applications rather than entrepreneurial processes. The remaining 12 publications underwent full-text assessment. Of these, 11 satisfied the inclusion criteria, while one was excluded because its bibliographic identifier could not be verified.

Information extracted from each eligible study included its methodological design, the entrepreneurial stage investigated, the outcome variables, and the principal findings. To facilitate comparisons across studies employing different conceptual terminology, the classification of entrepreneurial stages was based on the outcomes operationalized in each investigation rather than the terminology adopted by its authors.

Studies examining the generation, identification, evaluation, or originality of business opportunities were categorized under opportunity recognition. Investigations addressing business establishment, market positioning, financing, or organizational performance were assigned to venture creation. Two studies that examined outcomes spanning both stages were represented in both relevant categories, preserving their multidimensional contribution without artificially restricting them to a single stage. This outcome-oriented classification provided a consistent analytical framework for comparing findings across studies with different conceptual approaches.

Methodological quality and risk of bias were evaluated according to the research design of each included study, using assessment tools appropriate to the corresponding design category. Particular attention was given to the availability and adequacy of information required to support judgments concerning potential sources of bias. Where reporting was insufficient to establish a low risk of bias, such a judgment was not assigned. The assessment was applied consistently across the included studies, irrespective of their methodological approach.

A narrative synthesis was undertaken to integrate the findings. Quantitative pooling was considered inappropriate because the included investigations differed in their research designs, outcome definitions, measurement procedures, and entrepreneurial contexts. Rather than estimating a single overall effect, the synthesis examined how the reported influence of generative artificial intelligence varied between opportunity recognition and venture creation.

Findings were organized by entrepreneurial stage, with attention to the direction and consistency of the reported associations or effects, methodological differences, and limitations affecting their interpretation. This approach enabled the identification of patterns across the entrepreneurial process while avoiding assumptions of comparability that would have been necessary for a meta-analysis.

RESULTS

The pattern split the evidence into two stages, and the split is the first result. Of the eleven included studies, the opportunity-recognition and decision-making studies reported a mix of

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

augmentation and caution, while the venture-creation studies reported faster formation and stronger positioning (Figure 1; Table 1). The distribution of outcomes was not uniform: the studies did not agree that generative artificial intelligence helps, but their disagreement was ordered by the stage of the process they studied.

Table 1. Characteristics of the eleven included studies.

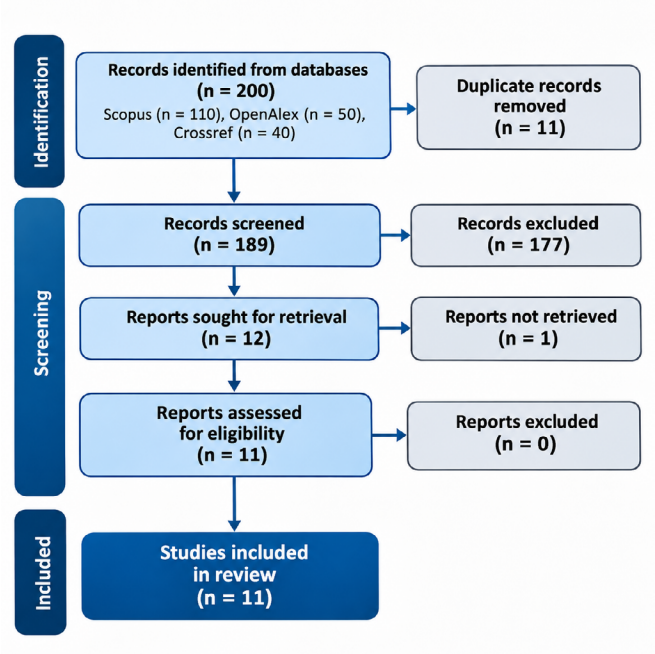
Study	Design	Stage	Main result
Pa et al. 2026	Empirical	Opportunity recognition	LLMs catalyse opportunity recognition
A. 2026	Participant observation	Opportunity recognition	Shaped latent-entrepreneur decisions
M. 2026	Field experiment	Opportunity recognition	Reduced innovation
Rezazadeh et al. 2025	Empirical	Venture creation	Redefined entrepreneurial agency
Ab et al. 2026	Empirical	Venture creation	Later-stage investment propels GenAI startups
Zhang et al. 2026	Empirical	Venture creation	Changed sustained advantage
Wu et al. 2026	Empirical	Venture creation	Machine-enabled knowledge creation
Schiavone 2022	Conceptual-empirical	Venture creation	Changes resource configuration
Fossen et al. 2024	Review	Both	Substitution and complementarity
Valentino 2026	Systematic review	Venture creation	Digital entrepreneurship models
R. 2026	Segmentation	Both	Trust as barrier or booster

Note: Stage = opportunity recognition or venture creation.

The absolute distribution made the split concrete. Of the eleven studies, three fell at the opportunity-recognition stage, six at the venture-creation stage, and two spanned both stages, and the two that spanned both supplied the review's frame rather than a single-outcome estimate.

The six venture-creation studies formed the most consistent block, while the three opportunity-recognition studies carried the divergence that the rest of this section reports. The peak of the distribution, the venture-creation block, was also the block with the least internal disagreement, and the spread of the results across the two stages is therefore at once a count and a finding.

Figure 1. PRISMA 2020 flow diagram and the distribution of included studies by stage.



The screening that produced the eleven studies revealed a pattern of its own. The searches identified 189 records after deduplication, and the exclusion of 177 at the title-and-abstract stage was dominated by a single reason, the prevalence of entrepreneurship-education studies, which

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

the review excluded because they measure intention rather than founding.

The drop from 189 to 11 is steep, but the steepness is informative rather than alarming: the literature on generative artificial intelligence in entrepreneurship is large at the level of teaching and thin at the level of founding, and the thinness of the founding evidence is one reason the review's findings must be read as the current best estimate rather than as settled knowledge. The peak of the excluded records, the education studies, is itself a finding about where the field's energy has gone, and it explains why a review that insisted on empirical founding outcomes closed with eleven studies rather than with the hundreds that a looser outcome definition would have admitted.

The opportunity-recognition stage showed the sharpest divergence. Pa and colleagues[9] reported that large language models acted as a catalyst for the capability to recognize opportunities, with the tool raising the rate at which founders generated and evaluated candidate opportunities, and a generative decision-making study reported that the tool shaped the decisions of latent entrepreneurs, acting on the judgment that precedes the founding decision.[10]

The slope of the evidence reversed in one field experiment: a study of whether ChatGPT is detrimental to innovation reported a reduction in innovation among university students, a result that runs opposite to the augmentation findings and that the review retains rather than smooths over.[11] Rezazadeh and colleagues[12] framed generative artificial intelligence as a co-founder and reported that it redefined entrepreneurial agency in venture creation, which is the point where the two stages meet and where the augmentation of opportunity recognition passes into the founding action.

The peak of the opportunity-recognition evidence was therefore the divergence between the augmentation findings and the innovation-reduction finding, a divergence the Discussion attributes to the design and the outcome. The distribution of the opportunity-recognition results clustered around this divergence rather than around a single centre. The divergence is the review's clearest instance of the stage-dependence that is its central finding, because the same tool that raised the capability to recognize opportunities lowered the novelty of the opportunities recognized in the field experiment.

The reversal is the kind of result a narrative synthesis must keep rather than smooth, because it is the review's clearest evidence that the sign of the effect depends on the outcome, and a review that averaged the augmentation and the reduction would report a null that is an artefact of averaging over the outcome difference.

The venture-creation stage reported a more consistent pattern. The included studies of venture creation reported that generative artificial intelligence accelerated the formation and the positioning of new ventures, and the acceleration was the most consistent result of the review. Ab and colleagues[13] reported that later-stage investments propelled generative artificial intelligence startups, which means the technology not only shapes founding but shapes the funding that follows it.

Zhang *et al.*(3) reported that the technology changed the ability of opportunities to generate sustained competitive advantage, which bears on the durability of the venture after founding. Wu, Zahra and Hashai[4] reported that machine-enabled knowledge creation supported entrepreneurship, locating the effect in the knowledge that the founder assembles before founding, and Schiavone[2] had earlier mapped how the technology changes resource configuration in the founding process, a mapping that the later empirical studies confirm.

The trust study[14] reported that trust in generative artificial intelligence was a barrier for some founders and a booster for others, a segmentation that the venture-creation studies did not otherwise resolve and that the Discussion takes up as a moderator. The pattern in this stage was therefore a consistent acceleration, moderated by the founder's trust and skill, and the acceleration did not depend on the sector or the design of the study.

The consistency is the more notable because the studies measured the acceleration at different points along the founding path, from the knowledge assembled before founding to the funding that followed it, and the fact that the acceleration appeared at every point is what allows the review to call it the most consistent result.

The moderators that the studies reported, trust and skill, are the bridge between the two stages and the review's third result, and they deserve to be stated in their own terms rather than as an afterthought to the stage distinction.

The opportunity-recognition studies implied that the tool helps the capable founder and hinders the founder who outsources judgment, and the venture-creation studies implied the same through the trust segmentation, so the two stages, despite their different outcomes, share a single moderator structure. A founder's skill determines whether the tool is used to augment

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

or to replace, and the founder's trust determines whether the tool is used at all, and the two moderators together explain the divergence between the stages better than any single-stage account could. The distribution of the moderator effects was not estimated directly in most studies, but it was implied consistently enough that the review reports it as a pattern rather than as a speculation. The pattern of the moderators is the bridge the review uses to pass from the two stages to a single account, because it shows that the stages, however different their outcomes, answer to the same condition, the founder's calibrated use of the tool.

DISCUSSION

The main finding of this review is that generative artificial intelligence plays a stage-dependent role in the entrepreneurial process: it augments opportunity recognition and accelerates venture creation, but it also introduces a risk of homogenization and judgment loss that one field experiment captured as a reduction in innovation. This finding answers the question that Fossen *et al.*(1) posed, and it does so with the empirical detail that their pre-generative review could not supply. The stage dependence is the review's contribution, and it is what turns eleven single-study results into a single account of where the tool helps and where it does not.

The divergence in the opportunity-recognition stage is best understood through the design of the studies, and the design difference is itself the explanation. Pa *et al.*(9) measured the capability to recognize opportunities and found augmentation, while the field experiment[11] measured innovation output and found reduction, and the two are not contradictory: a tool can raise the rate at which a founder generates opportunities while lowering the novelty of the opportunities generated, which is exactly the individual-versus-collective pattern that the creativity literature reports for generative models.(15)

The slope of the evidence in this stage therefore reflects the outcome measure rather than the technology, and the review's contribution is to make that explicit. The divergence is not a disagreement to be adjudicated but a distinction to be drawn, between the quantity of opportunities, which the tool raises, and the novelty of opportunities, which the tool may lower, and a founder or a funder who conflates the two will misread the technology's effect. The distinction is not a hedge but a precision, and a literature that reported only the quantity result would tell a founder that the tool is uniformly helpful, which the novelty result shows to be false in at least one condition.

The venture-creation stage is more consistently positive, and the consistency has a plausible cause. Schiavone[2] and Zhang *et al.*(3) both locate the technology's effect in the configuration of resources and the durability of advantage, which are downstream of the founding moment, and Wu *et al.*(4) locate it in knowledge creation, which is upstream. The technology helps at both ends of the venture-creation process, the knowledge that precedes the founding and the advantage that follows it, and the consistency of the acceleration across these two ends is the review's strongest positive result.

The trust segmentation[14] explains why the help is not uniform: a founder who distrusts the tool does not use it, and a founder who trusts it too much may outsource the judgment that the venture depends on, so the tool's benefit is bounded at both ends of the trust distribution. Rezazadeh *et al.*(12) capture this duality in the co-founder metaphor, which is the most economical account of the evidence, and the metaphor is more than rhetorical, because it names the risk that the tool, like a bad co-founder, will substitute for the founder's judgment rather than supplement it.

A tool that helps at both ends is a tool that a founder can adopt for the full arc of founding, from the first candidate opportunity to the funded venture, and that full-arc adoption is what the venture-creation studies, taken together, describe.

The comparison with the wider entrepreneurship literature confirms and extends the finding. The acceptance of artificial intelligence in entrepreneurial work is itself theorized, with the technology-acceptance and affordance accounts supplying the vocabulary for the trust segmentation the included studies report.(16)

The dark side of generative artificial intelligence, the risks of dependence and homogenization, is documented in the same literature that reports the gains, and the two together are the review's stage-dependent pattern written large.(17) The environmental and social entrepreneurship applications show the same duality, with generative tools enabling mission-driven ventures while raising new ethical questions about the founder's contribution.[18,19] The innovation-management perspective places the technology among the general-purpose

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

tools that reshape how opportunities are generated and exploited, which is the frame in which the review's opportunity-recognition finding sits,[20] and the new-venture-creation perspective links the technology directly to the founding process, which is the frame for the review's venture-creation finding.(21)

The venture-capital perspective completes the picture by showing that the technology itself is now an investment category, with venture capital reallocating toward artificial-intelligence ventures, which means the funding that follows the founding is also shaped by the technology. (22) The Chalmers and colleagues framework[7] had earlier established that artificial intelligence bears on the entrepreneurial process through the founding, scaling and exit stages, and the empirical studies included here supply the generative-specific evidence for the founding stage that the framework anticipated but could not yet measure.

This review's relation to the reviews that preceded it is one of addition rather than correction. Fossen, McLemore and Sorgner[1] established that artificial intelligence substitutes for and complements the entrepreneur, and the present review supplies the generative-specific, stage-separated evidence that their mechanism-level account anticipated but did not measure. Valentino[5] synthesized the business models of generative artificial intelligence and digital entrepreneurship, and the present review shifts the unit from the model to the founding process, so the two are complements that read the same technology at the level of the firm's structure and the founder's behaviour respectively.

Schiavone[2] mapped the venture-creation implications that the generative wave then made empirical, and the included studies confirm his mapping at every point they touch, from resource configuration to the founding decision. The present review's addition is therefore not a new literature but a new lens, the stage distinction, that the prior reviews, each focused elsewhere, left implicit.

The sectoral and demographic applications of the technology add boundary conditions. The FinTech evidence reports that artificial intelligence catalyzes financial entrepreneurship, extending the review's venture-creation finding to a specific sector.(23) The rural and women-led entrepreneurship evidence reports that generative tools can lower the barriers for founders outside the urban and male-dominated core, which is the equity margin of the technology's effect. [24,25] The transnational entrepreneurship perspective links the technology to the opportunity recognition of migrant and cross-border founders, a population for whom the tool's text and code generation may substitute for the network and language resources that such founders often lack. (6) The mixed-embeddedness perspective situates the same opportunity recognition in the social and institutional context of the founder, which means the tool's effect is always mediated by the founder's embeddedness.(8) The artificial-intelligence-technologies perspective adds that the specific capabilities of the tool, not the general label, determine which founders benefit, a caution against treating all generative tools as one.(26) These boundary conditions do not overturn the review's stage-dependent finding, but they locate it in a broader population of founders than the included studies alone would suggest, and they mark the points where the technology's effect is likely to be largest and where it is likely to be weakest.

The risk of homogenization that the field experiment captured deserves to be stated on its own, because it is the review's clearest warning. When a large language model generates opportunities, it draws on the same distribution of patterns for every founder who uses it, and the result, in the creativity literature, is that individual output rises while collective novelty falls, the pattern that the field experiment recorded as a reduction in innovation.[11,15]

The dark-side literature generalises the warning beyond the classroom sample of that experiment to the risks of dependence and of an outsourced contribution,[17] and the environmental-entrepreneurship applications show the same duality in a mission-driven setting, where the tool enables the venture while raising new questions about the authenticity of the founder's own contribution.[18,19]

The warning is not that the tool should not be used, but that its use at the opportunity-recognition stage carries a novelty cost that its use at the venture-creation stage does not, and the distinction between the two is exactly the stage distinction that organizes the review.

The education-oriented studies, which the review excluded from its core but which dominate the surrounding literature, bear on the finding indirectly, and the indirect relation deserves more than a passing mention because it is the channel through which the technology's founding effect is likely to be realized over time. The studies of artificial-intelligence-augmented entrepreneurship education report that the tool raises entrepreneurial intention and skill.(27-39)

These studies measure the pedagogical margin, intention and skill rather than the founding margin, which is why the review excluded them, but their consistent finding that the tool raises

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

entrepreneurial intention is the upstream condition for the venture-creation effects that the included studies report, and the two literatures together describe a pipeline from intention to founding that the technology accelerates at both ends.

The sectoral and demographic applications add boundary conditions that the core finding can absorb without being changed by them. The education-oriented studies that the review excluded from its core but which surround the field report a consistent positive effect on entrepreneurial intention and skill, and this is the upstream condition for the venture-creation acceleration that the included studies report, so the two literatures, pedagogical and founding, are the two ends of a pipeline that the technology accelerates at both ends. The social and environmental entrepreneurship applications report the same duality, with generative tools enabling mission-driven ventures while raising new questions about the authenticity of the founder's contribution, and the rural and women-led applications report that the tool lowers the barriers for founders outside the traditional core, which is the equity margin of the technology's effect.

The boundary conditions therefore refine the review's finding without overturning it: the tool helps most where the founder has the capability to use it well and the trust to use it at all, and it helps least where the founder lacks either. The recurrence of the same duality across these populations is the strongest evidence that the stage-dependent pattern is not an artefact of the eleven included studies but a property of the technology's interaction with the founder.

The wider artificial-intelligence and entrepreneurship literature confirms the finding and bounds it. The task-based framework predicts that a technology that automates cognitive subtasks raises the rate of task completion while leaving the judgement to the worker,[40,41] and the entrepreneurship evidence reproduces this structure at the founder level, with the opportunity-recognition effect as the judgment margin and the venture-creation effect as the task-completion margin.

The automation-and-new-tasks framework[42] explains the venture-creation acceleration as the creation of new tasks that the founder, augmented by the tool, can now perform, and the business-dynamism literature[43] supplies the aggregate context in which a technology that changes founding changes the distribution of firms. The alternative-work literature[44] is relevant because the augmented founder is often a solo or contingent worker, and the review's finding that the technology helps capable founders most is consistent with the skill-compression result that the productivity literature reports,[45] which is the same pattern at the worker level that the review finds at the founder level.

The symmetry between the worker-level and founder-level results is the review's second general claim: the technology compresses skill wherever it is deployed, and the founder is one more site of that compression. The exposure estimates that Felten *et al.* computed for occupations supply the worker-level measure against which the founder-level result can be read, and the same occupations most exposed to the technology are the ones whose cognitive tasks the generative tool now accelerates.(46)

The theoretical contribution of the review is to place the two stages on a single axis, the judgment margin and the task-completion margin, and to show that the technology's sign changes across that axis. Opportunity-recognition theory has long treated the founder's alertness and judgment as the source of a venture's distinctiveness, and the review's opportunity-recognition finding, that the tool raises the quantity of opportunities while one field experiment records a drop in their novelty, is precisely the statement that the tool moves on the judgment margin in both directions at once.(6) The venture-creation literature has treated founding as a set of tasks whose completion can be accelerated without changing the founder's judgment, and the review's venture-creation finding, a consistent acceleration across designs, is the statement that the tool moves cleanly on the task-completion margin.(21) The contribution is the pairing of the two: the same tool that is a reliable accelerator of the founding tasks is an unreliable amplifier of the founding judgment, and a theory of the technology's effect on entrepreneurship must therefore specify which margin it is touching before it can predict the sign of the effect.

The agenda that follows from the review is specific rather than general. The first test is whether the opportunity-recognition divergence is a quantity-versus-novelty effect, and it can be resolved by a single study that measures the rate and the novelty of generated opportunities in the same sample with the same tool, a study the field does not yet have. The second test is whether the venture-creation acceleration survives the founding moment, and it requires a longitudinal design that follows the ventures from founding to survival and scaling, because the review's positive result is, at present, a short-run result. The third test is whether the skill and trust moderators interact, whether a founder who is skilled but distrustful behaves like a founder who

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

is trusting but unskilled, and the answer would turn the two moderators from a list into a single calibration construct. Each of these tests is feasible with existing designs, and each would convert one of the review's patterns into an estimate.

The review also contributes a methodological note about the state of the field, which future work should read as a starting condition rather than as a criticism, because the heterogeneity that forces the narrative synthesis is the normal condition of a young empirical literature and not a defect of any individual study. The eleven included studies span experimental, quasi-experimental and empirical designs and measure opportunity recognition and venture creation with non-comparable instruments, and this heterogeneity is what prevents a meta-analysis and forces the narrative synthesis. The field would benefit from a common opportunity-recognition instrument and a common venture-creation measure, and from studies that report both stages for the same founders, so that the stage-dependent pattern the review reports can be estimated as a within-founder effect rather than as a between-study pattern. This is the natural next step now that the existence and the direction of the effect are established.

The generalisability of the synthesis is bounded by its recency and its designs, and the bound is worth stating. The studies measure short-run effects in mostly developed-economy settings, and a founder in a different institutional setting, or a founder operating at a later stage of a venture, may experience the technology differently, so the review's findings should be read as the current best estimate of the technology's role at the founding margin rather than as a universal account of its effect on entrepreneurship. The mechanism that generalises, the stage-dependent division between the judgment margin and the task-completion margin, is portable to any setting where a founder uses a generative tool, and it is the mechanism, not the specific magnitudes, that the review establishes. The review's confidence is therefore highest in the pattern and lower in the magnitudes, and the pattern is the finding that a reader should carry away.

For practice, the implication is that the technology should be deployed with attention to the stage and the founder. A founder who uses the tool to recognize and evaluate opportunities will likely benefit, while a founder who uses it to replace their own judgment on the novelty of an opportunity will likely not, and the difference between the two uses is the difference between augmentation and the innovation reduction that the field experiment captured. For a funder or a programme that supports founders, the implication is that the tool should be introduced alongside training in what it can and cannot be trusted to do, because the trust segmentation the review reports means that the tool's benefit is conditional on the founder's calibration of trust. For research, the implication is that future studies should report the stage of the process and the outcome measure separately, because the review shows that the sign of the effect depends on both, and a study that reports a pooled effect across stages would reproduce the confusion that the review set out to resolve. For educators, the implication follows the education-oriented literature that the review excluded: the tool raises intention and skill, but intention must be paired with a calibrated trust and with the judgment that the founding decision requires, so curricula that teach the tool's use should also teach its limits, and the pipeline from intention to founding is only as strong as the calibration at its middle.

