Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Born Small, Staying Small: Generative Artificial Intelligence and the Organization of New Firms

Vishesh Agarwal

vagarwal3@worldbankgroup.org

World Bank Group

Research Article

Keywords:

Posted Date: October 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11273168/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Born Small, Staying Small: Generative Artificial Intelligence and the Organization of New Firms

Abstract

How does generative artificial intelligence affect the organization of new firms? I study its diffusion by combining its post-2022 emergence with differences in exposure to large language models across UK industries. I find no statistically significant change in firm entry, but a shift in organization at founding: from 2025 onward, a one-standard-deviation increase in exposure was associated with 2.1 percent fewer founding directors and a 14.3 percent decline in the share of firms with three or more directors. These differences persisted six and twelve months after incorporation. The results show that generative AI exposure can reshape how firms are organized without a statistically detectable increase in firm creation.

Generative AI can perform or assist with a broad range of cognitive tasks with experimental and field evidence showing substantial productivity gains in knowledge-intensive work (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2025). These capabilities may also affect how economic activity is organized. Knowledge-based theories of organization emphasize that organizational structure responds to the costs of acquiring, using, and communicating problem-solving knowledge (Garicano, 2000; Garicano and Rossi-Hansberg, 2006), while research on entrepreneurial teams identifies access to complementary resources and capabilities as an important motive for adding team members (Forbes et al., 2006; Lazar et al., 2020). If generative AI allows an entrepreneur to perform tasks or access capabilities that would otherwise be supplied by another person, one response may be to form a firm with fewer people. But adjustment could occur along other margins as well: lower capability costs could change the number of firms that enter, the types of individuals who start firms, or the locations in which new firms are formed.

This paper studies how the organization of new firms changes with exposure to generative AI. I assemble administrative records for approximately 6.3 million private companies incorporated in the United Kingdom between 2018 and August 2026 and reconstruct the number of directors present when each company was formed using Companies House officer histories. I combine these records with the UK Department for Education's industry-level measure of exposure to large language models (Felten, Raj and Seamans, 2023; Department for Education, 2023). I compare changes in founding-team structure across 17 industries with different levels of LLM exposure before and after the widespread public availability of generative AI.

I find that greater LLM exposure is associated with a substantial shift toward leaner founding teams relative to the continuation of pre-existing industry trends. From 2025 onward, a one-standard-deviation increase in LLM exposure is associated with approximately a 2% reduction in the number of directors present at incorporation. The probability that a

firm is formed with a single director increases by 1.77 percentage points, while the probability of forming with three or more directors falls by 0.54 percentage points, or about 14% relative to its pre-2023 mean. The adjustment begins during 2023–24 and becomes larger thereafter. The main organizational response to greater LLM exposure therefore appears in the number of people with which new firms are formed.

The change in organization persists beyond incorporation. Six months after formation, firms in more exposed industries have approximately 2.2% fewer active directors relative to the continuation of their pre-existing trend, and the difference remains approximately 2.1% after twelve months. The probability of retaining three or more active directors is similarly lower at both horizons. By contrast, I find no statistically significant evidence that LLM exposure changes the number of new firms created, the geographic concentration of entrants, or the share of firms founded by individuals with prior entrepreneurial experience. The evidence is therefore concentrated in the organizational structure of entrants rather than in the extensive margin of entry or observable changes in entrant composition. These outcomes describe firms' formal director structures; directors are not equivalent to employees, and I do not interpret the estimates as effects on firm employment.

The identifying variation comes from differences in industry exposure to LLM capabilities over a long period spanning the public diffusion of generative AI. A raw event-study specification reveals that industries with different LLM exposure were already on different founding-team trajectories before 2023, making a conventional parallel-trends difference-in-differences design inappropriate. The preferred specification therefore allows each industry's LLM exposure to generate a linear differential trend over 2018–22 and estimates post-2022 deviations from the continuation of that trend. Pre-period tests provide no evidence of meaningful quadratic departures from the linear exposure-specific trends, and the results are stable to alternative starting dates for the pre-period, alternative definitions of founding directors, restrictions on industry classification, and the exclusion of individual industries in turn. Because treatment varies across only 17 industries, I use a restricted wild-cluster bootstrap with Webb weights for the preferred inference. The estimates should therefore be interpreted as departures from extrapolated exposure-specific pre-trends; the design cannot rule out other post-2022 shocks that are correlated with industry LLM exposure.