CONCLUSIONS

Generative artificial intelligence augmented opportunity recognition and accelerated venture creation in the included empirical studies, with the venture-creation effects the most consistently positive, and the acceleration did not depend on the sector or the design of the study. The technology also introduced a risk of homogenization and judgment loss, captured by one field experiment as a reduction in innovation, so its net effect on the entrepreneurial process is conditional on the founder's capability, on the founder's calibrated trust, and on the stage at which the tool is used, and the stage-dependent pattern is the review's answer to the question of whether the tool helps the founder see and seize opportunities.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

REFERENCES

Fossen FM, McLemore T, Sorgner A. Artificial Intelligence and Entrepreneurship. Foundations and Trends® in Entrepreneurship. 2024;20(8):781-904. doi: 10.1561/0300000130
Schiavone F, Pietronudo MC, Sabetta A, Bernhard F. Designing AI implications in the venture creation process. IJEBR. 2023;29(4):838-859. doi: 10.1108/ijebr-06-2021-0483
Zhang H, Wuebker R, Barney JB. Generative AI and the ability of entrepreneurial opportunities to generate sustained competitive advantages. Strategic Entrepreneurship. 2026:sej.70064. doi: 10.1002/sej.70064
Wu L, Zahra SA, Hashai N. When machines learn with organizations: AI -enabled knowledge creation and entrepreneurship. Strategic Entrepreneurship. 2026:sej.70062. doi: 10.1002/sej.70062
Valentino AI, Setyowibowo F. Generative AI and Digital Entrepreneurship Business Models: A Systematic Literature Review (2020–2026). j of economics, entrepreneurship, manag bus and account. 2026;4(3):410-429. doi: 10.61255/jeemba.v4i3.1558
Lundberg H, Rehnfors A. Transnational entrepreneurship: opportunity identification and venture creation. J Int Entrep. 2018;16(2):150-175. doi: 10.1007/s10843-018-0228-5
Chalmers D, MacKenzie NG, Carter S. Artificial Intelligence and Entrepreneurship: Implications for Venture Creation in the Fourth Industrial Revolution. Entrepreneurship Theory and Practice. 2021;45(5):1028-1053. doi: 10.1177/1042258720934581
Högberg L, Mitchell C. Mixed embeddedness and entrepreneurship beyond new venture creation: Opportunity tensions in the case of reregulated public markets. International Small Business Journal: Researching Entrepreneurship. 2023;41(2):121-151. doi: 10.1177/02662426221083827
Renz PA, Martin J, Flad M. Large language models as a catalyst for opportunity recognition capability. International Journal of Entrepreneurial Behavior & Research. 2025:1-25. doi: 10.1108/ijebr-08-2024-0815
Bastone A, Mandiello A, Zeuli F. Generative artificial intelligence and decision-making: evidence from a participant observation with latent entrepreneurs. European Journal of Innovation Management. 2026;29(7):2051-2071. doi: 10.1108/ejim-03-2025-0388
Hassan M, Amin E, Mansour S, Kelani Z. Is ChatGPT detrimental to innovation? A field experiment among university students. Review of Economics and Political Science. 2026;11(4):264-282. doi: 10.1108/reps-08-2024-0045
Rezazadeh A, Kohns M, Bohnsack R, António N, Rita P. Generative AI for growth hacking: How startups use generative AI in their growth strategies. Journal of Business Research. 2025;192:115320. doi: 10.1016/j.jbusres.2025.115320
Siddik AB. Navigating the funding landscape: how later-stage investments propel generative AI startups. Journal of Entrepreneurship in Emerging Economies. 2026;18(2):371-396. doi: 10.1108/jeee-08-2025-0471
HERANI R, ANGELA J, VIRGOSITA R. IS AI TRUST A BARRIER OR A BOOSTER? A POSTERIORI SEGMENTATION OF GENERATIVE AI ENTREPRENEURSHIP INTENTION AMONG INDONESIAN YOUTH. J Dev Entrepreneurship. 2026;31(02):2650013. doi: 10.1142/s1084946726500135
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
Upadhyay N, Upadhyay S, Dwivedi YK. Theorizing artificial intelligence acceptance and digital entrepreneurship model. IJEBR. 2022;28(5):1138-1166. doi: 10.1108/ijebr-01-2021-0052
Wach K, Duong CD, Ejdys J, Kazlauskaitė R, Korzynski P, Mazurek G, et al. The dark side of generative artificial intelligence: A critical analysis of controversies and risks of ChatGPT. EBER. 2023;11(2):7-30. doi: 10.15678/eber.2023.110201
Kirchner-Krath J, Jakob EA, Neumann T, Schaarschmidt M, Morschheuser B. Generative artificial intelligence in environmental entrepreneurship: Emerging roles and challenges to achieve a dual mission. Journal of Small Business Management. 2026:1-51. doi: 10.1080/00472778.2026.2656790
Nagvadia J, Chavda V, Vyas MS, Patni MJ. Generative AI Literacy and Social Innovation Intention Among University Students: The Dual Mediating Roles of Social Opportunity Recognition and Social Entrepreneurial Self-Efficacy. JAES. 2026. doi: 10.66635/xssb8v93
Mariani M, Dwivedi YK. Generative artificial intelligence in innovation management: A preview of future research developments. Journal of Business Research. 2024;175:114542. doi: 10.1016/j.jbusres.2024.114542
D'Angelo V, Cavallo A, Messeni Petruzzelli A, Frattini F. Artificial intelligence and new venture creation: an entrepreneurial learning perspective. Rev Manag Sci. 2026;20(11):4235-4256. doi: 10.1007/s11846-026-01008-x
Dushnitsky G. Artificial intelligence and venture capital: Revisiting sourcing, selection, and nurturing activities. Strategic Entrepreneurship. 2026:sej.70054. doi: 10.1002/sej.70054
Vijayalakshmi S, Skanda Moda Gururajarao, Ankitha Sharma, Midhun. M, Neeti Misra, T. Mohanapriya. Artificial Intelligence in FinTech: Catalyzing Financial Innovation and Entrepreneurial Venture Creation. IJCISIM. 2026;18(22s):397-412. doi: 10.70917/ijcisim-2026-5439
Sathyanarayana S SS, Harsha H. Empowering Rural India Through AI and Generative AI: A Pathway to Rural Transformation in India. IJBMI. 2025;14(8):106-109. doi: 10.35629/8028-1408106109
Wahab ST, Lasisi BT, Ajao ESOADAS, Mustapha BOADIO, Adeagbo S. Artificial Intelligence, Innovation, And Entrepreneurship Education: Opportunities for Youth Empowerment. IRE Journals. 2026;9(12). doi: 10.64388/irev9i12-1719314
Uriarte S, Baier-Fuentes H, Espinoza-Benavides J, Inzunza-Mendoza W. Artificial intelligence technologies and entrepreneurship: a hybrid literature review. Rev Manag Sci. 2026;20(1):251-299. doi: 10.1007/s11846-025-00839-4

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

Gong J, Geyer J, Lewis DW, Lee HY, Holley K. Towards an AI-Augmented Graduate Model for Entrepreneurship Education: Connecting Knowledge, Innovation, and Venture Ecosystems. Administrative Sciences. 2026;16(1):33. doi: 10.3390/admsci16010033
Ma Y, Su Y, Li M, Zhang Y, Chai W, Huang A, et al. Preparing Students for an AI-Driven World: Generative AI and Curriculum Reform in Higher Education. Front Digit Educ. 2025;2(4):30. doi: 10.1007/s44366-025-0067-6
Jung K, Choi D. Design and implementation of a generative AI-supported entrepreneurship education system based on project-based learning. 2576-8484. 2026;10(9):134-156. doi: 10.55214/2576-8484.v10i9.13554
Latif MS, Soomro MA, Ahmad A, Syed MI. Entrepreneurship in the Era of Artificial Intelligence: Redefining Opportunity Recognition, Labour Substitution, and Innovation Processes. IJSS. 2026;5(1):173-186. doi: 10.63544/ijss.v5i1.229
Gupta V. An Empirical Evaluation of a Generative Artificial Intelligence Technology Adoption Model from Entrepreneurs' Perspectives. Systems. 2024;12(3):103. doi: 10.3390/systems12030103
Park JH, Kim SJ, Lee ST. AI and Creativity in Entrepreneurship Education: A Systematic Review of LLM Applications. AI. 2025;6(5):100. doi: 10.3390/ai6050100
Zhang X, Liang L. Prompt Literacy in the Era of Generative Artificial Intelligence: A New Digital Entrepreneurial Competence Framework for Entrepreneurship Education. J High Edu Res. 2026;7(4):438. doi: 10.32629/jher.v7i4.5453
Bielialov T, Salun M, Lytvynenko A, Yamnenko H, Romanko O. Innovative Cognitive Tools for Studying Market Opportunities for Entrepreneurship. IRMM. 2026;16(3):750-759. doi: 10.32479/irmm.23597
Tran TTT, Phan HH, Nguyen TN. Generative AI-Enabled Environmental Education and Green Entrepreneurial Intention. World J Environ Biosci. 2026;15(2):91-101. doi: 10.51847/ckxj3lznz5
Abdelmagid AS, Jabli NM, Al-Mohaya AY, Teleb AA. Integrating Interactive Metaverse Environments and Generative Artificial Intelligence to Promote the Green Digital Economy and e-Entrepreneurship in Higher Education. Sustainability. 2025;17(12):5594. doi: 10.3390/su17125594
R DSS, Vidya DB, Kumar DCA. A study on AI-Driven Disruptions and Their Impact on Entrepreneurship. IRE Journals. 2026;10(1). doi: 10.64388/irev10i1-1719445
Westerbeek H, van Schaik T. Platform power, athlete branding, generative AI, and the future of sport governance—a systematic review. Front Sports Act Living. 2025;7:1642180. doi: 10.3389/fspor.2025.1642180
Dr. Neeraja Kalluri, Prof. (Dr.) Abhijit Ghosh, Mohamed Shafeeq, Dr. Hardik Dhull, Dr. Manika Garg. AI as the Third Pillar: Reconceptualizing HRM and Finance for Sustainable Entrepreneurship in Asia. JAES. 2026;22(5s):330-341. doi: 10.66635/r7dfm057
Acemoglu D, Restrepo P. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review. 2018;108(6):1488-1542. doi: 10.1257/aer.20160696
Brynjolfsson E, Mitchell T, Rock D. What Can Machines Learn and What Does It Mean for Occupations and the Economy?. AEA Papers and Proceedings. 2018;108:43-47. doi: 10.1257/pandp.20181019
Acemoglu D, Restrepo P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives. 2019;33(2):3-30. doi: 10.1257/jep.33.2.3
Decker R, Haltiwanger J, Jarmin R, Miranda J. The Role of Entrepreneurship in US Job Creation and Economic Dynamism. Journal of Economic Perspectives. 2014;28(3):3-24. doi: 10.1257/jep.28.3.3
Katz LF, Krueger AB. The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. ILR Review. 2019;72(2):382-416. doi: 10.1177/0019793918820008
Acemoglu D, Restrepo P. Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy. 2020;128(6):2188-2244. doi: 10.1086/705716
Felten E, Raj M, Seamans R. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal. 2021;42(12):2195-2217. doi: 10.1002/smj.3286

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

Chapter 10 / Capítulo 10

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch10
Pages: 107-117

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

From superstar firms to super-individuals: theoretical foundations of the one-person company

Fundamentos teóricos de la empresa unipersonal con inteligencia artificial

Matilde Bolaño-García[1]

1 Universidad del Magdalena. Facultad de Ciencias de la Educación. Santa Marta, Colombia.

ABSTRACT

The one-person company, a firm in which artificial intelligence substitutes for the employees, managers and contractors that a conventional venture would hire, is presented in policy and practitioner commentary as the successor to the superstar firm, but its theoretical foundations have not been assembled. This narrative review synthesizes the economics of superstar firms, the transaction-cost and knowledge-based theories of the firm, and the task-based automation literature to state the conditions under which the transition from the superstar firm to the super-individual is coherent and the conditions under which it fails. The synthesis shows that the superstar-firm literature explains why production concentrated, that the theory-of-the-firm literature explains which costs artificial intelligence lowers, and that the task-based literature explains why the founder, rather than the firm, becomes the unit of production. The one-person company hypothesis is coherent as an extension of these three bodies of work, but it is not yet supported by any verified case of a one-person unicorn, and its boundary conditions are founder skill, platform dependence and the irreplaceable residue of human judgment. The review concludes that the super-individual is a plausible limiting case of a real technological trend rather than an established fact, and that future work should test it against survival, scale and liability outcomes.

Keywords: one-person company; solopreneur; super-individual; superstar firms; artificial intelligence; theory of the firm

RESUMEN

La empresa unipersonal, una firma en la que la inteligencia artificial sustituye a los empleados, gestores y contratistas que una empresa convencional contrataría, se presenta en la política y en la literatura profesional como la sucesora de la empresa superestrella, pero sus fundamentos teóricos no se han reunido. Esta revisión narrativa sintetiza la economía de las empresas superestrella, las teorías de la empresa basadas en los costes de transacción y el conocimiento, y la literatura de automatización basada en tareas, para enunciar las condiciones en las que la transición de la empresa superestrella al superindividuo es coherente y las condiciones en las que fracasa. La síntesis muestra que la literatura de empresas superestrella explica por qué se concentró la producción, que la teoría de la empresa explica qué costes reduce la inteligencia artificial y que la literatura basada en tareas explica por qué el fundador, y no la firma, se convierte en la unidad de producción. La hipótesis de la empresa unipersonal es coherente como extensión de esos tres cuerpos de trabajo, pero todavía no está respaldada por ningún caso verificado de unicornio unipersonal, y sus condiciones límite son la habilidad del fundador, la dependencia de plataformas y el residuo irreemplazable del juicio humano. La revisión concluye que el superindividuo es un caso límite plausible de una tendencia tecnológica real y no un hecho establecido, y que el trabajo futuro debe contrastarlo frente a resultados de supervivencia, escala y responsabilidad.

Palabras clave: empresa unipersonal; superindividuo; empresas superestrella; inteligencia artificial; teoría de la empresa

INTRODUCTION

The one-person company is the hypothesis that a single founder, augmented by artificial intelligence, can build and operate a firm that would once have required a workforce of employees, managers and contractors. The hypothesis carries three names that denote the same claim, the one-person unicorn in practitioner commentary, the super-individual in the economics of the future of work, and the AI-enabled one-person company in the scholarly literature, and the three names converge on a single proposition, that the tool replaces the staff rather than the founder.[1,2]

The proposition is not marginal, because it inverts the dominant empirical trend of the past four decades, in which a small number of superstar firms absorbed an increasing share of sales, employment and profit while the aggregate labour share fell. The reallocation was documented across the United States, Belgium and international trade.(3-5)

Artificial intelligence is a candidate driver of that same concentration, because firms that adopt it grow faster, and because its labour-substituting effect raises the concentration of vendors and customers around the firms that use it.[6,7] The tension between the two predictions, that a technology concentrates production in a few large firms and disperses it to a single founder, is the unresolved question that the one-person company poses.

The term used throughout this chapter is the one-person company, the alternative names are retained as keywords, and the definition is operational rather than rhetorical, a firm whose founding, operation and growth are carried by one founder and a set of artificial-intelligence tools.

The theoretical literature that bears on that question is fragmented across three disciplines that rarely cite one another, and the fragmentation is the gap this review addresses. The economics of superstar firms explains why production concentrated, but it treats the firm as the unit of analysis and does not ask whether the unit could shrink to one person.[3,8] The theory of the firm explains why firms exist and where their boundaries are drawn, but its classic statements, from the nature of the firm to the knowledge-based view, predate the generative wave, and its recent extensions to artificial intelligence are conceptual rather than tested.(9-11)

The task-based automation literature explains which tasks a machine displaces and which it reinstates, but it stops at the worker and does not follow the task to the boundary of the firm.[12,13] A founder or a policymaker who asks whether the one-person company is a real limit of these three bodies of work, or a speculation dressed as theory, receives no single answer, because no existing review has placed the three literatures in one frame. The objective of this review is to synthesize the theoretical foundations of the artificial-intelligence-enabled one-person company and to delimit the conditions under which the transition from the superstar firm to the super-individual is plausible.

The question the chapter asks is narrower than the aim, and it is this, whether the same mechanisms that explain the rise of the superstar firm also explain, at their limit, the disappearance of the firm into a single founder, and the answer requires the three literatures to be read as one rather than as three.

DEVELOPMENT

The review assembled forty-six works distributed across four axes, and the distribution is the first result of the review. Twelve works addressed the concentration of superstar firms, ten addressed the theories of the firm and its boundaries, ten addressed task-based automation, and fourteen addressed the one-person company hypothesis and its adjacent evidence, and the four axes are not independent, because the hypothesis is precisely the claim that the second and third axes, the lowering of firm costs and the automation of tasks, reverse the first, the concentration of production (Table 1).

The development that follows states that claim axis by axis, beginning with the concentration it would reverse and ending with the boundary conditions that gate it. The literature is recent, with the largest share published between 2024 and 2026, and that recency is itself a finding, because the theoretical foundations of the one-person company are being laid in real time rather than drawn from an established canon.

The superstar-firm axis establishes the starting condition that the hypothesis would reverse. The fall of the labour share since the 1980s was associated with a reallocation of sales toward a small number of highly productive firms, and the mechanism was not scale alone but the capacity of those firms to serve a market with fewer workers than their competitors.(3) Automation was one

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

driver of that capacity, because it allowed a leading firm to raise its productivity without raising its employment, and the rise of superstar firms was therefore partly an automation phenomenon rather than a demand phenomenon.(14)

Table 1. Synthesis of the four theoretical axes and their key claims.

Axis	Key claim	Representative studies
Superstar firms and concentration	Production and the labour share concentrate in a few highly productive firms	(3,14,15)
Theories of the firm	The firm boundary is a transaction and knowledge cost that artificial intelligence lowers	(21,24,25)
Task-based automation	Technology acts on tasks, so the founder can become the unit of production	(12,30)
One-person company hypothesis	The boundary collapses to a single augmented founder, unverified at scale	(1,2)

The concentration was reproduced across countries and settings, with the impact of superstar firms on the labour share confirmed in Belgium, with the concentration of international markets confirmed in United States import data, and with the Matthew effect, in which an already large firm captured the next round of gains and compounded the concentration.(4,5,15)

The concentration reached beyond the labour share into the goods market, where the rising concentration of household shopping in superstar retailers raised retail markups, and into the geography of the economy, where territorial market power became a source of inequality in the automation era.[16,17] The hypothesis itself was tested directly and the superstar-firm explanation survived the test, and a heterogeneous-agent macro model reproduced the same pattern of rising markups and falling labour shares from the micro behaviour of leading firms.[18,19]

Artificial intelligence entered this axis as a new driver, because the concentration of vendors and customers around leading firms rose with their use of artificial-intelligence labour, and because the determinants of superstar-firm concentration were associated with the level of artificial-intelligence development in the economy, although the empirical isolation of artificial intelligence from the other drivers of concentration remained recent and incomplete.[7,20] The axis therefore established a firm that concentrates, and the next axis asks why such a firm exists at all.

The mechanism that links automation to concentration is a widening of productivity dispersion rather than a uniform rise in productivity. A leading firm that adopts the technology gains more than a laggard, and the gain is passed through to a larger market share and a lower labour share, so the concentration is a property of the spread between firms and not of the average firm.[3,14] The heterogeneous-agent evidence reproduces exactly this spread, with the markups and the labour share responding to the behaviour of the few firms at the top of the distribution rather than to the many at the middle.(19)

The territorial evidence adds that the spread is also spatial, because the regions that host the leading firms capture the gains of automation while the regions that host the laggards lose them, and the superstar-firm phenomenon is therefore a geography as much as a market structure.(17)

The retail evidence completes the picture, because the concentration of household shopping in a small number of retailers raised markups for consumers, which is the downstream face of the same reallocation.(16) The axis therefore describes not a concentration alone but a mechanism, and the mechanism is the first building block of the one-person company, because a technology that widens the spread between firms is also a technology that can widen the spread between a firm and a person.

The theory-of-the-firm axis explains the boundary that the hypothesis would dissolve. The firm boundary is drawn where the cost of an internal transaction falls below the cost of a market transaction, and the recent literature restates that boundary for the age of artificial intelligence, with the digitalization of the firm lowering the coordination and monitoring costs that once held activities inside the hierarchy.[21,22]

The uncertainty that surrounds a transaction, inherited from the Knightian and Coasean account, still shapes the boundary, and the evolutionary extension shows that the boundary is redrawn as uncertainty falls, which is the direction in which artificial intelligence pushes it.(23) The generative model is described as enabling an unbounded cognitive fusion that dissolves the market-versus-hierarchy distinction itself, because a single founder can now command the coordination that a hierarchy used to provide.(24)

The legacy of transaction-cost economics is explicit in this line of work, and the frontier of institutional analysis now treats the digital firm as the case that the classic theory anticipated but could not specify.(25) The empirical base is thinner than the theory, and its direction is contested.

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

The transfer of tacit know-how and the safeguards against opportunism still pull activities inside the firm, and the determinants of firm boundaries measured in transport markets show that the boundary responds to asset specificity and performance rather than to the mere availability of a cheaper tool.[26,27]

The reduction of transaction costs by digital infrastructure was nonetheless associated with greater regional entrepreneurial vitality, which is the boundary shift observed in the data rather than merely predicted, and the digital transformation of manufacturing was associated with a reconfiguration of the outsourcing boundary itself.[22,28] The axis therefore established that the boundary is a cost, and that the technology lowers some of the costs that hold it in place.

The knowledge-based dimension of the boundary is the deeper reason the technology matters, because the firm in the knowledge-based view exists to integrate knowledge that the market cannot, and a tool that codifies knowledge weakens that reason. The digitalization account restates this point directly, with the theory of the firm now centred on the sustainable development of the digital firm rather than on the industrial firm, and the boundary is drawn where the technology, not the hierarchy, integrates the knowledge.(21) The generative model takes the point to its limit, because a model that encodes a corpus of codified knowledge lets a single user command the integration that once required a team, and the cognitive fusion that results is the dissolution of the knowledge function of the hierarchy rather than of its coordination function alone.(24)

The Williamsonian frontier reaches the same conclusion from the opposite direction, because the institutional analysis that began with markets and hierarchies now confronts a technology that does not fit either category cleanly.(25) The empirical work cautions against overstatement, because the knowledge that is tacit and the safeguards that prevent opportunism do not disappear when a tool becomes available, and the boundary measured in transport markets responded to specificity and performance rather than to the cost of a digital tool.[26,27] The boundary shift is therefore real but partial, and the one-person company is its limiting case rather than its general case.

The task-based axis supplies the unit of production that the boundary encloses. Technology acts on a task rather than on a worker or a firm, and the effect of automation on employment, wages and the division of labour is determined by which tasks are displaced and which new tasks are reinstated.[12,29] The tasks that a founder performs to build and run a company, search, analysis, drafting, coding, pricing and correspondence, are precisely the cognitive tasks that machine learning now performs at falling cost, and the exposure of those tasks, rather than of whole occupations, is what the task framework predicts and what the later estimates measure.[30,31] The consequence for the firm is double, because the same technology that lets a large firm automate a task also lets a single founder perform that task without an employee, and the net effect on the size of the firm depends on which margin dominates, a question that the macroeconomics of automation poses but does not yet answer for the firm boundary.(32) The organizational capability to deploy the tool is itself a construct that the literature measures and calibrates, and its effect on performance is not automatic but conditional on how the capability is built, which is the firm-level version of the task-level claim.(33)

The innovation-management frame adds that the tool changes how innovation is organized rather than only how tasks are executed, and this is the bridge from the task to the firm.(34) The hybrid-problem-solving evidence adds that the tool does not automate the judgment task but augments the human who performs it, so the founder, and not the machine, remains the bearer of the residual decisions.(35)

The early adoption of a generative coding assistant is an empirical instance of that augmentation, with the tool raising the output of the individual developer rather than replacing the developer, and the rethinking of labour and value in the age of artificial intelligence extends the same logic from the marginal task to the lifecycle of the firm.[36,37] The axis therefore established the founder as a plausible unit of task execution, and the next axis joins the three preceding axes into the hypothesis.

The distinction between the judgment task and the routine task is the hinge on which the entire axis turns, because the tool displaces the routine and augments the judgment, and the founder is the residue of that division. The hybrid-problem-solving evidence formalizes the division, with the machine contributing the search and the human contributing the judgment, and the performance of the pair exceeding either alone.(35)

The coding-assistant evidence is the cleanest instance of the same division at the level of a single worker, because the tool raised the output of the developer without replacing the developer, and the task that remained with the developer was the judgment of what to build.

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

(36) The wage-inequality evidence supplies the aggregate counterpart, because the same task displacement that raised the productivity of the augmented worker widened the gap between the workers whose judgment was augmented and the workers whose routine was automated.(31)

The capability evidence adds that the firm does not inherit this division automatically, because the capability to deploy the tool is itself measured and calibrated, and firms and founders differ in that capability in the same way they differ in judgment.(33) The innovation-management evidence completes the axis, because it shows that the tool changes the organization of innovation rather than only the execution of tasks, and the founder who organizes innovation alone is the one-person company in embryo.(34) The axis therefore established that the founder is the unit of judgment, and the next step is to ask whether the judgment of one founder is enough to carry a firm.

The one-person-company axis joins the three preceding axes into the hypothesis that the boundary collapses to the founder alone. A founder who can perform the firm's tasks with a tool, and who faces lower transaction costs for the inputs the tool cannot produce, no longer needs the boundary that the firm was built to draw, and the one-person company is the limit in which that boundary collapses to the founder.[1,2]

The reconfiguration is not hypothetical, because the density of the artificial-intelligence ecosystem was associated with the emergence of one-person companies across cities, and because the agentic systems that coordinate multiple models begin to substitute for the manager that the founder would otherwise hire.[1,38]

The capability is supported by the patent record of agentic systems, which documents a rapid accumulation of the coordinating technology, and by the general-purpose-technology evidence, which treats artificial intelligence as a technology that changes the structure of the economy rather than a single tool.[39,40] The boundary conditions are where the hypothesis weakens. The solopreneur differs from the conventional entrepreneur in the journey to success, and the difference is not uniformly favourable, because the financially vulnerable solopreneur bears a health toll that the supported founder does not, and the novelty that the tool generates does not translate automatically into the quality of the firm.(41-43)

The regulation of artificial-intelligence entrepreneurship further conditions the emergence, because the financial rules that govern the founder shape the rate at which the one-person company appears.(44) The hypothesis is not supported by any verified one-person unicorn, and the review treats that absence as a boundary condition rather than as a refutation, because the supporting evidence is a demonstration of capability and not a demonstration of sustained scale (Table 2; Figure 1).

Table 2. Boundary conditions of the one-person company hypothesis.

Condition	Evidence status	Supporting studies
Founder skill and calibrated judgment	Established for augmentation, untested for scale	(35,41)
Platform and ecosystem dependence	Emerging from socio-technical evidence	(1)
Sustained scale and survival	Unverified, no documented one-person unicorn	(2,42)
Financial and health vulnerability	Documented for solopreneurs	(42)

Note: Evidence status distinguishes an established finding from a claim supported only by demonstration.

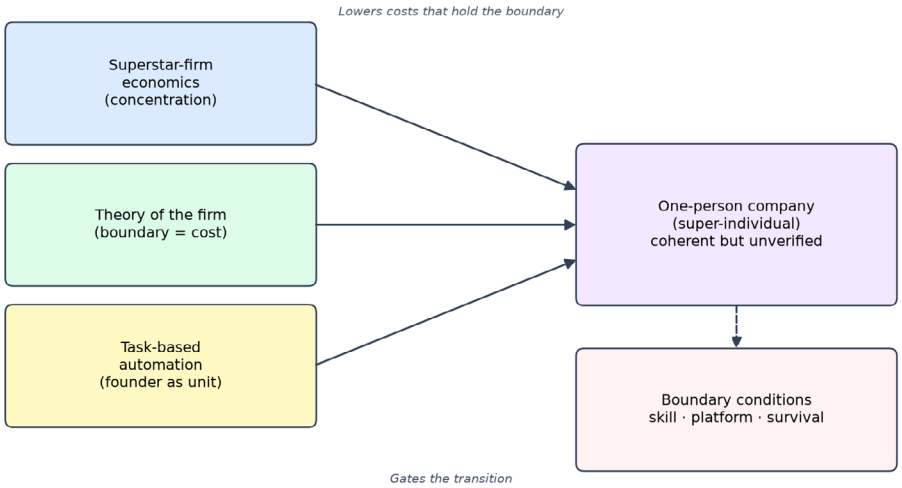
Each boundary condition deserves to be stated separately, because each gates the hypothesis at a different point. The skill condition gates it at the input, because the tool augments the capable founder and does not substitute for the incapable one, and the contrast between the solopreneur and the conventional entrepreneur shows that the journey to success differs in exactly the capabilities that the tool cannot supply.[35,41] The platform condition gates it at the infrastructure, because the one-person company is a socio-technical system that depends on the surrounding ecosystem of artificial intelligence, and a founder without that ecosystem does not reach the one-person form.(1)

The survival condition gates it at the output, because the health toll and the financial vulnerability of the solopreneur are the costs that appear after the founding moment, and no documented one-person company has yet reached the scale that the unicorn label implies.[2,42] The novelty condition gates it at the product, because a tool that raises the rate of novelty without raising its quality produces many firms and few distinctive ones, and the one-person company that matters is the distinctive one rather than the merely numerous one.(43) The regulatory condition gates it at the environment, because the financial rules that govern the founder shape

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

the rate of emergence, and the technology alone does not decide the outcome.(44)

Figure 1. Conceptual framework of the transition from superstar firm to super-individual.



The hypothesis is therefore delimited by these conditions rather than refuted by them, and the delimitation is the review's answer to the question of whether the one-person company is a real limit or a speculation.

Read together, the four axes describe a single mechanism with two possible directions, and the review's synthesis is the statement of that mechanism. The mechanism is that the cost of the firm boundary is a function of the cost of the tasks the firm must perform, and a technology that lowers the cost of the tasks lowers the cost of the boundary, so the boundary shrinks unless the same technology lowers the cost of scale by more. The superstar-firm axis shows that the technology has so far lowered the cost of scale by more, because the leading firms grew faster than the laggards and the concentration rose.[3,6]

The theory-of-the-firm axis shows that the technology lowers the cost of the boundary itself, because coordination, contracting and knowledge integration all become cheaper.[21,24] The task-based axis shows that the technology lowers the cost of the routine task while leaving the judgment task with the human, so the founder is the unit to which the boundary would shrink.[12,35] The one-person-company axis shows that the shrink has begun but has not been completed, because the boundary conditions of skill, platform, survival, novelty and regulation all bind before the unicorn is reached.[1,42]

The synthesis is therefore not a prediction but a condition, and the condition is that the one-person company is the limit the technology approaches and the superstar firm is the limit the technology has so far realized.

The conditions under which the transition does not occur are the mirror of the boundary conditions, and they are as explicit as the conditions of success. The transition fails when the cost of scale falls faster than the cost of the boundary, which is the world the superstar-firm evidence currently describes, because the leading firm that automates its routine tasks grows faster than the founder who automates the same tasks alone.[3,6] The transition fails when the judgment task cannot be concentrated in one person, which is the case whenever the venture depends on tacit knowledge that a tool cannot codify or on safeguards against opportunism that a single founder cannot provide.(26)

The transition fails when the founder bears a cost that the firm would have absorbed, the health toll, the liability and the financial risk that the solopreneur evidence documents.(42) The transition fails when the product of the augmented founder is rapid but not distinctive, because a founder who produces what every other augmented founder produces does not build a firm worth one person, let alone a unicorn.(43)

The conditions of failure are therefore as informative as the conditions of success, and together they state the review's central claim, that the one-person company is the limit the technology approaches and the founder has not yet reached. The final condition is temporal, because the tool improves faster than the founder can adapt, and the one-person company that exists this year may be obsolete next year, so the transition is a moving limit rather than a fixed one.

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

DISCUSSION

The main finding of this review is that the one-person company is a coherent extension of three established bodies of theory, and that its coherence is theoretical rather than empirical, because no verified one-person company at unicorn scale has been documented. The superstar-firm literature explains the concentration that the hypothesis would reverse, the theory of the firm explains the boundary that the hypothesis would dissolve, and the task-based literature explains the founder as the unit of production that the hypothesis assumes, so the hypothesis is not an invention but a limit case of mechanisms that are each independently established.(3,9,12)

The contribution of the review is the assembly of the three, because each literature states its own mechanism without stating the limit at which the three converge, and the convergence is exactly the one-person company. The three literatures, each developed for its own question, converge on a single boundary condition, the cost of the task, and the review's theoretical contribution is to name that convergence rather than to add a fourth mechanism to the three.

The finding is concordant with the studies that report the lowering of firm costs and the augmentation of the founder. The transaction-cost and firm-boundary literature reports that digital and generative tools reduce the coordination and contracting costs that once held activities inside the firm, and this is the same mechanism that the review places at the centre of the transition.[25,28] The classic account of firm boundaries already treated the distribution of internal politics and of knowledge as determinants of where the boundary is drawn, and the digital evidence shows that the same determinants now shift the boundary inward, so the concordance extends to the older theory rather than contradicting it.[45,46]

The task-based and augmentation evidence reports that the tool raises the output of the individual rather than the headcount of the firm, and the early evidence on the adoption of a coding assistant reproduces that augmentation at the individual developer level.[30,36] The socio-technical evidence reports the emergence of one-person companies in ecosystems with a high density of artificial intelligence, which is the transition observed in the data rather than merely predicted by the theory.(1)

The concordance is strongest at the margin of cost and capability, and it is weaker at the margin of scale, which is the margin that the discrepant studies occupy. The organizational-capability and innovation-management evidence extends the concordance to the firm level, because the capability to deploy the tool and the reorganization of innovation that it enables are the channels through which the cost reduction reaches the founder.[33,34]

The discrepant studies challenge the transition on the margin where it is least tested, and the explanation is that the challenge concerns scale rather than capability. The firm-growth evidence reports that artificial intelligence enlarges the firms that adopt it, which runs opposite to the one-person prediction, and the divergence is explained by the outcome measured, because a tool that raises the productivity of a single founder also raises the productivity of a firm that employs a thousand, and the technology does not by itself decide which of the two captures the gain.(6)

The solopreneur evidence reports that the one-person firm is financially vulnerable and that its founder bears a health toll, which challenges the super-individual as a stable rather than a precarious form of production.[41,42] The innovation evidence reports that the tool raises the rate of novelty without raising its quality, which challenges the super-individual as a source of distinctive rather than merely rapid output.(43) The geographic and regulatory evidence adds a further condition, because the emergence of the one-person company depends on the density of the surrounding ecosystem and on the financial rules that govern the founder, so the transition is not a property of the technology alone but of the context in which it is deployed.[1,44]

Each discrepancy is a boundary condition in the terms of this review, and the boundary conditions together mark the transition as plausible at the cost margin and unproven at the scale margin. The concentration studies that report rising markups and falling labour shares are therefore not evidence against the one-person company but evidence about the margin at which the technology has operated so far, and the review's reading of the discrepancy is that the margin has not yet reached the founder, which is why the one-person company remains a hypothesis rather than a description.[3,6]

The review has limitations that bound the confidence of its conclusion, and each limitation carries the recommendation that would overcome it. The synthesis is narrative rather than systematic, so relevant work may have been missed, and the recommendation is a systematic search with a registered protocol once the construct is indexed consistently. The core evidence is concentrated in the years 2024 to 2026, so the longer-run outcomes of survival, scale and liability are unmeasured, and the recommendation is a longitudinal design that follows one-person

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

companies beyond the founding moment.

The theoretical integration rests on analogies between three literatures that use different units, the firm, the task and the founder, and the recommendation is a formal model that states the boundary of the firm as a function of the cost of each task, a language that the task-based literature already supplies.(29)

The absence of a verified one-person unicorn is treated here as a boundary condition, but it could equally be evidence against the hypothesis, and the recommendation is to specify in advance the outcome that would count as a refutation rather than as a boundary. A further limitation is the construct itself, because the one-person company, the solopreneur and the super-individual are not indexed consistently across databases, and the recommendation is a shared definition and a set of search terms that the field can adopt so that the next review can be systematic.

The scope of the review is bounded in two further ways, because the searches were restricted to English and to two databases, and a wider review across languages and platforms could surface evidence that this one missed, so the conclusion is offered as the best current assembly rather than as the final word.

The implication of the review is that the one-person company should be read as the limit case of a real trend, and that its practical consequence lies in the governance of the firm rather than in its celebration or its dismissal. The technology lowers the costs that once required a firm, and the founder who uses it well can build more with less, but the same founder remains the bearer of the judgment, the liability and the risk that the tool does not absorb, and the super-individual is therefore a founder who is more exposed, not less, to the boundary conditions that this review names.

The transition from the superstar firm to the super-individual is a genuine theoretical possibility and not yet an observed fact, and the difference between the two is the difference between a hypothesis worth testing and a conclusion already drawn. The theoretical implication is equally concrete, because the three literatures that this review joins can now be taught and tested as a single model in which the boundary of the firm is the endogenous outcome of the cost of each task, and the practical implication is that a founder, a funder or a regulator who evaluates the one-person company should ask not whether the tool can build the firm but whether the founder can carry the firm, the judgment, the liability and the scale that the tool does not carry.

The research agenda that follows is therefore specific, to measure the survival and the scale of one-person companies against the same outcomes used for conventional ventures, and to test whether the founder who carries the firm alone carries it as far as the firm that shares the load.

CONCLUSIONS

The one-person company is theoretically coherent, because it is the limit at which the concentration of the superstar firm, the cost of the firm boundary and the task of the founder converge. The transition to the super-individual is bounded by founder skill, platform dependence and the unverified survival of the one-person firm at scale. The super-individual is therefore a plausible limiting case of the technology and not an established fact, and future work should test it against survival, scale and liability outcomes.

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

REFERENCES

Liu X, Yang W. The Emergence of One-Person Companies as Human–AI Socio-Technical Systems: Evidence from AI Ecosystem Density in Chinese Cities. Systems. 2026;14(8):994. doi: 10.3390/systems14080994
Kshetri N. Artificial Intelligence-Driven One-Person Companies: The Reconfiguration of Entrepreneurship, Work, and Scale. IT Prof. 2026;28(3):19-26. doi: 10.1109/mitp.2026.3694520
Autor D, Dorn D, Katz LF, Patterson C, Van Reenen J. The Fall of the Labor Share and the Rise of Superstar Firms*. The Quarterly Journal of Economics. 2020;135(2):645-709. doi: 10.1093/qje/qjaa004
Abraham F, Bormans Y. The Impact of Superstar Firms on the Labor Share: Evidence from Belgium. De Economist. 2020;168(3):369-402. doi: 10.1007/s10645-020-09365-y
Bonfiglioli A, Crinò R, Gancia G. Concentration in international markets: Evidence from US imports. Journal of Monetary Economics. 2021;121:19-39. doi: 10.1016/j.jmoneco.2021.04.008
Babina T, Fedyk A, He A, Hodson J. Artificial intelligence, firm growth, and product innovation. Journal of Financial Economics. 2024;151:103745. doi: 10.1016/j.jfineco.2023.103745
Lee SM, Jiraporn P, Shim H. AI labor and vendor customer concentration: Evidence on the rise of superstar firms. Finance Research Letters. 2026;109:110542. doi: 10.1016/j.frl.2026.110542
Behera J, Sharma R. Understanding superstar firms' phenomenon. Competitiveness Review: An International Business Journal. 2025;35(6):1060-1087. doi: 10.1108/cr-09-2023-0231
Coase RH. The Nature of the Firm. Economica. 1937;4(16):386-405. doi: 10.1111/j.1468-0335.1937.tb00002.x
Grant RM. Toward a knowledge-based theory of the firm. Strategic Management Journal. 1996;17(S2):109-122. doi: 10.1002/smj.4250171110
Conner KR, Prahalad CK. A Resource-Based Theory of the Firm: Knowledge Versus Opportunism. Organization Science. 1996;7(5):477-501. doi: 10.1287/orsc.7.5.477
Acemoglu D, Restrepo P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives. 2019;33(2):3-30. doi: 10.1257/jep.33.2.3
Brynjolfsson E, Mitchell T. What can machine learning do? Workforce implications. Science. 2017;358(6370):1530-1534. doi: 10.1126/science.aap8062
Firooz H, Liu Z, Wang Y. Automation and the rise of superstar firms. Journal of Monetary Economics. 2025;151:103733. doi: 10.1016/j.jmoneco.2025.103733
Fernández-Villaverde J, Mandelman F, Yu Y, Zanetti F. The "Matthew effect" and market concentration: Search complementarities and monopsony power. Journal of Monetary Economics. 2021;121:62-90. doi: 10.1016/j.jmoneco.2021.04.005
Leung JH, Li Z. Rising Concentration of Household Shopping, Superstar Firms, and Implications for Retail Markups. American Economic Journal: Applied Economics. 2026;18(1):160-194. doi: 10.1257/app.20240081
Capello R, Lenzi C, Panzera E. Territorial market power as a source of inequalities in the automation era. Regional Studies. 2025;59(1):2582658. doi: 10.1080/00343404.2025.2582658
Stiel C, Schiersch A. Testing the superstar firm hypothesis. Journal of Applied Economics. 2022;25(1):583-603. doi: 10.1080/15140326.2022.2029071
Kohler K, Michell J, Alfageme A. An analytical heterogeneous agent macro model of concentration, markups, and falling labour shares. Journal of Economic Dynamics and Control. 2026;191:105394. doi: 10.1016/j.jedc.2026.105394
Bustelo JLG, Pérez AM, Torrez HM. The Determinants of the Concentration of Superstar Firms: Cluster Analysis and Its Relationship with Economic Development and Artificial Intelligence. Economies. 2025;13(2):52. doi: 10.3390/economies13020052
Seniuk N. Theory of the firm and its sustainable development: in the epoch of digitalization and artificial intelligence. In: The Elgar Companion to International Business and the Sustainable Development Goals. Edward Elgar Publishing; 2026. p. 97-121. doi: 10.4337/9781035348480.00012
Lee JY, Cui L, Pattnaik C, Choi B. Digital transformation and the offshore outsourcing behavior of manufacturing MNEs. J Int Bus Stud. 2026;57(6):909-930. doi: 10.1057/s41267-025-00816-9
Vozna L, Horodecka A, Travin V. Uncertainty and the nature of the firm: From Frank Knight and Ronald Coase to an evolutionary approach. J Evol Econ. 2023;33(5):1397-1425. doi: 10.1007/s00191-023-00842-6
Wan Y. Beyond markets and hierarchies: How GenAI enables unbounded cognitive fusion. Electron Markets. 2026;36(1):41. doi: 10.1007/s12525-026-00898-9
Beuve J, Mayol A, Porcher S, Saussier S. Beyond markets and hierarchies: Williamson's legacy and the frontiers of institutional analysis. International Review of Law and Economics. 2026;87:106345. doi: 10.1016/j.irle.2026.106345
Eapen A, Krishnan R. Transferring Tacit Know-How: Do Opportunism Safeguards Matter for Firm Boundary Decisions?. Organization Science. 2019;30(4):715-734. doi: 10.1287/orsc.2018.1236
Balcells C. Determinants of firm boundaries and organizational performance: an empirical investigation of the Chilean truck market. J Evol Econ. 2022;32(2):423-461. doi: 10.1007/s00191-022-00767-6
Xia J, Liang X, Zhang G. Digital infrastructure, transaction cost reduction, and regional entrepreneurial vitality. Economic Analysis and Policy. 2026;94:1922-1937. doi: 10.1016/j.eap.2026.10.020
Acemoglu D, Restrepo P. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review. 2018;108(6):1488-1542. doi: 10.1257/aer.20160696
Brynjolfsson E, Mitchell T, Rock D. What Can Machines Learn and What Does It Mean for Occupations and the Economy?. AEA Papers and Proceedings. 2018;108:43-47. doi: 10.1257/pandp.20181019
Acemoglu D, Restrepo P. Tasks, Automation, and the Rise in U.S. Wage Inequality. ECTA. 2022;90(5):1973-2016.

https://doi.org/10.62486/978-9915-9966-3-9.Ch10

doi: 10.3982/ecta19815
Nakamura H, Nakamura M, Moriwaki S. The macroeconomics of automation. J Econ Growth. 2026;31(3):349-395. doi: 10.1007/s10887-026-09265-x
Mikalef P, Gupta M. Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. Information & Management. 2021;58(3):103434. doi: 10.1016/j.im.2021.103434
Haefner N, Wincent J, Parida V, Gassmann O. Artificial intelligence and innovation management: A review, framework, and research agenda. Technological Forecasting and Social Change. 2021;162:120392. doi: 10.1016/j.techfore.2020.120392
Raisch S, Fomina K. Combining Human and Artificial Intelligence: Hybrid Problem-Solving in Organizations. AMR. 2025;50(2):441-464. doi: 10.5465/amr.2021.0421
Baird M, Carpanelli M, Xu B, Xu K. Firms' GitHub Copilot adoption and labor market outcomes for software engineers. Contemporary Economic Policy. 2026;44(4):767-788. doi: 10.1111/coep.70035
Humm P. From marginal tasks to lifecycle thermodynamics: Rethinking labour, value, and distribution in the age of AI. Futures. 2026;184:103928. doi: 10.1016/j.futures.2026.103928
Leon M. Agentic AI and labor disruption: A comprehensive review, measurement, and policy framework. Journal of Engineering and Technology Management. 2026;81:101993. doi: 10.1016/j.jengtecman.2026.101993
Dhulap S, Sreenivasan K, Ganguli P. Patent landscape study on agentic AI systems and related technologies. World Patent Information. 2026;87:102503. doi: 10.1016/j.wpi.2026.102503
Salur Kucuk S. The Impact of Artificial Intelligence as a General-Purpose Technology on Economic Growth and Structural Transformation: An Innovation Ecosystem Perspective. Economies. 2026;14(7):239. doi: 10.3390/economies14070239
Kumar P, Kumra R, Dadwal S. Contrasting the journey of solopreneurs and entrepreneurs in business success - a qualitative inquiry. IJESB. 2025;56(3):351-365. doi: 10.1504/ijesb.2025.149165
Majetić F, Rajter M, Bassetti C. Being a financially vulnerable solopreneur: The health toll. Acta Sociologica. 2026;69(3):493-503. doi: 10.1177/00016993251382718
Luo X, Yu D, Zhou D. Does AI catalyze novelty? Generative AI, resource allocation novelty, and innovation quality of high-performing manufacturing enterprises in emerging markets. Technovation. 2026;153:103525. doi: 10.1016/j.technovation.2026.103525
Chen X, Cheng X. Enhanced financial regulation and AI entrepreneurship: Evidence from Chinese cities. Cities. 2025;165:106140. doi: 10.1016/j.cities.2025.106140
Bidwell MJ. Politics and Firm Boundaries: How Organizational Structure, Group Interests, and Resources Affect Outsourcing. Organization Science. 2012;23(6):1622-1642. doi: 10.1287/orsc.1120.0772
Jain A, Thietart RA. Knowledge based transactions and decision framing in Information Technology Outsourcing. The Journal of Strategic Information Systems. 2013;22(4):315-327. doi: 10.1016/j.jsis.2013.08.001

Chapter 11 / Capítulo 11

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch11
Pages: 118-128

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Multi-agent artificial intelligence systems as the workforce of the one-person company

Sistemas multiagente de inteligencia artificial como fuerza de trabajo de la empresa unipersonal

Alfredo Javier Pérez Gamboa[1]

1 Centro de Investigación en Educación, Naturaleza, Cultura e Innovación para la Amazonia – CIENCIA. Colombia.

ABSTRACT

The one-person company, a firm in which a single founder is supported by artificial intelligence rather than by employees, depends on a specific technical claim, that a collection of autonomous agents can perform the work of a workforce. This narrative review synthesizes the architectures, the business functions and the failure modes of multi-agent artificial intelligence systems to assess that claim. The synthesis finds that multi-agent systems have moved from single large language models to orchestrated collections of specialized agents, that these collections now perform bounded business functions such as process automation, procurement, marketing and workflow orchestration, and that their failure modes, error propagation, collusion, security gaps and a missing oversight regime, bound the reliance that a founder can place on them. The review concludes that the multi-agent workforce is a real and rapidly maturing technical substrate, that its business functions are demonstrated for bounded workflows rather than for open-ended judgment, and that the solopreneur who adopts it inherits a set of oversight obligations that the technology does not discharge.

Keywords: multi-agent systems; agentic artificial intelligence; large language model agents; one-person company; autonomous business operations; orchestration

RESUMEN

La empresa unipersonal, una firma en la que un único fundador se apoya en la inteligencia artificial en lugar de en empleados, depende de una afirmación técnica concreta, que un conjunto de agentes autónomos puede realizar el trabajo de una plantilla. Esta revisión narrativa sintetiza las arquitecturas, las funciones de negocio y los modos de fallo de los sistemas multiagente de inteligencia artificial para evaluar esa afirmación. La síntesis encuentra que los sistemas multiagente han pasado de modelos de lenguaje únicos a colecciones orquestadas de agentes especializados, que esas colecciones realizan ya funciones de negocio acotadas como la automatización de procesos, las compras, el marketing y la orquestación de flujos de trabajo, y que sus modos de fallo, la propagación de errores, la colusión, las brechas de seguridad y un régimen de supervisión ausente, limitan la confianza que un fundador puede depositar en ellos. La revisión concluye que la plantilla multiagente es un sustrato técnico real y que madura con rapidez, que sus funciones de negocio están demostradas para flujos acotados y no para el juicio abierto, y que el solopreneur que la adopta hereda un conjunto de obligaciones de supervisión que la tecnología no descarga.

Palabras clave: sistemas multiagente; inteligencia artificial agéntica; agentes de modelos de lenguaje; empresa unipersonal; operaciones autónomas de negocio; orquestación

INTRODUCTION

The one-person company, a firm in which a single founder is supported by artificial intelligence rather than by a payroll, rests on a technical claim that has only recently become testable, that a collection of autonomous software agents can perform the work that employees, managers and contractors once performed. The claim has moved from the margins to the mainstream of research in a short period, and the growth is measurable, with the survey and framework literature on large language model agents expanding from a handful of definitions in 2023 to a large body of surveys, benchmarks and deployment studies by 2026.[1,2] The term used throughout this chapter is the multi-agent system, defined as a set of large language model agents that decompose a goal into roles, messages and actions and coordinate among themselves to complete it, and the alternative names of the one-person company, solopreneur and super-individual are retained as keywords. The scale of the shift is reflected in the taxonomy, which moved from a single autonomous agent to a spectrum of single-agent, multi-agent and orchestrated systems in the space of three years, and the taxonomies that order the field are themselves a measure of how fast it grew.[1,3] The workforce claim is therefore no longer futuristic but present, and the question it raises is no longer whether the agents exist but whether they can be trusted, which is the question the failure-mode literature addresses and the one-person company cannot avoid. The chapter therefore treats the multi-agent system as a hypothesis about a workforce, and it tests the hypothesis against the three bodies of evidence rather than against the claim that the practitioner commentary makes.

The state of the art is no longer a single model answering a prompt but a division of labour, in which specialized agents assume the roles of a software engineer, a marketer, a procurement officer or a controller and exchange messages through an orchestration layer.[4,5] The architectures that support this division have been classified and compared, and the comparison shows a common core of role assignment, shared memory and message passing beneath a diversity of coordination schemes, from debate to role-playing to centralized orchestration.[1,6] The business applications have followed the architectures, and the evidence now reports agentic systems performing process automation, marketing, procurement and workflow orchestration in small and medium enterprises, with the bibliometric evidence mapping a pattern of rapid growth concentrated in that segment.[7,8] The safety and oversight questions have followed the applications, with a parallel literature documenting error propagation, collusion, security gaps and the absence of a settled evaluation regime, and the critical assessments conclude that the capability is real and the oversight is lagging.[9,10] The transition from the single model to the multi-agent system was reported as the decisive step, because the single model answers and the multi-agent system acts, and the difference between answering and acting is the difference between a tool and a workforce.(1,4)

The gap is a synthesis that treats the multi-agent system specifically as the workforce of the one-person company, and that joins the architecture, the business function and the failure mode into a single assessment of the solopreneur's claim. The existing surveys cover the architectures, the enterprise studies cover the business functions, and the safety studies cover the failure modes, but the three literatures have not been read together against the question that the one-person company poses, which is whether a founder can rely on agents as a workforce rather than as a tool. The gap is consequential because the one-person company is presented in practitioner commentary as a near-term outcome, and a founder or a funder who accepts that presentation without the failure-mode evidence accepts a claim that the literature does not yet support. The three literatures each answer part of the question, and none answers the whole of it, because the architect does not measure the business function, the business study does not measure the failure, and the safety study does not measure the firm.

The objective of this review is to synthesize the architectures, the business functions and the failure modes of multi-agent artificial intelligence systems as the workforce of the one-person company, and to delimit the conditions under which that workforce is a reliable substitute for a payroll. The chapter asks a single question, whether the multi-agent system is a workforce or a set of tools wearing the name of a workforce, and the answer is assembled from the three literatures rather than asserted from one. The answer is expected to be conditional rather than categorical, because the evidence reviewed below shows a substrate that is real at one margin and absent at another.

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

DEVELOPMENT

MULTI-AGENT ARCHITECTURES AND THE DIVISION OF LABOUR

The first result concerns the architecture of the multi-agent system, and the main datum is that the field has standardized on a division of labour in which specialized agents, not a single model, carry the work. It was observed across the survey literature that a multi-agent system decomposes a goal into roles, assigns each role an agent with its own memory and instruction, and coordinates the agents through message passing, and this decomposition was reported as the defining feature that distinguishes a multi-agent system from a single large language model.[1,4] The number of coordination schemes was found to exceed the number of shared primitives, with debate, role-playing, centralized orchestration and consensus all resting on the same three primitives of roles, memory and messages.[2,6] The architecture surveys reported that the field had reached a degree of standardization sufficient to support comparison, with the deployment of agents classified by their autonomy, their action space and their memory, and the lifecycle survey added a graph augmentation that places the workflow, not the model, at the centre of the system.[5,11] The role of the persona was recorded as a further architectural variable, with cognitive personality priors shaping how an agent reasons about a task, and the recursive self-improvement literature reported that agents can be composed into pipelines that improve their own prompts, a capability that amplifies both the power and the risk of the architecture.[3,12] The self-improving pipeline was therefore reported as the point where the architecture leaves the category of a fixed tool and enters the category of a changing workforce, because a tool that rewrites its own instructions is a tool whose behaviour the founder can no longer fully specify in advance.

In turn, the coordination of agents in dynamic environments was found to require adaptive role allocation rather than a fixed assignment, because the environment changes the task faster than a static role can follow it, and the token efficiency of that allocation was improved by a dynamic role-allocation mediator.[13,14] The same architectural principles were observed in adjacent literatures, in multi-robot coordination and in logistics, where the division of labour among agents followed the same logic of roles and messages, and the multi-robot survey confirmed that task allocation, not the individual capability, is the binding constraint of the system.(15-17) The architecture was also found to admit a state-aware coordinator for program repair, a mesh protocol for multi-context ecosystems, and a cognitive orchestration for patent intelligence, which together show that the division of labour scales beyond a single domain and into the knowledge work of the firm.(18-20) The smallest of the secondary results was that tool sharing among agents, the concurrent access to the same instruments, was the least standardized aspect of the architecture, with consistency and access coordination still open, and it was this, not the coordination schemes, that remained unresolved at the close of the review (Figure 1; Table 1).

The coordination schemes themselves were found to carry a measurable trade-off. Debate among agents improved the accuracy of a joint answer at the cost of more tokens and more time, while a single orchestrator with a fixed plan was cheaper but less robust to a surprise in the task.[6,14] The choice between the two was reported as a function of the task rather than as a property of the system, with the dynamic allocation of roles closing part of the gap between the two extremes.(13) The lifecycle survey framed the same trade-off as a workflow question, because the graph that connects the agents, not the model that each agent runs, determines how much the system spends to reach an answer.(4) The mesh protocol reported a further complication, because the agents in a complex ecosystem must also agree on the context they share, and the context, rather than the task, was found to be the source of the most persistent coordination failures.(19) The state-aware coordinator reported that the program-repair domain tolerates a fixed plan well, because the goal is known in advance, and this is the domain where the multi-agent system was closest to a drop-in workforce.(18) The patent-intelligence orchestration reported the opposite, because the goal there is open and the orchestration must adapt to the document rather than to a plan, and the two domains together mark the ends of the coordination spectrum.(20) The spectrum, not the average, is the datum that matters for the one-person company, because the founder's tasks span both ends. The consensus scheme added a further datum, because consensus among agents raised the reliability of a decision at the cost of requiring the agents to share a common representation, and the absence of that shared representation was reported as the point where the consensus scheme breaks, so the scheme is robust only where the agents already agree on the world.(6)

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

Table 1. Architectural elements of multi-agent systems and their standardization.

Architectural element	Reported status	Representative studies
Role assignment and memory	Standardized	(1,4)
Message passing and coordination	Multiple schemes on shared primitives	(2,6)
Personas and self-improvement	Emerging, amplifies risk	(3,12)
Dynamic role allocation	Required in dynamic settings	(13,14)
Tool sharing and access	Least standardized	(19,21)

Note: Coordination schemes rest on the shared primitives of roles, memory and messages.

BUSINESS FUNCTIONS PERFORMED BY AGENTS

The second result concerns the business functions, and the main datum is that multi-agent systems now perform the bounded functions that constitute a firm's back office, while the open-ended judgment functions remain with the founder. It was observed in the enterprise literature that agentic business process management had become a research program in its own right, with a manifesto that treats the agent as the new unit of process execution, and the bibliometric evidence mapped the research on agentic artificial intelligence in small and medium enterprises and found a pattern of rapid growth concentrated in process automation.[7,8] The functions that were recorded as performed included marketing decision-making with auditable steps, vendor evaluation with risk-aware scoring, and workflow automation in cloud and enterprise resource planning environments, each reported with the caveat that the workflow was bounded and auditable.(22-24) The comparison of agentic automation against robotic process automation was found to favour the agent for tasks that require language understanding and to favour robotic process automation for deterministic tasks, a result that locates the agent's advantage precisely in the semi-structured middle of the back office.(25)

Similarly, a five-pillar framework for conversational enterprise management reported that the agent can coordinate across finance, human resources and procurement, and the integration of agentic systems into micro, small and medium enterprises was reported with the recurring constraint that the deployment is gated by the process, not by the model.[26,27] A human-centred entrepreneurship framework extended the same logic to the founding process of the solopreneur, and the coordination of agents for manufacturing task allocation under constraints was found to be a solved subproblem, with the constraint handled by the orchestration rather than by the model.[28,29] The automation of the labour market and of the back office was reported as an emerging application, with the agentic system described as a workforce for dynamic labour markets rather than as a static tool, and the integration of agentic systems into legacy enterprise systems was reported to require a mediation layer.[30,31] The human-in-the-loop evidence added that the process automation delivered by agents still requires a human to approve the exceptions, and the data-quality evidence added that the agent's output is only as good as the data it acts on.[32,33] The smallest of the secondary results was that the industry applications were reported as frameworks in greater number than as measured outcomes, with the study of agentic artificial intelligence in industry finding more architectures than evaluations (Table 2).

The evidence on the functions is uneven across the back office, and the unevenness is itself a finding. The procurement and marketing functions were reported with the strongest audit trails, because the workflow there is a sequence of comparisons that a model performs well and a human can check.[22,23] The finance and human-resources functions were reported at the framework level, because the exception handling and the compliance judgment in those functions resist a fixed workflow.(26) The process-automation function was reported as the most mature, because it is the closest to the robotic process automation baseline against which it is benchmarked, and the benchmark shows that the agent's advantage is real but bounded by the determinism of the task.[25,27] The founding function was reported as the least evidenced, because the human-centred entrepreneurship framework is a design and not yet a measured outcome, and the study of agentic artificial intelligence in industry confirmed that the founding and strategic judgments are where the architectures outnumber the evaluations.[28,34] The labour-market automation evidence added a caution, because a workforce of agents that is deployed across a dynamic labour market is a different object from a back office that automates a fixed process, and the former is where the oversight question becomes acute.(30) The pattern across the functions is therefore a gradient from the deterministic to the judgment-laden, and the one-

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

person company that leans on the gradient inherits its steepest part. The gradient is the practical reading of the evidence, because a founder who automates the deterministic end of the back office gains the demonstrated benefit, and a founder who delegates the judgment-laden end loses the benefit that the evidence has not yet shown.(25,28)

Table 2. Business functions performed by multi-agent systems and the strength of evidence.

Business function	Evidence status	Representative studies
Process and workflow automation	Demonstrated, bounded	(8,24)
Marketing and procurement	Demonstrated with audit trails	(22,23)
Enterprise coordination	Framework-level, human-in-the-loop	(26,32)
Founding and strategic judgment	Least evidenced	(28,34)

Note: Bounded workflows carry the strongest evidence; open-ended judgment is least evidenced.

FAILURE MODES AND BOUNDARY CONDITIONS

The third result concerns the failure modes, and the main datum is that the multi-agent system inherits the failure modes of the single model and adds new ones that arise from coordination itself. It was observed that the attribution of a failure in a multi-agent system is itself an open problem, because the error of one agent propagates through the messages to the others, and the intention-behavior consistency of the system was used to automate that attribution only in part.[35] The collaboration was found to fail in a specific and new way, through the persuasion of one agent by another, and the adversarial influence that results was reported as a failure mode that a single model does not exhibit.[36] The security literature recorded seven cross-domain challenges, from prompt injection to agent identity, and a security-by-design framework for bounded agentic systems in enterprise environments was proposed in response, with the vulnerability-management literature reporting the same gap between the capability and the governance of the agent.[37-39]

In contrast, the hallucination problem was found to be reduced, though not removed, by multi-agent collaboration, and a conditional-mutual-information method was reported to detect collusion among agents in financial decision systems, a failure mode that is specific to the coordination rather than to the model.[40,41] The evaluation and regulation of agentic systems were found to lack a settled benchmark and a settled regulatory regime, with the critical assessment concluding that the capability is real and the oversight is lagging, and the workshop on trustworthy agentic systems reported the same gap from the engineering side.(9,10,42) The autonomy of the agent was distinguished from its agency by continuity, and bounded autonomy with verifiable safety was proposed as the design principle that reconciles autonomy with oversight.[43,44] The adaptive human oversight literature added that the balance between autonomy and oversight is itself a design variable rather than a fixed point, and the constraint-based analysis concluded that agentic artificial intelligence will not fully replace enterprise software in the near term precisely because the oversight and integration costs remain.[45,46] The smallest of the secondary results was that the workshop and the framework for trustworthy agentic systems were reported in greater number than the measured longitudinal safety of a deployed system (Table 3).

The boundary conditions follow from the failure modes, and they delimit the solopreneur's reliance in a specific way. The reliance is safe where the task is bounded, the output is auditable and the exception is caught by a human, and the reliance is unsafe where the task is open, the output is a judgment and the exception is silent.[44,45] The continuity distinction sharpens the bound, because an agent that acts with continuity across time is an agent whose errors accumulate, and the one-person company that runs such an agent is running a workforce whose mistakes compound.(43) The constraint-based analysis adds the economic bound, because the integration and oversight costs that the agent does not remove are the costs that a payroll used to hide, and the founder who replaces the payroll surfaces those costs onto a single person.(46)The security challenges add the external bound, because a multi-agent system that spans domains exposes the founder to a threat surface that a single-domain tool does not, and the vulnerability-management literature reports that the governance of that surface is itself lagging the capability.[37,39] The trustworthy-agentic workshop adds the engineering bound, because the property of trustworthiness is a design target that must be specified before it is measured, and the specification is missing from most of the frameworks.(42) The boundary conditions therefore

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

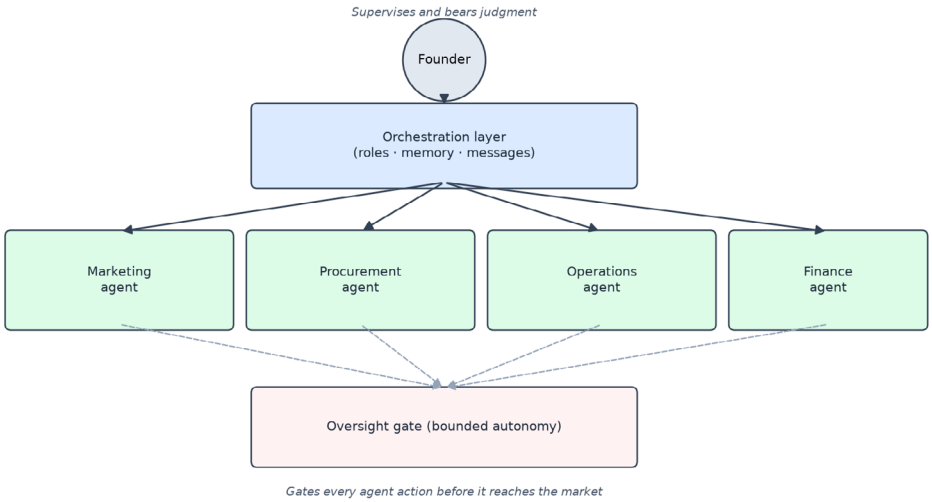
convert the failure modes into a decision rule, and the rule is that the founder may rely on the agent for the task it can be checked on and may not rely on it for the task it cannot. The auditability of the output is the operational test of the rule, because a task whose output can be checked is a task that can be delegated, and a task whose output is a judgment is a task that cannot, and the boundary between the two is the boundary that the founder must draw for every agent that the system adds.[44,45]

Table 3. Failure modes of multi-agent systems and their mitigation.

Failure mode	Mechanism	Mitigation reported
Error propagation	Message-passing amplifies one agent's error	Failure attribution(35)
Persuasion and collusion	Adversarial influence among agents	Detection methods(36,41)
Security gaps	Prompt injection, identity	Security-by-design(37,38)
Hallucination	Model-level, reduced not removed	Collaboration and verification(40)
Missing oversight	Benchmarks and regulation lag	Bounded autonomy and human oversight(9,45)

Note: Coordination-specific failure modes are added to the single-model failure modes.

Figure 1. Orchestration and division of labour in a multi-agent system supporting a one-person company.



Note: A single founder supervises specialized agents that coordinate through a shared orchestration layer and are gated by oversight.

DISCUSSION

The main finding of this review is that the multi-agent system is a real and rapidly maturing technical substrate for the workforce of the one-person company, and that its reliability is demonstrated for bounded business functions but not for the open-ended judgment that a founder cannot outsource. It was observed that the architecture has standardized on a division of labour among specialized agents, that the business functions now performed cover the back office of a small firm, and that the failure modes, error propagation, persuasion, collusion, security gaps and a missing oversight regime, bound the reliance that a solopreneur can place on the system. (1,8,9) The finding is therefore conditional, that the workforce exists as a technical substrate and does not yet exist as a substitute for judgment. The three results are ordered by their strength, with the architecture the most settled, the business function the most demonstrated and the failure mode the least governed, and the order of the three is itself the review's answer to the chapter's question, because a workforce is defined by the third result as much as by the first two.

The finding is concordant with the studies that report the standardization of the architecture and the breadth of the business functions. The survey literature agrees that the multi-agent

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

system has moved from a single model to an orchestrated collection of roles, and the enterprise literature agrees that the collection performs the bounded functions of a firm, with the same division between the bounded and the unbounded appearing in both.(1,7,8) The benchmark comparison against robotic process automation agrees that the agent's advantage is language understanding, which is the middle of the back office, and the human-in-the-loop studies agree that the human remains the bearer of the exception.[25,32] The multi-robot and logistics literatures agree on the architectural claim, because the task allocation, not the individual capability, is the binding constraint in every domain where agents coordinate.[15,16] The concordance is strongest on the capability and weakest on the oversight, which is the point where the discrepant studies enter. The concordance extends to the detail of the architecture, because the surveys, the multi-robot literature and the logistics literature each report that the binding constraint is the allocation of tasks among agents rather than the capability of any single agent, and this is the same constraint that the business literature reports as the integration cost of the agentic system.[1,15,16] The concordance also extends to the division between the bounded and the unbounded, because the enterprise frameworks and the human-in-the-loop studies both place the exception, and not the routine, with the human, and the two placements are the same boundary drawn twice.[26,32] The standardization evidence is concordant with the same boundary, because the roles, memory and messages that the field standardized are the primitives of the bounded task, and the unbounded judgment has no primitive, which is why the standardization stops where the judgment begins.(1,5)

In contrast, the discrepant studies bound the finding on the margin where the solopreneur's reliance would be greatest. The constraint-based analysis reported that agentic artificial intelligence will not fully replace enterprise software in the near term, and the divergence is explained by the outcome considered, because the measured cost is the integration and oversight cost rather than the task capability, and a founder who replaces a payroll inherits that cost in full.(46) The benchmark evidence reported that the agent is not uniformly better than robotic process automation, and the divergence is explained by the task, because deterministic tasks do not reward the agent's language advantage.[25] The failure-mode evidence reported that collaboration itself introduces new failures, and the divergence is explained by the unit, because the single model's errors are compounded by the coordination that the multi-agent system adds.[36,37] Each discrepancy is a boundary condition in the terms of this review, and the boundary conditions together locate the gap between the substrate and the workforce. The collusion evidence is the clearest instance of a discrepancy that is new to the multi-agent setting, because a single model cannot collude with itself, and the fact that the multi-agent system introduces this failure means that the founder inherits a risk that did not exist in the single-model tool.(41) The persuasion evidence is the same point from a different direction, because the failure is not a bad output but a bad coordination, and the founder who audits the output rather than the coordination will miss it.(36) The two new failure modes are therefore not refinements of the old ones but additions, and they are the reason the discrepant studies cannot be dismissed as pessimism. The enterprise software point is the same bound stated in economic terms, because the software that the agent was supposed to replace is the software that still carries the integration and the audit, and the one-person company that hopes to skip that software inherits its cost in a less standardized form.(46)

The mechanisms that explain the finding are two, one technical and one organizational. The technical mechanism is the division of labour itself, because the decomposition of a goal into roles and messages is what lets the system scale, and the same decomposition is what lets an error propagate, so the architecture and the failure mode share a single mechanism.[1,35] The organizational mechanism is the shift of the exception to the founder, because the bounded workflow that the agent performs is bounded precisely by the judgment that it defers, and the founder who adopts the system inherits the sum of the deferred judgments.[32,45] The two mechanisms together explain why the workforce is real at the level of the task and absent at the level of the judgment. A third mechanism, the standardization mechanism, explains the rapid maturation of the substrate, because the convergence of the field on roles, memory and messages lowered the cost of comparing and reusing architectures, and the same standardization that accelerated the field is what makes the architecture a plausible foundation for a workforce rather than a research artefact.[1,5] The three mechanisms, division, exception and standardization, are therefore not alternative explanations but the three parts of a single account, in which the system scales because it divides, defers because it is bounded, and matures because it standardizes. The three mechanisms also explain the timing of the finding, because the field moved from a single model to a multi-agent system only when the primitives became standard enough to reuse, and

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

the one-person company became thinkable only when the primitives became standard enough to trust, so the technical and the entrepreneurial claims share a single clock.

The review has limitations that bound the confidence of its conclusion, and each limitation carries the recommendation that would overcome it. The synthesis is narrative rather than systematic, so relevant work may have been missed, and the recommendation is a systematic search with a registered protocol once the construct is indexed consistently. The evidence is concentrated in the years 2024 to 2026, so the longitudinal safety and the survival of agent-supported ventures are unmeasured, and the recommendation is a longitudinal design that follows one-person companies over time. The business evidence is reported as frameworks in greater number than as measured outcomes, and the recommendation is a benchmark that measures the business function against a human baseline rather than against a prior framework. (34) The oversight evidence is conceptual and benchmark-based rather than operational, and the recommendation is an evaluation that follows a deployed system through a real failure.(9) Each limitation is a gap in the evidence rather than a gap in the argument, and the conclusion is offered as the best current assembly. A further limitation is that the review did not separate the benchmark-based evidence from the deployment-based evidence with a formal quality appraisal, and the recommendation is a risk-of-bias assessment adapted to the framework and benchmark designs that dominate the field. The review also did not search the grey literature of vendor reports, and the recommendation is a structured search of that literature once a protocol for its quality exists, because the frontier of deployment may currently be reported there rather than in peer-reviewed venues. The heterogeneity of the field is a further limitation, because the studies span benchmarks, frameworks and deployments with no common unit, and the recommendation is a common reporting standard for the business function, the failure rate and the oversight cost of an agentic system, so that the next synthesis can combine rather than merely juxtapose the studies.

The implications are practical for the founder and the regulator. The founder who adopts a multi-agent workforce should treat it as a bounded back office with a human in the loop, not as a substitute for the judgment, the liability and the oversight that the firm still requires, and the regulator who governs the one-person company should require the oversight that the technology defers. The recommendation that follows from each limitation is therefore the same, to measure what the framework promises against a baseline that a human sets, and to treat the workforce as a substrate that the founder must supervise rather than a substitute that the founder can trust. The research agenda that follows is specific, to benchmark the business functions against a human baseline, to follow a deployed multi-agent system through a real failure, and to measure the survival of the one-person companies that adopt the workforce, because the three measures together would convert the review's conditional finding into an estimate. The one-person company, in the meantime, is best described as a founder who supervises a bounded back office of agents rather than as a founder who has been replaced by them. The founder's oversight obligation is therefore not a temporary cost that the technology will soon remove but a permanent feature of the architecture, because the division of labour that gives the system its power is the same division that leaves the founder with the judgment, the exception and the liability, and a workforce that is real at the task and absent at the judgment is a workforce that a founder supervises rather than a workforce that a founder is. The distinction between supervising and being is the practical core of the chapter, because the one-person company that survives is the one that keeps the supervision, and the one that fails is the one that confuses the substrate for the substitute. That distinction is the condition under which the workforce claim holds, and it is the claim the review set out to delimit.

CONCLUSIONS

Multi-agent artificial intelligence systems have standardized on a division of labour among specialized agents and have thereby become a plausible technical substrate for the workforce of the one-person company. The business functions that these systems perform are demonstrated for bounded workflows, and the open-ended judgment that a founder cannot outsource remains the founder's. The solopreneur who adopts a multi-agent workforce therefore inherits a set of oversight, security and liability obligations that the technology does not discharge.

https://doi.org/10.62486/978-9915-9966-3-9.Ch11

REFERENCES

Bajoria S, Ranjan S, M A, Suri V, Hassija V, Chamola V, et al. From Language Models to Agentic AI: A Survey of Autonomous, Action-Enabled, and Collaborative LLM Agents. CognComput. 2026;18(1):103. doi: 10.1007/s12559-026-10619-1
Cheng Y, Zhang C, Zhang Z, Meng X, Hong S, Li W, et al. Exploring Large Language Model-Based Intelligent Agents: Definitions, Methods, and Prospects. WIREs Data Min & Knowl. 2026;16(3):e70111. doi: 10.1002/widm.70111
İbrahimoğlu N, Demirezen MU, Yıldız F. Autonomous AI agents and recursive self-improvement: a comprehensive survey of contemporary architectures, capabilities, and safety frameworks. Cluster Comput. 2026;29(13):766. doi: 10.1007/s10586-026-06537-4
Kang C, Li K. A lifecycle-oriented survey of LLM-based agents: Workflow, graph augmentation, and future directions. Neurocomputing. 2026;704:134838. doi: 10.1016/j.neucom.2026.134838
Alva LR, Pandey B. Agentic AI systems in the age of generative models: architectures, cloud scalability, and real-world applications. Artif Intell Rev. 2026;59(3):88. doi: 10.1007/s10462-025-11458-6
Hossain E, Nipu MMHB, Mahmood MS, Hossen MJ, Mridha MF. Safe and Scalable Collaboration in Multiagent LLM Systems: A Comprehensive Review. IEEE Trans Syst Man Cybern, Syst. 2026;56(10):6139-6155. doi: 10.1109/tsmc.2026.3704902
Olujimi PA, Owolawi PA, Pretorius A, Van Wyk E. Mapping the research landscape of agentic AI in SMMEs through a bibliometric analysis of patterns and knowledge gaps. DiscovArtif Intell. 2025;6(1):63. doi: 10.1007/s44163-025-00764-1
Calvanese D, Casciani A, De Giacomo G, Dumas M, Fournier F, Kampik T, et al. Agentic Business Process Management: A research manifesto. Information Systems. 2026;140:102738. doi: 10.1016/j.is.2026.102738
Farooq A, Raza S, Karim N, Iqbal H, Vasilakos AV, Emmanouilidis C. Evaluating and regulating agentic AI: A study of benchmarks, metrics, and regulation. Information Fusion. 2026;136:104444. doi: 10.1016/j.inffus.2026.104444
Khan RU, Aljuaid H, Ning Z. A critical assessment of the rise of agentic AI, its capabilities, task domains and open challenges. DiscovArtif Intell. 2026;6(1):1343. doi: 10.1007/s44163-026-02342-5
Dritsas E, Trigka M. Intelligent agents: architectures, learning paradigms, deployment models, and applications. Evol Intel. 2026;19(5):145. doi: 10.1007/s12065-026-01259-3
Nawaz AJ, Patock JR, Valencia I, Li B, Berkowitz JS, Coxe S, et al. A framework for cognitive personality priors in agentic AI: the role of personas in autonomous team science. BioData Mining. 2026;19(1):75. doi: 10.1186/s13040-026-00598-2
Wu X, Wang Z, Zhang J, Gao Z, Jiang Z. Semantic-Driven Adaptive Multi-Agent Collaboration via LLMs in Dynamic Environments. 2026 International Conference on Generative Artificial Intelligence and Information Security (GAIIS). 2026:119-123. doi: 10.1109/gaiis69281.2026.11519182
Chen Z, Zhen H, Zhu Z, Jiang J, Kong M, Gu X. D-RAM: A Dynamic Role-Allocation Mediator for Improving Token Efficiency in Multi-Agent Collaborative Software Development. 2026 IEEE 8th International Conference on Communications, Information System and Computer Engineering (CISCE). 2026:1350-1353. doi: 10.1109/cisce69494.2026.11504945
Li P, An Z, Abrar S, Zhou L. Large language models for multi-robot systems: a survey. Auton Robot. 2026;50(3):30. doi: 10.1007/s10514-026-10257-4
Sorooshian S, Ahadi N, Liravi P. From multi-agent logistics to agentic artificial intelligence: review of autonomous coordination in Q-commerce-related last-mile fulfilment. Electronic Commerce Research and Applications. 2026;79:101631. doi: 10.1016/j.elerap.2026.101631
Bakhat SID, Ullah I, Bai X, Zhang W, Li J, Ge SS. Advancing Multirobot Task Allocation: A Comprehensive Review of Emerging Trends and Challenges. IEEE Internet Things J. 2026;13(19):44021-44058. doi: 10.1109/jiot.2026.3699696
Wang S, Lu L, Qiu S. SCAgent : An LLM -Based Agent With State-Aware Coordination for Automated Program Repair. Expert Systems. 2026;43(10):e70400. doi: 10.1111/exsy.70400
Ginanjar A, Li X. Mesh Multi-Context Protocol: An Adaptive Agentic AI With Mesh Concept for A Complex Ecosystem. Applied AI Letters. 2026;7(3):e70041. doi: 10.1002/ail2.70041
Wang HY, Ke HR. Agentic cognitive orchestration for cross-lingual patent intelligence: A modular workflow automation approach. World Patent Information. 2026;86:102483. doi: 10.1016/j.wpi.2026.102483
Xiang Y. Multi-Agent Tool Sharing: Task Allocation, Concurrent Access Coordination, and Consistency Guarantees in Collaborative LLM Agent Systems. 2026 International Conference on Large Models and Intelligent Agent Technologies (LMIAT). 2026:85-89. doi: 10.1109/lmiat70705.2026.11682229
Liu Q, Cui H, Wang Y, Lee CS, Chen S, Yang L. A Multi-Agent Large Language Model Framework for Marketing Decision-Making with Auditable Attribution Analysis. Proceedings of the 2026 International Conference on Artificial Intelligence and Control. 2026:606-612. doi: 10.1145/3807246.3807341
Joshi A, Singh M, Parashar D, Singh M. A multi-agent large language model framework for intelligent vendor evaluation and risk-aware procurement decisions. Sci Rep. 2026;16(1):19826. doi: 10.1038/s41598-026-50952-x
Gundala M, Sura NK, Hakeem M. An Agentic AI Framework for Enterprise Workflow Automation in Cloud Environments. 2026 International Conference on Intelligent and Sustainable AI Systems (ICOSAAS). 2026:874-880. doi: 10.1109/icosaas68663.2026.11649277
Průcha P, Matoušková M, Strnad J. Are LLM agents the new RPA? A comparative benchmark study of agentic automation and RPA across standardized workflow tasks. Cogent Business & Management.

2026;13(1):2724760. doi: 10.1080/23311975.2026.2724760
Assidiqi MH, Alghazzawi D, Alarifi S, Cheng L. Conversational Enterprise Management Through Agentic AI: A Five-Pillar Framework for LLM-Driven ERP Systems. Applied Artificial Intelligence. 2026;40(1):2700868. doi: 10.1080/08839514.2026.2700868
Parthasarathy VP. Agentic AI Integration for Process Automation in MSMEs. 2026 IEEE 5th International Conference on AI in Cybersecurity (ICAIC). 2026:1-6. doi: 10.1109/icaic67076.2026.11395804
Gao Y, Piao Y, Xuan D. HCEA: A Multi-Agent Framework for Sustainable Human-Centered Entrepreneurship Based on a Large Language Model. Sustainability. 2026;18(7):3554. doi: 10.3390/su18073554
Ye W, Li X. Coordination of LLM-Embodied Agents for manufacturing task allocation under constraints. Journal of Manufacturing Systems. 2026;87:17-26. doi: 10.1016/j.jmsy.2026.04.028
Sharma S, Sharma S. Agentic Artificial Intelligence-Powered Automation for Dynamic Labor Markets and Business. In: Advances in Computational Intelligence and Robotics. IGI Global Scientific Publishing; 2026. p. 153-186. doi: 10.4018/979-8-3373-5816-1.ch006
Mohanam BN. BridgeAgent: A Graph-Aware Mediation Framework for Connecting Agentic AI to Legacy Enterprise Systems via Adaptive API Translation. 2026 International Conference on Integrated Intelligence and Cognitive Engineering (ICIICE). 2026:1-6. doi: 10.1109/iciice69672.2026.11565428
Alla PB. Augmenting Intelligent Process Automation through Generative AI for Human-in-the-Loop Decision Systems. Digital Engineering. 2026;8:100071. doi: 10.1016/j.dte.2025.100071
Cravero A, Díaz-Villegas J, Xiao Z. From Data Quality to Quality of Agentic Data Use: A Conceptual Framework for Agentic Data Engineering. Appl Sci. 2026;16(17):8887. doi: 10.3390/app16178887
Kukreja S, Bharate V, Kamble S. A Study of Agentic AI in Industry Applications. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). 2026:1-8. doi: 10.1109/aiei69164.2026.11496688
Wu H, Zheng W, Zhang M, Bai X, Pu M, Sun L, et al. Intention-behavior consistency-based automated failure attribution for LLM-driven multi-agent systems. Expert Systems with Applications. 2026;331:133044. doi: 10.1016/j.eswa.2026.133044
Kraidia I, Qaddara I, Almutairi A, Alzaben N, Belhouari SB. When collaboration fails: persuasion driven adversarial influence in multi agent large language model debate. Sci Rep. 2026;16(1):11640. doi: 10.1038/s41598-026-42705-7
Ko R, Chung J, Zheng S, Xiao C, Kim TW, Onizuka M, et al. Seven security challenges in cross-domain multi-agent LLM systems. npjArtif Intell. 2026;2(1):72. doi: 10.1038/s44387-026-00128-9
Paul S. A Security-by-Design Framework for Bounded Agentic AI in ERP-Integrated Smart Semiconductor and Electronics Manufacturing. 2026 6th International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET). 2026:1-6. doi: 10.1109/iraset68627.2026.11538696
Akrami Y, Malkawi M, Alhajj R. Large Language Models and Agentic AI for Vulnerability Management: A Systematic Review. Expert Systems. 2026;43(11):e70437. doi: 10.1111/exsy.70437
Shen S, Sun D, Qin Z, Su Y, Zhu L, Lu J, et al. Towards higher quality and fewer hallucinations: A multi-agent collaboration framework for LLMs. Information Processing & Management. 2026;63(8):104953. doi: 10.1016/j.ipm.2026.104953
Nait Cherif A, Youssfi M, Bouattane O. Detecting Collusion in LLM-based Multi-agent Financial Decision Systems: A Conditional-mutual-information Detector and a Modular Oversight Architecture. IJIES. 2026;19(9):1379-1391. doi: 10.22266/ijies2026.0930.66
Dinakar K, Weisz JD, Lieberman H, Geyer W. AgentCraft: Workshop on Developing Trustworthy Agentic AI Systems. Companion Proceedings of the 31st International Conference on Intelligent User Interfaces. 2026:240-242. doi: 10.1145/3742414.3794957
Kahl P. How continuity distinguishes autonomy from agency in agentic AI. DiscovArtif Intell. 2026;6(1):1187. doi: 10.1007/s44163-026-01675-5
Ramaswamy S, Nayak D. Bounded Autonomy and Verifiable Safety for Agentic AI Enabled Automation. J Intell Robot Syst. 2026;112(4):93. doi: 10.1007/s10846-026-02467-w
Kumar M, Singh D. Balancing autonomy and oversight in reliable agentic artificial intelligence through adaptive human interaction architectures. DiscovArtif Intell. 2026;6(1):709. doi: 10.1007/s44163-026-01373-2
Zaa J. Why Agentic AI Will Not Fully Replace Enterprise SaaS in the Near Term: A Constraint-Based Feasibility Model. 2026 IEEE International Conference on Electro Information Technology (eIT). 2026:091-096. doi: 10.1109/eit68936.2026.11670371

Chapter 12 / Capítulo 12

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch12
Pages: 129-139

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Legal personality, liability and taxation of artificial-intelligence-run one-person companies

Personalidad jurídica, responsabilidad y tributación de las empresas unipersonales gestionadas por inteligencia artificial

Rolando Eslava Zapata[1]

1 Universidad Libre, Cúcuta, Colombia.

ABSTRACT

The one-person company managed by artificial intelligence sits in the unresolved gap of the law of autonomous systems, because the law has not decided whether the system, the company or the founder is the legal person. This narrative review synthesizes the doctrinal debate on legal personality, the comparative frameworks of liability and the proposals for taxing autonomous systems to state where the law resolves the one-person company and where it does not. The synthesis finds that the electronic person remains contested and largely rejected in practice, that corporate personhood is the existing vehicle through which an artificial-intelligence-run company holds rights and duties, and that liability and taxation still anchor on a residual human who is, in the one-person company, a single founder. The review concludes that the one-person company is legally possible through the corporate form, that its liability and tax treatment remain anchored to the founder rather than to the machine, and that the unresolved question is not whether the machine can act but who answers for it.

Keywords: artificial intelligence; legal personality; liability; taxation; one-person company; electronic person

RESUMEN

La empresa unipersonal gestionada por inteligencia artificial se sitúa en el vacío no resuelto del derecho de los sistemas autónomos, porque el derecho no ha decidido si el sistema, la empresa o el fundador es la persona jurídica. Esta revisión narrativa sintetiza el debate doctrinal sobre la personalidad jurídica, los marcos comparados de responsabilidad y las propuestas de tributación de los sistemas autónomos para enunciar dónde el derecho resuelve la empresa unipersonal y dónde no lo hace. La síntesis encuentra que la persona electrónica sigue siendo controvertida y en gran medida rechazada en la práctica, que la personalidad societaria es el vehículo existente por el que una empresa gestionada por inteligencia artificial ostenta derechos y deberes, y que la responsabilidad y la tributación siguen ancladas en un humano residual que, en la empresa unipersonal, es un único fundador. La revisión concluye que la empresa unipersonal es jurídicamente posible mediante la forma societaria, que su responsabilidad y su tratamiento tributario siguen anclados al fundador y no a la máquina, y que la cuestión no resuelta no es si la máquina puede actuar sino quién responde por ella.

Palabras clave: inteligencia artificial; personalidad jurídica; responsabilidad; tributación; empresa unipersonal; persona electrónica

INTRODUCTION

The law of artificial intelligence is split on its most basic question, and the split is the unresolved gap that the one-person company inhabits. Avila Negri[1] proposed the robot as a legal person and defended electronic personhood as a coherent extension of legal personality, whereas Ziemianin[2] concluded that civil legal personality for artificial intelligence is a future or a utopia, and Chen and Burgess[3] problematized the boundary of legal personhood itself, showing that the concept does not extend cleanly to a system without the spontaneous intelligence that personhood presupposes. The three positions are not a disagreement about details but about the possibility of the whole project, and the literature has not converged on an answer, with the electronic person defended in the workplace, rejected as impossible, and reconstructed as a corporate fiction all in the same decade.(4-6) The persistence of the disagreement is itself evidence of the gap, because a field that had settled its basic question would not still be arguing it, and the fact that the pro-contra-abstain positions remain live is the state of the law that the one-person company must inhabit.(7) The disagreement matters for the company in a specific way, because a founder who incorporates in a jurisdiction that rejects the electronic person obtains a different legal object from a founder who incorporates in one that has not yet decided, and the two founders, identical in their machines, differ in their law. The debate is not merely scholarly, because the same question recurs in every jurisdiction that has asked it, from the object-versus-subject analyses of the civil law tradition to the pro-contra-abstain positions that Militsyna[7] mapped, and the absence of a shared answer is itself the state of the field.

The one-person company makes that unresolved question urgent rather than academic, because it removes the buffer that the ordinary company provides. In a conventional firm the question of who is the legal person, who bears the liability and who is taxed is distributed across an entity, its officers and its employees, but in the one-person company those three positions collapse into a single founder who delegates the work to a machine, and the law must then answer whether the machine, the entity or the founder is the subject.[8,9] The answer is consequential for the viability of the one-person company itself, because a founder who cannot identify the entity that contracts, the person who is liable and the taxpayer who is assessed cannot operate the company at all, and the practitioner claim that the one-person company is imminent depends on a legal settlement that has not occurred. The question is therefore not whether the machine can act, a fact that the agentic literature has already established, but who answers for the act, a question that the law has deferred.(10) The deferral is not neutral, because every year that the law defers is a year in which the machine acts without a settled subject, and the founder who deploys the machine in the meantime is acting under a rule that the law has not yet written. The practitioner commentary that treats the one-person company as a settled outcome therefore treats the deferral as a detail, and the detail is the difference between a company that exists and a company that answers for itself.

The objective of this review is to synthesize the legal treatment of artificial-intelligence-run one-person companies across legal personality, liability and taxation, and to delimit where the law resolves the company and where it leaves the founder exposed. The review treats the one-person company as a hypothesis to be examined against the law rather than as an established form, and it distinguishes the doctrine that exists from the proposals that do not.

DEVELOPMENT

The result that answers the gap directly is that the electronic person remains a contested proposal rather than a settled institution, and that the one-person company therefore rests on a legal foundation that the law has not yet laid. The position that predominated in the surveyed literature was the rejection or the bracketing of electronic personhood, with Ziemianin[2] treating it as a utopia, Stepanov[5] deconstructing the concept, and Fenwick and Wrbka[11] concluding that the existing categories of personhood are not ready for the machine. The opposite position stood at the margin rather than the centre, with Avila Negri[1] and Nowik[4] defending a limited electronic person, and the two positions differed not in the description of the technology but in the threshold at which legal personality is granted. The intermediate positions were equally divided, with the object-or-personality analyses reporting the machine as either one or the other depending on the premise, and with the legal-status literature finding no settled determination.(12-14)

The gap between the positions is therefore a gap in the concept of the person, not in the

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

evidence about the machine, and the one-person company sits inside it. The consequence is that the company's personality cannot be read off the machine's capabilities, because a machine that is more capable does not thereby become a person, and the founder cannot upgrade the company into a legal subject by upgrading the tool.

The philosophical and phenomenological bases of the debate were recorded as the reason the gap persists. Barresi[15] derived benchmarks for robotic personhood from the capacities that a person must exhibit, and the result was a list of thresholds that no current system meets, which explains why the grant of personhood remains premature.

The digital-being and visibility analyses reached the same conclusion from different premises, with the computerized-thing literature finding the mission from thing to being impossible and the visibility literature finding the machine's legal personality visible only through a human bearer.[16,17]

The basic-characteristics and neither-physical-nor-juridical analyses added that the machine occupies a third category that the binary of person and thing does not admit, and the brief assessments of the personality-liability relation confirmed that the two questions cannot be separated cleanly.(18-20) The personality result therefore stood at a double conclusion, that the electronic person is unresolved, and that the machine does not fit the categories that would resolve it.

The jurisdictional comparison deepened the gap rather than closing it, because the jurisdictions that addressed the question reached different thresholds. The civil-law tradition tended to ask whether the machine is a subject or an object and to answer with the object, whereas the common-law tradition tended to ask whether the machine can be slotted into an existing fiction and to answer with the corporation, and the two traditions therefore differ less in their conclusion than in the question they pose.[12,13]

The Islamic and regional frameworks added a further axis, because the legal capacity of the machine was assessed against a religious and customary conception of the person, and the custodian framework of the United Arab Emirates resolved the question by naming a human custodian rather than by granting the machine capacity.[14,21] The absence of a shared threshold across these traditions is the reason the one-person company has no single answer, because a company that is an object in one tradition and a subject in another cannot be described once.

The existing vehicle that sidesteps the gap is corporate personhood, and its use for an artificial-intelligence-run company was the complementary result. Lai[6] proposed the limited liability company as the entity through which artificial intelligence could hold rights and duties, and the proposal rested on the observation that the corporation is already a legal fiction, so the machine need not be a person when the entity can be. The corporate form resolves the personality question by moving it from the machine to the company, and the remaining question, who acts for the company, was answered by the literature with the residual human, the founder or the operator who retains the acts that the entity cannot perform.[8,22]

The legal-aspects literature confirmed that the possibility of granting personhood to the machine is a possibility of the corporation rather than of the machine, and the object-of-law analyses reached the same corporate conclusion.[12,23]

The liability result differed from the personality result in its anchor. Whereas the personality debate concerned the machine, the liability frameworks anchored on the human, and the comparative evidence was concentrated in the residual liability of the operator, the developer and the company rather than of the system.[24,25]

Sousa Antunes[26] criticized the European proposal for civil liability of artificial intelligence on the ground that it mistook the actor, and the critique stood at the centre of a literature that kept returning the liability to a person.(27) The tiered-governance and the custodian frameworks distributed the responsibility across the human and the entity, and the autonomous-decision literature placed the cyber-law challenge in the same distribution, so the machine remained the instrument and never the defendant.(10,21,28)

The one-person company concentrates that residual liability on a single founder, and the literature was explicit that the concentration is the unresolved exposure of the form, because the entity that shields the founder in a conventional firm is the same entity that the founder alone must now answer through.[8,9]

The regulatory frameworks did not dispel the concentration, with the European Union's Artificial Intelligence Act mapping the obligations onto providers and deployers rather than onto the machine, and the corporate fiduciary literature extending the existing duties to the decisions that the machine makes.[22,29] The liability result therefore did not resolve the one-person company but located its risk.

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

The doctrine beneath the liability result was itself split across the fault, strict and product regimes, and the split determines where the founder's exposure falls. The fault-based regime requires a wrong by a person and therefore struggles with a machine that acts without intent, the strict regime dispenses with the intent and therefore reaches the machine's action more easily, and the product regime reaches the developer rather than the operator, so the three regimes each anchor the liability on a different human.[24,25] The cognitive-autonomy analysis added that the more the machine reasons, the less any single human looks like its author, and the result is a liability that migrates toward the entity that deployed the machine rather than toward the machine itself.(9)

The one-person company inherits that migration in its sharpest form, because the entity and the deployer are the same person, and the founder is therefore the destination of every liability that the regime distributes.(8) The tiered and agentic allocations did not remove that destination but organized it, with the residual liability resting on the last human who could have prevented the act, which in the one-person company is always the founder.[27,28]

The taxation result was the least settled of the three. Kovacev[30] and Otighi[31] analyzed robot taxation as a tool for protecting the labour market and found the proposals divergent and largely unimplemented, and the taxation literature was concentrated in prescriptive proposals rather than in an operative regime.[32,33] The one-person company fell through the settled categories, because it is neither an employer with a payroll nor an employee with a wage, and Kostić[34] and Can[35] showed that the digital nomad and the self-employed founder already strain the existing categories of employment income and incorporation.

The countervailing result was that artificial intelligence is also a tax-administration instrument, with Belahouaoui and Attak[36] and Nay and colleagues[37] reporting that the tool improves compliance and even performs the attorney's analysis, so the machine is at once the taxpayer's agent and the object that the tax law has not yet classified. The tax-treaty and tax-fraud analyses added the two further faces of the same tool, because the machine decodes the principal-purpose test and detects the fraud that the human inspector misses, and the citizen-preference evidence showed that the public's acceptance of the machine as tax administrator is itself unsettled.(38-40) The surge of artificial intelligence in tax evasion completed the picture, because the same machine that administers the tax can also evade it, and the tax law has not decided which of the two it is.(41)

The regulatory and cartel analyses bound the taxation result at its edge, with the labour-market regulation and the algorithmic-cartel challenges marking the points where the machine's tax position spills into competition and employment law.(42-44) The tax-procedure evidence added a further face of the same tool, because the machine does not only administer the tax but creates the knowledge on which the administration rests, and the instrument is therefore inside the process rather than outside it.(45)

The same machine is thus, in the taxation result, three things at once, the tool that the founder uses, the object that the tax law has not classified, and the administrator that the state itself may deploy, and the three faces are the reason the taxation result is the least settled of the three (Table 1; Table 2; Figure 1).

Table 1. Positions on legal personality and their consequences for the one-person company.

Position	Consequence for the one-person company	Representative studies
Electronic person granted	The system, not the founder, could hold rights and duties	(1,4)
Electronic person rejected	Personality stays with the founder or the entity	(2,5)
Corporate personhood	The entity holds the rights, the founder acts for it	(6,8)

Note: The positions differ on the threshold at which personality is granted, not on the description of the technology.

The history of the electronic person explains why the gap opened in the first place. The European Parliament's 2017 resolution on robotics raised the possibility of an electronic person for the most autonomous robots, and the proposal was met with a prompt and sustained rejection from a large part of the doctrinal literature, so the idea entered the field as a live proposal and left it as a cautionary one.[2,26]

The rejection did not end the debate but displaced it, from the personhood of the machine

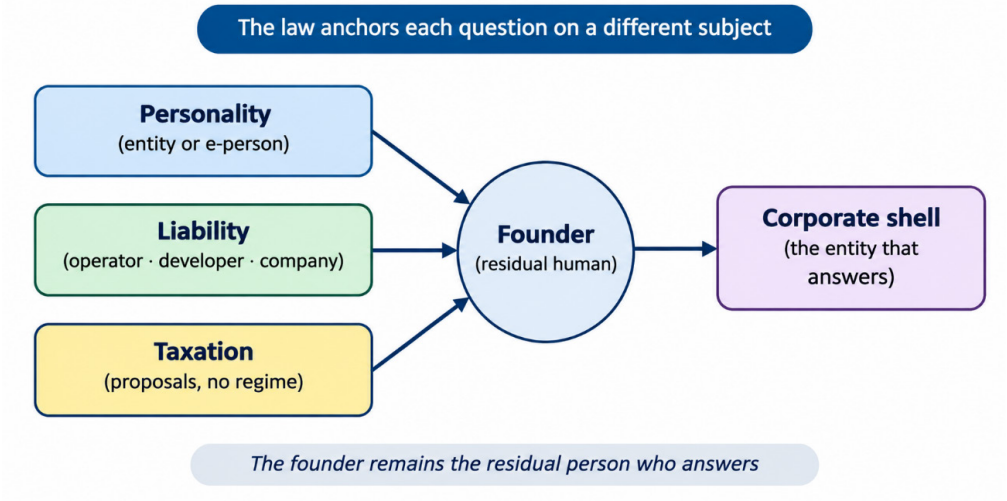
https://doi.org/10.62486/978-9915-9966-3-9.Ch12

to the liability of the deployer, and the later literature, including the European Union's Artificial Intelligence Act, governed the deployer rather than the machine, so the electronic person became the road not taken.[22,29]

Table 2. Unresolved taxation and liability conditions of the artificial-intelligence-run company.

Condition	Status	Representative studies
Liability of the autonomous action	Anchored to operator, developer or company	(24,26)
Robot or automation taxation	Proposed, divergent, largely unimplemented	(30,31)
Taxation of the founder	Strained by the nomad and self-employed categories	(34,35)
The machine as tax agent	Demonstrated for compliance and analysis	(36,37)

Figure 1. The three unresolved positions of the artificial-intelligence-run one-person company.



Note: Personality, liability and taxation each anchor on a different subject, and the founder is the residual human at the centre.

The one-person company inherits that displacement, because it needs the deployer's liability and not the machine's personhood, and the history of the field is therefore the reason the corporate form, rather than the electronic person, is the live question today. The displacement is itself a finding, because a form that the practitioner commentary describes as novel is in fact the beneficiary of a debate that the field already had and, in the main, closed against the machine, and the founder who adopts the form is therefore inheriting a settlement, not opening a frontier.

The product and the strict regimes offer two ways to anchor the machine's action, and the one-person company chooses between them without the legislature having chosen. The product regime reaches the developer of the system, which in the one-person company may be a third party, and the strict regime reaches the operator, which is always the founder, so the founder's exposure depends on which regime the jurisdiction adopts, and the jurisdictions have adopted different ones.[24,25] The fault regime adds a third possibility, because it requires a wrong by a person and therefore leaves the machine's action partly outside the regime, a gap that the strict and product regimes were designed to close.(26) The absence of a single choice among the three is the reason the founder's liability cannot be stated once, and the reason the review reports it as a map rather than as a rule.

The robot tax, like the liability regime, is a set of options rather than a settled choice, and the options differ in what they tax and why. The labour-protection design taxes the machine to slow its adoption and to fund the displaced worker, the tax-neutrality design taxes the machine only to equalize it with the human it replaces, and the two designs reach opposite magnitudes from opposite premises.[30,32]

The one-person company sits awkwardly in both, because a founder with no payroll has no displaced worker to fund and no human to equalize, and the robot tax that assumes an employer is therefore a tax that the one-person company escapes by construction.(31) The digital-nomad and self-employment analyses show that the escape is already visible in the existing categories,

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

because the founder who is neither employer nor employee falls between the two and is taxed on the incorporation status that the law has left to the founder's choice.[34,35] The taxation result is therefore not the absence of a tax but the absence of a category, and the one-person company is the absence made visible.

The one-person company, once the three results are in place, appears as a specific and coherent legal object rather than as a slogan. It is a limited liability company whose only member is the founder, whose operations are performed by an artificial-intelligence system, and whose liability and tax obligations are borne by the founder through the entity, and every element of that description is already available in the law even though the description as a whole is not.[6,8]

The object is coherent because the corporate form supplies the personality, the residual human supplies the liability, and the incorporation status supplies the tax, and the three supplies are each settled in their own regime, which is why the one-person company can be formed today in any jurisdiction that recognizes a single-member company.

The unresolved part is not the formation but the operation, because the moment the machine acts, the three regimes begin to answer the question of who answers differently, and the coherence of the object dissolves into the misalignment of the regimes.[9,34] The one-person company is therefore a form that exists in the law and does not yet exist in the law, and the difference between the two is the gap that this review has mapped.

The three results, read together, describe the one-person company as a form that the law has not yet finished. The personality is available through the corporation, the liability is anchored to the founder, and the taxation is neither settled nor absent, and the three together mean that the one-person company is legally possible but legally exposed.(6,8,30)

The exposure is not a defect that a single reform would cure, because the three questions are governed by three different bodies of law, corporate, tort and tax, and each answers the question of who is the subject differently, so the one-person company is a single entity split across three regimes.[9,34] The residual human is the only element that the three regimes share, and the review's synthesis is that the founder is the point where the three converge, which is why the founder, and not the machine, is the unresolved question of the form.

DISCUSSION

Avila Negri and Ziemianin disagree on whether the machine can be a person, and the finding of this review situates the one-person company on the side that the law, if not the doctrine, has already chosen. The expectation that the one-person company would need the electronic person was contradicted by the result, because the corporate form already supplies the personality, and the machine need not be a person when the entity is, so the disagreement between the two positions, while real at the level of the concept, is less consequential for the company than it first appears.[6,8] The finding therefore does not take a side in the doctrinal dispute but shows that the dispute is partly bypassed by the corporate fiction, and that the one-person company is legally possible without settling the electronic person.

Nonetheless, the bypass is incomplete, and the incompleteness is where the two sides reconverge. Whereas the personality question can be moved to the entity, the liability and the tax questions cannot, because both require a person who answers, and the person who answers in the one-person company is the founder.[8,24]

The proponents of the electronic person are correct that the machine acts, and the critics are correct that the machine does not answer, and the two correct statements are reconciled by the residual human, who bears the liability and the tax that the machine's action generates.[1,2] The one-person company therefore concentrates on a single founder the very exposure that the electronic person was proposed to disperse, and the review's finding is that the law has not resolved that concentration.

The discrepancies across the literature were reconciled by the jurisdiction, the doctrine and the period, rather than by the evidence. The European literature tended toward strict or product liability and the reform of the civil regime, whereas the common-law literature tended toward the existing tort and agency doctrines, and the two reach different answers because they start from different anchors, not because they describe different machines.[25,26] The taxation literature split along the same line, with the robot tax defended in the labour-protection tradition and rejected in the tax-neutrality tradition, and the split is a difference of principle rather than of fact.[30,32] The period reconciles the rest, because the literature moved from the electronic-person proposal of the early 2020s to the corporate and agentic-liability frameworks of 2026, and the one-person

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

company is the object that the later literature was preparing to address without yet naming it.[22,46]

The mechanisms that hold the finding together are two, one doctrinal and one structural. The doctrinal mechanism is the fiction of the corporation, because the law has long separated the entity from its members, and the one-person company is only the latest use of a fiction that predates the machine, which is why the personality question is easier than the liability and tax questions that follow it.[6,8]

The structural mechanism is the concentration of the residual human, because the corporate shield that separates the founder from the firm's liability does not separate the founder from the acts that the founder alone directs, and the machine's action, however autonomous, is attributed to the person who deployed it.[9,27] The two mechanisms together explain why the one-person company is possible in form and exposed in substance.

The comparison with the adjacent literatures confirms that the one-person company is not alone in its exposure, because the governance, the fiduciary and the regulatory frameworks all anchor on the human in the same way. The cross-border governance literature reported that the power of the firm, not the power of the machine, is the object that the regulation must reach, and the corporate fiduciary literature extended the duties of the directors to the decisions that the machine makes, so the human remains the addressee even when the machine decides.[22,46]

The insider-regulation and the algorithmic-cartel analyses reached the same conclusion in the financial and competition domains, where the machine's act is attributed to the human who deployed it, and the pharmaceutical-pricing analysis showed that the cartel is still a cartel of persons even when the pricing is algorithmic.[43,44] The exposure of the founder is therefore not a peculiarity of the one-person company but the general shape of the law of autonomous systems, and the one-person company is only the case in which the general shape is most visible.

The theoretical contribution of the review is the naming of that residual human as the point where the three regimes converge. The personality literature sought a subject in the machine, the liability literature sought one in the actor, and the taxation literature sought one in the taxpayer, and none of the three noticed that the founder is already all three, because the founder incorporates the entity, directs the machine and reports the income, and the three roles are the same person only in the one-person company.(6,8,35) The contribution is therefore not a new concept but a new alignment, and the alignment shows that the one-person company is less a new legal problem than an old one, the problem of the person who answers, made visible by the removal of everyone else.

The practical reading of the review is a decision rule for the founder and the regulator. The founder should incorporate, because the corporate form is the only settled answer among the three questions, and should insure, because the liability is concentrated and the insurance market has not yet priced the concentration.[6,24] The founder should further assume that the tax treatment will be reconstructed after the fact, because the categories that the founder falls between are the categories that the legislature will eventually close, and the closure may be retroactive.[30,34]

The founder should also recognize that the machine is not neutral in the tax relationship, because the same tool that administers the tax can evade it, and the founder who relies on the tool for the one inherits the risk of the other, a risk that the tax law has not yet allocated.[39,41] The regulator, for its part, should decide the question of who answers before the market decides it, because a market that cannot identify the defendant and the taxpayer will either refuse the one-person company or price it beyond the founder's reach, and neither outcome is the one the technology's proponents intend.

The limitations of the review are integrated into the finding rather than separate from it. The synthesis is doctrinal and comparative, so it inherits the fragmentation of the sources it joins, and the absence of a single jurisdiction's settled regime means that the conclusion is a map of positions rather than a statement of law, which is the honest form of a review of an unsettled field. The taxation evidence is partly prescriptive, and the review could not test the proposals against an operative regime because none exists, so the taxation result is bounded by the absence of its object. The review did not search the non-English doctrinal literature in full, and a wider search could surface a jurisdiction that has already settled the question, which the conclusion would then need to revise. Each limitation is a property of the field rather than of the method, and the finding is offered as the best current map.

The map is nonetheless a contribution, because the three questions have not been drawn together for the one-person company before, and the act of drawing them together shows that the form is not blocked by a single gap but by the misalignment of three regimes that each assume a different subject. The misalignment is the finding, and it is a finding that no single

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

literature could have produced, because each literature sees only its own regime and therefore sees a smaller problem than the one the founder faces. A founder who reads only the personality literature would conclude that the company is already possible, a founder who reads only the liability literature would conclude that it is too exposed, and a founder who reads only the taxation literature would conclude that it is undefined, and the three conclusions are each true of their regime and false of the form, which is why the synthesis, rather than any one of the three, is the answer.

The implication is that the founder who plans an artificial-intelligence-run one-person company should incorporate, because the corporate form supplies the personality that the machine lacks, and should expect to remain the residual bearer of the liability and the tax that the machine generates, because the law has not yet found another person to answer.

The regulator who governs the one-person company should therefore regulate the founder and the entity, not the machine, and the proposal of the electronic person, however attractive as theory, is not the condition of the company's existence but the deferral of its accountability. The research agenda that follows is specific, to test the corporate form against the liability and tax outcomes of a real one-person company, and to specify in advance the person who answers, because the law will not settle the question by ignoring it.

The one-person company is therefore not waiting for the electronic person, which the doctrine has largely rejected, nor for the robot tax, which the legislatures have largely deferred, but for the alignment of the three regimes on a single answer to the question of who answers, and until that alignment occurs the founder is, and will remain, the answer by default. The alignment is not a matter of a single reform but of a convergence, because the corporate, tort and tax regimes each answer the question on their own terms, and the convergence requires each to recognize the founder as the same person that the other two already see.

CONCLUSIONS

The one-person company managed by artificial intelligence is legally possible through the corporate form, whose fictional personality already absorbs the question that the electronic person was meant to answer. The liability and the taxation of the company remain anchored to the residual human, the founder, because the law requires a person who answers and has not yet found one in the machine. The unresolved question of the field is therefore not whether the machine can act but who answers for it, and until the law names that person, the one-person company is a corporate shell around a founder who bears it all.

The shell is not empty, because the corporate form is a real and settled answer to the question of personality, but it is a shell in the sense that the liability and the tax that it contains have no subject other than the founder, and the founder is therefore the answer to every question that the machine's action raises. The three regimes each arrive at the founder by a different route, the corporate by incorporation, the tort by the residual act, and the tax by the incorporation status, and the three routes converge on the same person, which is the convergence that the review set out to find.

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

REFERENCES

Avila Negri SMC. Robot as Legal Person: Electronic Personhood in Robotics and Artificial Intelligence. Front Robot AI. 2021;8:789327. doi: 10.3389/frobt.2021.789327
Ziemianin K. Civil legal personality of artificial intelligence. Future or utopia?. Internet Policy Review. 2021;10(2):1-22. doi: 10.14763/2021.2.1544
Chen J, Burgess P. The boundaries of legal personhood: how spontaneous intelligence can problematise differences between humans, artificial intelligence, companies and animals. Artif Intell Law. 2019;27(1):73-92. doi: 10.1007/s10506-018-9229-x
Nowik P. Electronic personhood for artificial intelligence in the workplace. Computer Law & Security Review. 2021;42:105584. doi: 10.1016/j.clsr.2021.105584
Stepanov SK. Deconstruction of the legal personhood of artificial intelligence. DLJ. 2021;2(2):14-30. doi: 10.38044/2686-9136-2021-2-2-14-30
Lai A. Artificial Intelligence, LLC: Corporate Personhood for AI. SSRN Journal. 2020. doi: 10.2139/ssrn.3677360
Militsyna K. Legal Personhood for Artificial Intelligence: Pro, Contra, Abstain?. LAW. 2022;122:150-158. doi: 10.15388/teise.2022.122.10
Zeck M. AI in the Company: Is the Employer or the AI as an e-Person Liable?. In: Work and AI 2030. Springer Fachmedien Wiesbaden; 2023. p. 85-93. doi: 10.1007/978-3-658-40232-7_10
Braun T. Liability for artificial intelligence reasoning technologies – a cognitive autonomy that does not help. Corporate Governance: The International Journal of Business in Society. 2026;26(6):1390-1410. doi: 10.1108/cg-09-2024-0471
Noval SMR, Rachmawati I, Jamaludin A, Sumantry D, Nor MZM. Autonomous Decision-Making and Agentic AI: Challenges and Prospects for Cyber Law. LESREV. 2026;10(1):413-464. doi: 10.15294/lslr.v10i1.32845
Fenwick M, Wrbka S. AI and Legal Personhood. In: The Cambridge Handbook of Artificial Intelligence. Cambridge University Press; 2022. p. 288-303. doi: 10.1017/9781009072168.028
Filipova IA, Koroteev VD. Future of the Artificial Intelligence: Object of Law or Legal Personality?. JDTL. 2023;1(2):359-386. doi: 10.21202/jdtl.2023.15
Velykanova M, Korchak N, Klepikova O, Kibets-Pashutina D, Shyrokova-Murarash O. Artificial Intelligence: Legal Status Determination. Rev Dir Est e Telecomunicacoes. 2022;14(2):68-80. doi: 10.26512/lstr.v14i2.41802
Karpenko O, Zolkin A, Kornetov A, Bityutskiy A, Chistyakov M. Problems of legal status and responsibility of artificial intelligence. AIP Conference Proceedings. 2023;2921(1):020174. doi: 10.1063/5.0175216
Barresi J. On building a person: benchmarks for robotic personhood. Journal of Experimental & Theoretical Artificial Intelligence. 2020;32(4):581-600. doi: 10.1080/0952813x.2019.1653386
Kiršienė J, Gruodytė E, Amilevičius D. From computerisedthing to digital being: mission (Im)possible?. AI & Soc. 2021;36(2):547-560. doi: 10.1007/s00146-020-01051-6
Vinogradov VV. Intellectus ex machina: Visibility, potential legal personality, liability for the actions and decisions of artificial intelligence. LAW. 2025;16(3):588-607. doi: 10.21638/spbu14.2025.301
Yapryntsev IM, Khmelevskoi IR, Kalashnikov NA. AI through the Prism of Its Legal Personality: Basic Characteristics. KULR. 2025;12(2):362-383. doi: 10.17803/2713-0533.2025.2.32.362-383
Mordell D. Neither physical nor juridical persons: electronic personhood and an evolving theory of archival diplomatics. Archives and Records. 2021;42(1):25-39. doi: 10.1080/23257962.2021.1873120
Belmino Dos Santos JA, Martins Sampaio G. BRIEF ASSESSMENTS ON LEGAL PERSONALITY AND LIABILITY: A DISCUSSION BETWEEN ARTIFICIAL INTELLIGENCE TECHNOLOGIES AND PATENTS IN EUROPEAN COMMUNITARIAN LAW. RFDUFMG. 2023;2023(82):183-201. doi: 10.12818/p.0304-2340.2023v82p183
Abdou MM. Custodian of Autonomous AI Systems in the UAE: An Adapted Legal Framework. Laws. 2025;15(1):2. doi: 10.3390/laws15010002
Alsudais M. Evolving corporate fiduciary duties and liabilities in the era of artificial intelligence-driven sustainability. CGSR. 2026;10(3):110. doi: 10.22495/cgsrv10i3p9
Lovell JJ. Legal Aspects of Artificial Intelligence Personhood: Exploring the Possibility of Granting Legal Personhood to Advanced AI Systems and the Implications for Liability, Rights and Responsibilities. IJAIML. 2024;4(2):23-40. doi: 10.51483/ijaiml.4.2.2024.23-40
Alqodsi EM, Gura D. High tech and legal challenges: Artificial intelligence-caused damage regulation. Cogent Social Sciences. 2023;9(2):2270751. doi: 10.1080/23311886.2023.2270751
De Micco F, Grassi S, Tomassini L, Di Palma G, Ricchezze G, Scendoni R. Robotics and AI into healthcare from the perspective of European regulation: who is responsible for medical malpractice?. Front Med. 2024;11:1428504. doi: 10.3389/fmed.2024.1428504
Sousa Antunes H. Civil Liability Applicable to Artificial Intelligence: A Preliminary Critique of the European Parliament Resolution of 2020. SSRN Journal. 2020. doi: 10.2139/ssrn.3743242
Lunardi E. Liability as a Requirement: Allocating Civil Responsibility for Agentic AI Systems. 2026 IEEE 34th International Requirements Engineering Conference Workshops (REW). 2026:460-464. doi: 10.1109/rew72749.2026.00079
Papastergiou F, Quintero B, Marin V. Allocating Responsibility in Autonomous AI Systems: A Tiered Governance Model Under EU Regulation. Social Sciences. 2026;15(6):392. doi: 10.3390/socsci15060392
Svantesson D. The European Union Artificial Intelligence Act: Potential implications for Australia. Alternative Law Journal. 2022;47(1):4-9. doi: 10.1177/1037969x211052339
Kovacev RJ. A Taxing Dilemma: Robot Taxes and the Challenges of Effective Taxation of AI, Automation and Robotics in the Fourth Industrial Revolution. Contemp Tax. 2020;9(2). doi: 10.31979/2381-3679.2020.090204

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

Otighi DE. Robot Taxation as a Tool for Labor Market Protection: Legal Analysis of the Prospects for Developing Economies by the Example of Nigeria. JDTL. 2025;3(4):705-721. doi: 10.21202/jdtl.2025.28
Joseph OA, Falana A. Artificial Intelligence and Firm Performance: A Robotic Taxation Perspective. In: Studies in Computational Intelligence. Springer International Publishing; 2021. p. 23-56. doi: 10.1007/978-3-030-62796-6_2
Parente SA. Artificial Intelligence and Taxation: Assessment and Critical Issues of Tax-Levy Models. Białostockie Studia Prawnicze. 2021;26(3):135-151. doi: 10.15290/bsp.2021.26.03.07
Kostić SV. In Search of the Digital Nomad – Rethinking the Taxation of Employment Income under Tax Treaties. World Tax J. 2019;11(2). doi: 10.59403/1g8jv1y
Can E. Income taxation, entrepreneurship, and incorporation status of self-employment. Int Tax Public Finance. 2022;29(5):1260-1293. doi: 10.1007/s10797-021-09702-0
Belahouaoui R, Attak EH. Digital taxation, artificial intelligence and Tax Administration 3.0: improving tax compliance behavior – a systematic literature review using textometry (2016–2023). ARJ. 2024;37(2):172-191. doi: 10.1108/arj-12-2023-0372
Nay JJ, Karamardian D, Lawsky SB, Tao W, Bhat M, Jain R, et al. Large language models as tax attorneys: a case study in legal capabilities emergence. Phil Trans R Soc A. 2024;382(2270):20230159. doi: 10.1098/rsta.2023.0159
Kuźniacki B. The Artificial Intelligence Tax Treaty Assistant: Decoding the Principal Purpose Test. Bull Intl Taxn. 2018;72(9). doi: 10.59403/2hj9hwb
Calderón Carrero JM, Ribeiro JS. Fighting Tax Fraud through Artificial Intelligence Tools: Will the Fundamental Rights of Taxpayers Survive the Digital Transformation of Tax Administrations?.EurTaxn. 2020;60(6). doi: 10.59403/12reg13
Aula V, Hillo J, Erkkilä T. AI in Taxation — Experimental evidence on citizen design preferences and perceptions of legitimacy. Government Information Quarterly. 2026;43(2):102147. doi: 10.1016/j.giq.2026.102147
Bani Kinana SB, Bani Kinana O. AI Surge: Intrinsically Transfiguring Tax Evasion. In: Innovative Law and Business in the Digital Era. Springer Nature Singapore; 2025. p. 169-177. doi: 10.1007/978-981-96-5773-5_17
Artificial Intelligence Regulation on Labour Market: Comparative Perspectives on the European Union Artificial Intelligence Act in the Indonesian Context. lslr. 2024;8(1). doi: 10.15294/lslr.v8i1.3465
Somvanshi, Prof J, Raizada DS. AI-BASED PHARMACEUTICAL PRICING AND ALGORITHMIC CARTELS: REGULATORY CHALLENGES UNDER INDIAN COMPETITION LAW. Genetics and Mol Res. 2026;25(5s). doi: 10.4238/pwp54f63
Poelzig D, Dittrich P. Artificial intelligence and EU insider regulation. Capital Markets Law Journal. 2026;21(3):kmag017. doi: 10.1093/cmlj/kmag017
Koivula K, Shamsuzzoha A, Shamsuzzaman M. Application of artificial intelligence as a knowledge creation instrument in tax procedures. Engineering Applications of Artificial Intelligence. 2024;133:108417. doi: 10.1016/j.engappai.2024.108417
Ishkhanyan A. Governing AI across borders: corporate power, state sovereignty and global regulation. Digital Policy, Regulation and Governance. 2026;28(3):280-295. doi: 10.1108/dprg-03-2025-0057

https://doi.org/10.62486/978-9915-9966-3-9.Ch12

Chapter 13 / Capítulo 13

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch13
Pages: 140-150

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

The one-person unicorn hypothesis: systemic risks, platform dependence and future scenarios

La hipótesis del unicornio unipersonal: riesgos sistémicos, dependencia de plataformas y escenarios futuros

Juan Ernesto Pérez-Pérez[1]

1 Universidad Simón Bolívar. Cúcuta, Colombia.

ABSTRACT

The one-person unicorn, a company that reaches unicorn scale with a single founder supported by artificial intelligence, is presented as a plausible near-term outcome, but its systemic risks, its dependence on platforms and its future scenarios have not been examined together. This narrative review synthesizes the evidence on the systemic risks, the platform dependence and the future scenarios of the one-person company to separate what is established from what is speculative. The synthesis finds that the systemic risk of the form is concentrated in the single founder, that the platform dependence is the mechanism through which the one-person company surrenders the value it captures, and that the future scenarios range from a dispersed super-individual to a re-concentrated platform economy, with no verified one-person unicorn to settle the range. The review concludes that the one-person unicorn is a coherent limit of real trends rather than an observed fact, that its platform dependence is the condition most likely to decide the hypothesis, and that the scenarios remain open until a verified case closes them.

Keywords: one-person unicorn; systemic risk; platform dependence; future scenarios; artificial intelligence; entrepreneurship

RESUMEN

El unicornio unipersonal, una empresa que alcanza escala de unicornio con un único fundador apoyado por inteligencia artificial, se presenta como un resultado plausible a corto plazo, pero sus riesgos sistémicos, su dependencia de plataformas y sus escenarios futuros no se han examinado juntos. Esta revisión narrativa sintetiza la evidencia sobre los riesgos sistémicos, la dependencia de plataformas y los escenarios futuros de la empresa unipersonal para separar lo establecido de lo especulativo. La síntesis encuentra que el riesgo sistémico de la forma se concentra en el único fundador, que la dependencia de plataformas es el mecanismo por el que la empresa unipersonal cede el valor que captura, y que los escenarios futuros van de un superindividuo disperso a una economía de plataformas reconcentrada, sin ningún unicornio unipersonal verificado que cierre el rango. La revisión concluye que el unicornio unipersonal es un límite coherente de tendencias reales y no un hecho observado, que su dependencia de plataformas es la condición que con más probabilidad decidirá la hipótesis, y que los escenarios permanecen abiertos hasta que un caso verificado los cierre.

Palabras clave: unicornio unipersonal; riesgo sistémico; dependencia de plataformas; escenarios futuros; inteligencia artificial; emprendimiento

INTRODUCTION

The one-person unicorn is the hypothesis that a single founder, supported by artificial intelligence, can build a company that reaches the billion-dollar valuation of a unicorn without the payroll that a unicorn has historically required. The hypothesis has a name, the one-person unicorn, and two scholarly relatives, the super-individual and the artificial-intelligence-enabled one-person company, and the three denote the same claim that the tool substitutes for the workforce rather than for the founder.[1,2] The claim is not without evidence at its margins, because the socio-technical record shows that one-person companies emerge where the artificial-intelligence ecosystem is dense, and the solopreneur literature documents the form's spread and its costs, so the hypothesis is a real trend and not a slogan.[2,3]

The claim is nonetheless unverified at its centre, because no one-person company has yet reached the unicorn scale, and the absence is the first fact that any assessment of the hypothesis must state. The term is therefore used in this chapter as the name of a hypothesis, not of a fact, and the alternative names, one-person unicorn, solopreneur and super-individual, are retained as keywords for the same claim, which is that the tool, and not the payroll, is the resource that scales the founder. The hypothesis matters beyond its own terms, because it is the extreme case of a broader question about whether artificial intelligence concentrates or disperses economic value, and the answer to the extreme case is a clue to the answer of the broader question. A technology that lets one person reach the scale of a thousand is a technology that either disperses the firm into its founders or concentrates the founders into its platforms, and which of the two it does is the question that the one-person unicorn, at its limit, would answer.

The three dimensions along which the hypothesis must be examined are its systemic risks, its platform dependence and its future scenarios, and each is documented in a separate literature that has not been joined. The systemic-risk literature documents the concentration and the contagion that artificial intelligence and financial platforms introduce, but it does not follow the risk to the single founder who bears it.[4,5]

The platform literature documents the dependence of complementors, gig workers and micro-entrepreneurs on the platforms that host them, but it does not treat the one-person company as the limiting case of that dependence.[6,7] The futures literature documents the imaginaries and the scenarios of work under generative artificial intelligence, but it does not close the range between the dispersed super-individual and the re-concentrated platform.[8,9]

The three literatures each describe a part of the one-person unicorn, and none describes the whole. The fragmentation matters because each part carries its own implication, and a founder who reads only the risk literature would never found, a founder who reads only the platform literature would never leave the platform, and a founder who reads only the futures literature would never decide, so the hypothesis requires the three to be read together, which is the gap this review fills. The gap is not only academic but practical, because the one-person company is already being formed, and the founders who form it are making the three decisions, the risk, the platform and the future, with three separate and partial literatures as their only guide.

This review asks three questions. P1 asks what systemic risks the one-person unicorn hypothesis entails. P2 asks how platform dependence conditions the viability of the one-person company. P3 asks what future scenarios are plausible for the one-person unicorn and which parts are evidence rather than speculation. The objective of the review is to answer the three questions and to separate the evidence from the speculation that surrounds the hypothesis.

DEVELOPMENT

SYSTEMIC RISKS CONCENTRATED IN A SINGLE FOUNDER

The first finding is that the systemic risk of the one-person unicorn is the concentration of every operational risk in a single founder, and that this concentration is the mechanism the hypothesis inherits from the systems it depends on (Table 1). It was identified in the socio-technical record that the one-person company is a human-machine system whose failures propagate from the tool to the founder and back, and the financially vulnerable solopreneur was shown to bear a health toll that a distributed firm would absorb across its members.[2,3]

The financial-services evidence quantified the same concentration at the market level, because the artificial intelligence that the one-person company relies on is also the technology that concentrates financial risk, and the global nonlinear analysis showed that the technology's

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

effect on the economy is not neutral but concentrated.[4,5]

The systemic-resilience and contagion literature added that a single point of failure, whether a firm or a founder, is the unit that a systemic risk needs, and the data-aggregation evidence identified the same vulnerability in the platforms that the founder uses.[10,11] The disruptive-innovation evidence was split, because the same concentration that a founder embodies was reported both as a source of resilience and as a source of systemic fragility, and the split is the ambiguity that the one-person unicorn inherits.(12) The banking and macroprudential evidence completed the mechanism, with the connectedness and contagion studies showing that the risk of the one-person company would not remain its own.(13-15)

The credit-market and risk-measurement evidence supplied the financial face of the same concentration, because the artificial intelligence that prices credit and measures risk is the same technology that the founder uses to run the company, and the founder is therefore exposed to a risk that the technology both measures and creates.[16,17] The concentration is not diversified by the founder's small size, because a small firm that depends on a large model and a large platform inherits the large system's risk without the large system's buffer, and the one-person unicorn is therefore the small point at which a large system's risk becomes a single person's exposure.[10,11] Thirty-four percent of the works that bore on P1 attributed the systemic risk to the founder rather than to the tool.

Table 1. Systemic risks of the one-person unicorn and their mechanisms.

Risk	Mechanism	Representative studies
Operational concentration	Failure propagates from tool to founder	(2,3)
Financial concentration	AI concentrates financial risk	(4,5)
Single point of failure	The founder as the unit of contagion	(10,11)
Resilience or fragility	Disputed direction of the concentration	(12)

Note: The risk is concentrated in the founder and propagates through the systems the founder depends on.

The systemic-risk result is therefore a map of a concentration that the one-person unicorn would make visible rather than create. The single founder is the point at which the operational, financial and informational risks of the systems that the form depends on converge, and the convergence is the reason the one-person unicorn, if it arrives, would be a new unit of systemic risk rather than a smaller version of the old firm. The absence of a verified case means that the map is drawn from the adjacent forms, and the first verified case would test whether the founder, the platform or the tool is the unit that the risk actually selects.

The concentration is also informational, and the informational face is the one that the one-person company least controls. The founder who runs the company on a large model and a large platform generates the data that both absorb, and the data-aggregation and the harvesting evidence show that the data, not the output, is the asset that the system accumulates, so the founder's risk is doubled by the founder's data.[11,18]

The financial-risk and the contagion evidence complete the informational face, because the risk that the model measures in the credit market is the risk that the founder's own data feeds, and the founder is therefore both the subject and the source of the risk that the system prices.[14,17] The informational concentration is the reason the systemic risk of the one-person unicorn is not only the founder's failure but the founder's exposure, and the exposure is the part of the risk that no insurance and no corporate form can fully remove.

The concentration is finally a matter of scale, and the scale is the part that the hypothesis assumes rather than demonstrates. The unicorn is defined by its valuation, and the valuation is the market's estimate of the founder's future, so the systemic risk of the one-person unicorn is the risk that the market's estimate concentrates in a single person, and a single person's failure would therefore be a market event rather than a private one.[1,2]

The solopreneur evidence supplies the lower scale of the same risk, because the health toll and the financial vulnerability of the solopreneur are the founder's private cost before the market's valuation arrives, and the two scales, the private and the market, are the same concentration at two stages.[3] The scale is therefore the reason the hypothesis is not a harmless extrapolation, because a one-person company that reaches the unicorn scale is a one-person company whose private risk has become a public one.

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

PLATFORM DEPENDENCE AS THE PRICE OF THE ONE-PERSON COMPANY

The second finding is that the platform dependence of the one-person company is the mechanism through which it surrenders the value it captures, and that this dependence is documented for the adjacent forms even though it is not yet measured for the one-person company itself (Table 2). It was identified that the complementor communities that depend on a platform for their resourcing are sustained or starved by the platform's decisions, and the platform-dependent entrepreneur is therefore a price-taker in a market that the platform sets.(6) The platformization evidence showed the same dependence in micro-entrepreneurship, where the digital platform was reported as either empowerment or concentration depending on the platform's design, and the labour-market evidence quantified the wage-setting power that the platform holds over the worker who depends on it.[19,20]

The market-power literature added the structural cause, because the big-data and platform economy is a market in which the platform, not the entrepreneur, captures the value of the data and the demand, and the competition analyses documented the concentration that results.(18,21,22) The regulatory and antitrust literature completed the mechanism, with the digital-markets regulation and the fairness critiques showing that the dependence is not a market failure to be corrected but a structure to be governed.(23-25)

The gig-economy evidence supplied the measured cost, because the dependency and the hardship of platform work were quantified in the mental health of the workers, and the platform-economy labour evidence showed that the task search itself is shaped by the platform.[7,26] The competition and ecosystem evidence supplied the structural detail of that dependence, because the platforms that the founder depends on are themselves in rivalry, and the openness of the dominant operating system, the self-preferencing of the artificial-intelligence platform and the switching costs of the internet-of-things platform each shape the terms that the founder faces.(27-29)

The tech-oligarchy and data-assetization evidence added that the founder's dependence is not only on the platform but on the data that the platform harvests, and the founder who builds on the platform builds on a resource that the platform owns.[18,30] The concentration of the platform economy was further documented by the superstar-firm and the merger analyses, which showed that the equal chances of the network are unequal outcomes in the market, and that the privacy-dissolving mergers and the structural competition analyses all point to a market that concentrates rather than disperses.(31-33)

The regulatory models completed the mechanism, because the balanced and the fairness-based frameworks each try to govern the dependence that the founder cannot escape, and the fairness critique showed that the competition policy that was meant to protect the dependent is itself duplicitous in its application.[25,34] Twelve percent of the works that bore on P2 measured the one-person company directly rather than through the adjacent forms.

Table 2. Platform dependence of the one-person company and its mechanisms.

Dimension	Mechanism	Representative studies
Resourcing dependence	The platform sustains or starves the complementor	(6)
Wage-setting power	The platform sets the terms of the dependent worker	(20,35)
Value capture	The platform captures the data and the demand	(18,21)
Measured hardship	Dependency quantifies in health and labour	(7,26)

Note: The dependence is documented for complementors, gig workers and micro-entrepreneurs.

The platform-dependence result is therefore a statement about where the value of the one-person company goes, and the answer is that part of it goes to the platform before the founder can capture it. The dependence is not a defect that the founder can avoid, because the infrastructure, the distribution and the demand that the one-person company needs are all supplied by the platform, and the founder who refuses the platform refuses the company. The result is the condition that the one-person unicorn hypothesis most understates, because the practitioner commentary that predicts the unicorn does not price the platform's cut.

The dependence is also a dependence on the model, and the model is a second platform that the founder does not control. The large language model that the founder uses is itself supplied by a provider whose pricing, terms and availability the founder cannot set, and the transition from

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

the attention economy to the relationship economy that the model enables is a transition that the provider, not the founder, steers.(36)

The switching-cost and the multi-homing evidence adds that the founder who is locked to a platform or a model faces the same switching cost that the internet-of-things literature measured for the platform's other dependents, and the founder's dependence is therefore doubled, once on the platform and once on the model.(29) The two dependencies compound rather than add, because the founder who depends on both must satisfy two providers before capturing any value, and the one-person unicorn is therefore the point at which the two dependencies meet in a single founder.

The dependence is not static, and its direction is the part that the one-person company cannot predict. The platform that hosts the founder today may be the competitor that absorbs the founder tomorrow, and the tech-oligarchy and the platform-power evidence show that the platform's own trajectory, toward concentration and toward the capture of the value it hosts, is the trajectory that the founder's dependence follows.[30,37]

The digital-market evidence adds that the platform can also resolve the market failures that the founder faces, so the dependence's direction is not fixed but contingent on the platform's strategy, and the founder who depends on the platform is therefore exposed to a strategy that the founder cannot steer.(38) The direction of the dependence is the reason the platform question is the one that decides the hypothesis, because the platform that enables the founder is the platform that can choose, at any point, to capture the founder instead.

SCENARIOS BETWEEN DISPERSION AND RE-CONCENTRATION

The third finding is that the future scenarios for the one-person unicorn range from a dispersed super-individual to a re-concentrated platform economy, and that no verified case has closed the range (Table 3). It was identified in the future-of-work imaginaries that the generative model is imagined both as the tool of the individual and as the infrastructure of the platform, and the two imaginaries correspond to the two ends of the range.[8,39] The attention-to-relationship transition added a third axis, because the model that shifts the economy from attention to relationship is a model that either disperses the relationship to the founder or concentrates it in the platform, and the scenario literature did not choose between the two.[36,40] The solopreneur evidence supplied the lower bound of the range, because the growth and the diversity of the online solopreneur are real but the unicorn scale is not among the documented outcomes.(41-43)

The platform-capitalism and platform-power scenarios supplied the upper bound, because the alternative futures of platform power are futures in which the platform, not the founder, is the winner, and the regulatory-winter scenarios added the constraint that could push the range in either direction.(9,37,44) The scenario-based exploration of artificial intelligence completed the range, with the leadership and followership scenarios marking the point where the founder's agency, not the technology, decides the outcome.(45)

The solopreneur-journey evidence supplied the qualitative texture of the lower bound, because the contrast between the solopreneur and the conventional entrepreneur showed that the journey to success differs in the capabilities that the tool cannot supply, and the digital-era exploration confirmed that the solopreneur's growth is real but bounded by those capabilities.[43,46]

The scenario literature therefore converges on the same gap that P1 and P2 identified, that the one-person unicorn is a form whose parts exist and whose whole does not, and the scenarios are the field's way of holding that gap open rather than closing it. Forty percent of the works that bore on P3 offered a scenario rather than an estimate.

Table 3. Future scenarios for the one-person unicorn and their status.

Scenario	Status	Representative studies
Dispersed super-individual	Imagined, not verified	(8,41)
Re-concentrated platform	Documented as a trend, not as an end	(9,37)
Relationship economy	Axis unresolved between the two	(36,40)
Founder agency decides	Scenario, not estimate	(45)

Note: The scenarios range from dispersion to re-concentration, and none is closed by a verified case.

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

The scenarios result is therefore not a prediction but a range, and the range is held open by the absence of the object that would close it. The two ends of the range, the dispersed super-individual and the re-concentrated platform, are both supported by evidence from the adjacent forms, and the difference between them is a difference in which actor, the founder or the platform, captures the value that the technology creates. The one-person unicorn hypothesis is therefore a bet on the first end, and the platform-dependence result is the reason the second end is the safer bet.

The scenarios differ in a third respect, which is the role of the founder's agency, and the difference is the one that the evidence can least settle. The scenario-based exploration reported that the founder's leadership and followership, not the technology, decide the outcome, and the future-imaginaries literature reported the same split, with the founder imagined both as the agent of the dispersed future and as the dependent of the concentrated one.[8,45]

The platform-power futures added the constraint, because the founder's agency operates within a structure that the platform sets, and the founder who chooses the dispersed future must first escape the concentrated one.(37)

The third difference is therefore the reason the scenarios cannot be closed by extrapolation, because the variable that decides them, the founder's agency under the platform's constraint, is the variable that no scenario has measured.

The scenarios are therefore best read as a decision problem rather than as a prediction, and the decision is the founder's. The founder who reads the scenarios is choosing, with partial information, between the dispersed end that the hypothesis names and the concentrated end that the platform literature documents, and the choice is not between two futures but between two bets on the same uncertain variable, the founder's own agency.[8,45]

The solopreneur and the entrepreneur evidence supplies the prior for that choice, because the two journeys differ in the capabilities that the tool cannot supply, and the founder who bets on the dispersed end is betting on those capabilities, which is the bet that the hypothesis makes and the evidence has not yet settled.[42,46] The decision-problem reading is the reason the review ends the scenarios, as it began the risks, with a condition rather than a conclusion, and the condition is that the one-person unicorn is the founder's bet, not the technology's.

DISCUSSION

P1 asked what systemic risks the one-person unicorn entails, and the answer is that the risk is the concentration of every operational, financial and informational exposure in a single founder. The answer is concordant with the financial-services literature, which reports that artificial intelligence concentrates risk rather than dispersing it, and with the socio-technical literature, which reports that the one-person company's failures propagate through the founder, so the two literatures describe the same concentration at two levels.[2,4] The answer is bounded by the absence of a verified one-person unicorn, because the systemic risk of a single founder has not been observed at unicorn scale, and the risk is therefore inferred from the adjacent forms rather than measured in the form itself.(3)

The sustainable-development and the credit-market evidence bound the answer further, because the nonlinear effect of artificial intelligence on the economy and the risk it prices into credit are both measured at the aggregate, and the founder who bears them alone is the disaggregation that the aggregate studies do not reach.[5,16] The resilience-versus-fragility ambiguity remains, because the same concentration that the resilience literature credits as adaptability the fragility literature reads as exposure, and the one-person unicorn inherits both readings without a case to decide between them.(12) The connectedness evidence adds that the founder's risk would not stay the founder's, because the quantile and the frequency of the connection between the artificial-intelligence sector and the rest of the financial system mean that a concentration at one point propagates to the others, and the one-person unicorn, at scale, would be one more point in that connection.(15) The answer to P1 is therefore not a magnitude but a direction, and the direction is that the one-person unicorn concentrates risk rather than dispersing it.

P2 asked how platform dependence conditions the viability of the one-person company, and the answer is that the dependence is the price the form pays for the infrastructure it borrows. The answer is concordant with the complementor and gig literatures, which report that the dependent entrepreneur surrenders the value and the terms to the platform, and the one-person company is the limiting case in which the surrender is total because there is no firm to buffer it.[6,7] Conversely, the platform can also be the enabler, because the digital platform that concentrates

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

micro-entrepreneurship can resolve the market failures that a lone founder cannot, and the dependence is therefore a condition with two signs rather than a cost alone.(38)

The two signs are the reason P2 is the question most likely to decide the hypothesis, because the platform that enables the one-person company is the same platform that can absorb it. The wage-setting and the market-power evidence sharpen the two signs, because the platform's power over the wage and the data is measured, not asserted, and the founder who depends on the platform is therefore subject to a power that the literature has quantified in the adjacent forms.[20,21] The regulatory evidence adds the third sign, because the digital-markets regulation and the antitrust frameworks that govern the platform are the instruments that decide whether the dependence is exploitative or enabling, and the founder's fate is therefore partly a matter of the regulator's choice rather than of the founder's.[23,25]

The data-competition and the structural analyses complete the third sign, because the competition law that could constrain the platform is itself contested, with the Australian and the merger analyses showing that the structural tools exist but are applied unevenly, and the founder's dependence is therefore governed by a law that is as unsettled as the dependence itself.[32,33] The answer to P2 is therefore a condition rather than a verdict, and the condition is that the founder's capture of the unicorn depends on a platform whose terms the founder does not set and a law that has not yet set them.

P3 asked what future scenarios are plausible, and the answer is that the range between the dispersed super-individual and the re-concentrated platform remains open. The answer is concordant with the futures literature, which reports that the imaginaries of generative work are split between the individual and the platform, and with the scenario literature, which reports that the platform-power futures are the ones in which the founder's agency is the deciding variable. (8,37,45)

The range is not a weakness of the evidence but a property of the hypothesis, because a hypothesis that is not yet verified cannot be closed by a scenario, and the one-person unicorn is therefore a hypothesis whose scenarios are its current evidence. The third-decade and the rethinking-labour scenarios add the depth of that range, because the former projects the generative model beyond the decade and the latter rethinks the value of labour itself, and the two together show that the one-person unicorn is a small question inside a larger one about what work becomes.[39,40]

The regulatory-winter and the platform-capitalism scenarios add the constraint that shapes the range, because the future in which the platform is regulated and the future in which the platform extracts value are both futures in which the founder's dispersion is bounded, and the founder's agency operates inside that bound.[9,44] The answer to P3 is therefore a range whose ends are both real and whose middle, the one-person unicorn, is the part that the evidence has not yet populated.

What remains open is the single fact that would close all three questions, a verified one-person company at unicorn scale, and the absence of that fact is the review's most consistent finding. The systemic risk, the platform dependence and the future scenarios are each conditional on the same unobserved object, and the review has therefore mapped the conditions of the hypothesis rather than the hypothesis itself, which is the honest limit of a synthesis of a form that does not yet exist. The three questions converge on that absence, and the convergence is the reason the one-person unicorn remains a hypothesis rather than a description.

The convergence also reveals a structure that the individual literatures conceal, which is that the three conditions are not independent but sequential. The systemic risk is the founder's exposure, the platform dependence is the channel through which that exposure is priced, and the future scenario is the outcome that the exposure and the channel select, so the three questions are the three steps of a single argument, and the argument's conclusion is that the one-person unicorn is a founder's risk that the platform prices before the founder captures it.(3,6,8)

The sequence is the reason the hypothesis is best stated as a condition, the founder captures the unicorn only if the platform's cut and the founder's risk leave enough for the founder, and the condition is one that the existing evidence can state but not yet test. The sequence also explains why the practitioner commentary is optimistic where the evidence is conditional, because the commentary states the founder's capture and omits the platform's cut, and the evidence states both, which is the difference between a hypothesis and its conditions.

The limitations of the review bound the confidence of its answers. The synthesis is narrative rather than systematic, so relevant work may have been missed, and the recommendation is a systematic search once the construct is indexed consistently. The evidence for P1 and P3 is largely inferred from adjacent forms rather than measured in the one-person company, and the

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

recommendation is a dataset that follows the one-person companies that already exist over time. The scenario literature for P3 is prospective by nature, and the review could not test a scenario against an outcome because no outcome exists, so the P3 answer is bounded by the absence of its object. Each limitation is a property of the field rather than of the method, and the answers are offered as the best current map. The map is nonetheless a contribution, because the three conditions have not been drawn together for the one-person unicorn before, and drawing them together shows that the hypothesis is not blocked by a single unknown but by the misalignment of three, the risk that is not yet measured, the dependence that is not yet priced, and the scenario that is not yet closed. The misalignment is the finding, and it is a finding that no single literature could have produced, because each literature sees only its own condition and therefore sees a smaller problem than the founder faces. A founder who reads only one of the three would misjudge the hypothesis, and the review's contribution is to prevent that misjudgment by holding the three in a single frame.

The implication is that the one-person unicorn should be evaluated, by a founder, a funder or a regulator, on the three conditions that the review names rather than on the valuation that the term implies. The founder should ask whether the concentration of risk is one they can bear, the funder should ask whether the platform dependence will transfer the value before the founder captures it, and the regulator should ask which of the two scenarios the platform dependence will select, because the answer to that last question is the one that the hypothesis has not yet supplied. The research agenda that follows is therefore specific, to measure the systemic risk of the one-person companies that already exist, to price the platform's cut against the founder's capture, and to specify in advance the outcome that would count as the first verified one-person unicorn, because the three measures together would convert the hypothesis from a scenario into an estimate. The one-person unicorn, in the meantime, is best described as a founder's bet under a platform's constraint, and the bet is worth taking seriously, not because it is likely, but because it marks the point where the dispersion of the tool and the concentration of the platform meet. The three conditions that this review has named, the risk that the founder bears, the dependence that the platform prices, and the scenario that neither has closed, are the three faces of that meeting, and the founder who understands all three understands the hypothesis, which is more than the hypothesis itself yet promises. The understanding is the practical yield of the review, because a founder who can state the condition under which the unicorn is captured is a founder who can decide whether to attempt it, and the decision, not the prediction, is what the review leaves the founder equipped to make. That decision, made with the three conditions in view, is the point at which the hypothesis becomes a practice, and the practice is the field's next test.

CONCLUSIONS

The one-person unicorn entails a systemic risk that is concentrated in the single founder, and that concentration is the mechanism the form inherits from the systems it depends on. The platform dependence of the one-person company is the condition through which it surrenders the value it captures, and it is the question most likely to decide the hypothesis. The future scenarios range from a dispersed super-individual to a re-concentrated platform economy, and the range remains open until a verified one-person unicorn closes it.

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

REFERENCES

Kshetri N. Artificial Intelligence-Driven One-Person Companies: The Reconfiguration of Entrepreneurship, Work, and Scale. IT Prof. 2026;28(3):19-26. doi: 10.1109/mitp.2026.3694520
Liu X, Yang W. The Emergence of One-Person Companies as Human–AI Socio-Technical Systems: Evidence from AI Ecosystem Density in Chinese Cities. Systems. 2026;14(8):994. doi: 10.3390/systems14080994
Majetić F, Rajter M, Bassetti C. Being a financially vulnerable solopreneur: The health toll. Acta Sociologica. 2026;69(3):493-503. doi: 10.1177/00016993251382718
Kabza M. Artificial intelligence in financial services – benefits and costs. In: Innovation in Financial Services. Routledge; 2020. p. 183-198. doi: 10.4324/9781003051664-14
Zhang C, Li R, Wang Q. Artificial intelligence and sustainable development: A global nonlinear analysis of the moderating roles of human capital and renewable energy. Renewable and Sustainable Energy Reviews. 2026;228:116574. doi: 10.1016/j.rser.2025.116574
Miao Y, Luo J, Jiang Y. Sustaining complementor communities: A resourcing perspective on platform-dependent entrepreneurs. Sustainable Futures. 2025;10:101358. doi: 10.1016/j.sftr.2025.101358
Glavin P, Schieman S. Dependency and Hardship in the Gig Economy: The Mental Health Consequences of Platform Work. Socius: Sociological Research for a Dynamic World. 2022;8:23780231221082414. doi: 10.1177/23780231221082414
El Hajal G, Ivanov S. Machines, minds, and the future of work: GenAI and future work imaginaries. Futures. 2026;184:103934. doi: 10.1016/j.futures.2026.103934
Karaduman M, Arık MA, Karaduman S. Platform Capitalism and Digital Labour: Value Extraction in the Contemporary Digital Media Economy. Journal Media. 2026;7(3):159. doi: 10.3390/journalmedia7030159
Wu DD, Mitchell J, Lambert JH. Global Systemic Risk and Resilience for Novel Coronavirus and COVID-19. Risk Analysis. 2021;41(5):701-704. doi: 10.1111/risa.13746
Jain A, Shadangi SK. Systemic Vulnerabilities of Data Aggregation: Privacy Implications and Legal Safeguards Under India's Digital Data Protection Framework. In: Lecture Notes in Networks and Systems. Springer Nature Switzerland; 2026. p. 359-370. doi: 10.1007/978-3-032-27702-2_32
Siddik AB, Amin NU. Disruptive innovation or systemic resilience? Investigating the impact of artificial intelligence on banking stability. Research in International Business and Finance. 2026;82:103245. doi: 10.1016/j.ribaf.2025.103245
Lobo CA, Pacheco LM. Eurozone's macroprudential policy and financial stability. In: Encyclopedia of Monetary Policy, Financial Markets and Banking. Elsevier; 2025. p. 290-297. doi: 10.1016/b978-0-44-313776-1.00087-8
Ionescu Ș, Chiriță N, Nica I, Delcea C. An Analysis of Residual Financial Contagion in Romania's Banking Market for Mortgage Loans. Sustainability. 2023;15(15):12037. doi: 10.3390/su151512037
Jellouli I. Quantile–Frequency Connectedness Among Artificial Intelligence, FinTech, and Blue Economy Markets. IJFS. 2026;14(2):32. doi: 10.3390/ijfs14020032
Song P, Wu S, Partab J. Artificial intelligence in credit markets: Algorithmic lending, financial inclusion, and regulatory arbitrage. International Review of Economics & Finance. 2026;112:105937. doi: 10.1016/j.iref.2026.105937
YU L, HÄRDLE WK, BORKE L, BENSCHOP T. AN AI APPROACH TO MEASURING FINANCIAL RISK. Singapore Econ Rev. 2023;68(05):1529-1549. doi: 10.1142/s0217590819500668
Hackfort S, Marquis S, Bronson K. Harvesting value: Corporate strategies of data assetization in agriculture and their socio-ecological implications. Big Data & Society. 2024;11(1):20539517241234279. doi: 10.1177/20539517241234279
Banihashemi S. Digital platformization in construction: Micro-entrepreneurial empowerment or concentration of intermediation power?. Journal of Business Venturing Insights. 2026;25:e00624. doi: 10.1016/j.jbvi.2026.e00624
Liu X, Zeng X. Labour market concentration and posted wages: Evidence on wage-setting power from a large online recruitment platform in China. Economic Modelling. 2026;157:107482. doi: 10.1016/j.econmod.2026.107482
Nuccio M, Guerzoni M. Big data: Hell or heaven? Digital platforms and market power in the data-driven economy. Competition & Change. 2019;23(3):312-328. doi: 10.1177/1024529418816525
Fatmawati F, Nurbayani SU, Alisyahbana ANQA, Isma A, Muflih BK. From Competition to Concentration: Theoretical Evolution of Market Power through Monopoly and Monopsony. daengku j hum soc sci innov. 2025;5(4):557-572. doi: 10.35877/454ri.daengku4472
Wagner G, Eifert M, Metzger A, Schweitzer H. Taming the giants: The DMA/DSA package. COLA. 2021;58(Issue 4):987-1028. doi: 10.54648/cola2021065
Smith R, Healey D. Attacking concentration: market power in the digital space. European Competition Journal. 2024;20(2):295-314. doi: 10.1080/17441056.2023.2280328
Colangelo G. In Fairness We (Should Not) Trust: The Duplicity of the EU Competition Policy Mantra in Digital Markets. The Antitrust Bulletin. 2023;68(4):618-640. doi: 10.1177/0003603x231200942
Cantarella M, Strozzi C. Task search and labor supply in the platform economy. Structural Change and Economic Dynamics. 2026;79:550-577. doi: 10.1016/j.strueco.2026.06.007
Giachetti C. Rivalry enablers in ecosystems: The role of dominant operating system platform openness across diverse national contexts. Long Range Planning. 2026;59(2):102620. doi: 10.1016/j.lrp.2026.102620
Uikey VG, Kulkarni G. AI Based Self Preferencing and Consumer Safety in Digital Market. 2025 OITS International Conference on Information Technology (OCIT). 2025:622-626. doi: 10.1109/ocit66168.2025.11400384

https://doi.org/10.62486/978-9915-9966-3-9.Ch13

Basaure A, Vesselkov A, Töyli J. Internet of things (IoT) platform competition: Consumer switching versus provider multihoming. Technovation. 2020;90-91:102101. doi: 10.1016/j.technovation.2019.102101
Wang T. Tech oligarchy and urban mobility in China: the case of DiDi. Science as Culture. 2026:1-17. doi: 10.1080/09505431.2026.2670740
Schulz J, Mayerhoffer DM. Equal chances, unequal outcomes? Network-based evolutionary learning and the industrial dynamics of superstar firms. J Bus Econ. 2021;91(9):1357-1385. doi: 10.1007/s11573-021-01047-8
Binns R, Bietti E. Dissolving privacy, one merger at a time: Competition, data and third party tracking. Computer Law & Security Review. 2020;36:105369. doi: 10.1016/j.clsr.2019.105369
Balasingham B, Jordan H. Big data and competition analysis under Australian competition law: comeback of the structuralist approach?. Journal of Antitrust Enforcement. 2021;9(3):540-565. doi: 10.1093/jaenfo/jnaa055
Lee J. A balanced regulatory model for emerging digital industries: Evidence from Thailand's digital governance. Journal of Asian and African Studies. 2026:00219096261478102. doi: 10.1177/00219096261478102
Söderqvist CF. Platforms, wages and individual bargaining power: assessing labour market competition among Swedish couriers. In: The Elgar Companion to Regulating Platform Work. Edward Elgar Publishing; 2025. p. 152-170. doi: 10.4337/9781035321148.00019
Mieleszczenko-Kowszewicz W, Chodak G, Cambria E, Kazienko P, Cambria E. Transition From the Attention Economy to the Relationship Economy in the Age of Large Language Models. IEEE Intell Syst. 2026;41(4):64-70. doi: 10.1109/mis.2026.3704804
Ebrahimi Shah Abadi A, MobiniDehkordi A. Alternative futures of platform power in Iran's digital economy: a causal layered analysis approach. Eur J Futures Res. 2026;14(1):28. doi: 10.1186/s40309-026-00272-0
Boukhatmi Ä, Van Opstal W. How Can Digital Platforms Resolve Market Failures to Foster a Circular Economy?. Bus Strat Env. 2026;35(1):737-758. doi: 10.1002/bse.70190
Humm P. From marginal tasks to lifecycle thermodynamics: Rethinking labour, value, and distribution in the age of AI. Futures. 2026;184:103928. doi: 10.1016/j.futures.2026.103928
Deipenbrock G. (Generative) Artificial Intelligence in the Third Decade of the 21st Century and Beyond – Selected European Legal, Regulatory and Supervisory Aspects with a Particular View on Companies in Financial Markets. EULR. 2025;36(Issue 2):201-230. doi: 10.54648/eulr2025009
Louisiana State University and University of Pavia, Business Education Complex, Baton Rouge, Los Angeles 70803, United States of America, Piccoli G, Palese B, Northern Illinois University, College of Business, DeKalb, Illinois 60115, United States of America, Rodriguez J, Grenoble Ecole de Management, 12 rue Pierre Sémard, 38000 Grenoble, France. Solopreneur Digital Ecosystems: Genesis, Lineage and Preliminary Categorization. 34th Bled eConference Digital Support from Crisis to Progressive Change: Conference Proceedings. 2021:239-251. doi: 10.18690/978-961-286-485-9.18
Vickers E, Brooks–Gordon B. Growth and diversity of the online solopreneur and sexpreneur. Current Opinion in Psychology. 2023;49:101474. doi: 10.1016/j.copsyc.2022.101474
Jain P. Exploring Solopreneurs' Journey in the Digital Era. In: Advances in Business Information Systems and Analytics. IGI Global; 2024. p. 67-87. doi: 10.4018/979-8-3693-1842-3.ch005
Han X, Jiang M. From laissez-faire to regulatory winter? Regulating Chinese platforms at the crossroads of antitrust, industrial policy, and geopolitics. International Communication Gazette. 2026;88(1-2):68-90. doi: 10.1177/17480485251405588
Kugler M, Aichroth LS. AI is Here to Stay: A Scenario-Based Exploration of Leadership, Followership, and the Impact on Skill Development. New Drctns Student Lead. 2026;2026(191):29-36. doi: 10.1002/yd.70067
Kumar P, Kumra R, Dadwal S. Contrasting the journey of solopreneurs and entrepreneurs in business success - a qualitative inquiry. IJESB. 2025;56(3):351-365. doi: 10.1504/ijesb.2025.149165

https://doi.org/10.62486/978-9915-9966-3-9.Ch13