The paper is related to a rapidly developing literature on generative AI and entrepreneurship. Marchesi and Tang (2025) study the commercialization of ChatGPT and find greater firm creation in industries in which AI can support early-stage entrepreneurial tasks. Using Chinese firm registrations, Cai et al. (2025) find that areas with greater AI-related human capital experience increased entry of small firms and that new firms are formed with smaller founding teams. Bena, Bian and Giannetti (2026) document increased startup formation in industries more exposed to GenAI, together with smaller entrants, greater geographic dispersion, and changes in founder characteristics. I complement this work by studying the formal organization of firms at incorporation in UK administrative data and following that organizational structure after entry. The long pre-2023 history is also important in this setting because differences in founding-team size across exposure levels were trending before the diffusion of generative AI and must be accounted for explicitly. In the UK setting studied here, I find evidence of organizational compression but no statistically significant effects on firm entry, geography, or prior founder experience.

The paper also contributes to research on knowledge, technology, and the organization of production. In Garicano (2000) and Garicano and Rossi-Hansberg (2006), organiza-

tional structure depends on how knowledge is acquired and allocated across individuals, and changes in the technology for acquiring knowledge can alter the organization of work. A related entrepreneurship literature studies how founding teams are assembled and emphasizes resource seeking and access to complementary capabilities as motives for adding team members (Forbes et al., 2006; Lazar et al., 2020). I add to these literatures by showing that industries more exposed to a technology that expands access to cognitive capabilities experience a shift toward smaller organizations at the point of firm formation, and that this difference remains visible after firms begin operating. The results are consistent with technological access to knowledge and problem-solving capabilities changing the number of people assembled to organize entrepreneurial activity, rather than operating solely through whether entrepreneurship occurs. Research on digital technologies similarly emphasizes complementarities between technology and organizational design, while recent work links AI investment to firm growth and product innovation (Bresnahan, Brynjolfsson and Hitt, 2002; Cockburn, Henderson and Stern, 2018; Babina et al., 2024).

The rest of the paper is organized as follows. Section I describes the UK company-registration setting, the Companies House data, and the measure of LLM exposure. Section II presents the empirical strategy. Section III reports the main results, while Section IV examines the persistence of organizational differences and other margins of adjustment. Section V presents robustness and falsification exercises, and Section VI concludes.

I. Data and Institutional Setting

I combine the statutory register of UK companies maintained by Companies House with its officer appointment histories. The analysis covers private limited companies incorporated between January 1, 2018 and August 31, 2026. I retain companies with a valid incorporation date, a usable primary five-digit Standard Industrial Classification (SIC) code, and a match to the industry-level measure of exposure to large language models described below. Because firms can revise their reported activities, I assign each company its earliest observed primary SIC code. This prevents post-incorporation reclassification from determining treatment status. The resulting sample contains 6,289,111 companies across 17 broad industries.

The principal outcomes describe the formal director structure in place when a company is created. I classify an officer as a founding director when the recorded appointment date is on or before the company's incorporation date. The baseline count includes both natural-person and corporate director appointments; excluding corporate directors makes little difference because only 0.33 percent of companies have one at founding. From these records, I construct the number of founding directors and indicators for having exactly one, at least two, or at least three founding directors. These variables measure the legal governance structure established at incorporation. They do not measure employment, contracting relationships, equity ownership, or informal collaboration, and a director need not be an employee or a cofounder in the economic sense. For the persistence analysis, I count directors active on each company's six- and twelve-month anniversary, defined as officers appointed by that date whose recorded resignation occurs afterward or is not observed. This count includes both founding directors who remain and directors appointed after incorporation; I separately measure post-incorporation additions and founding-director retention. Each horizon is restricted to incorporation cohorts with complete follow-up in the September 2026 officer snapshot.

Table 1—Summary Statistics for Newly Incorporated Companies

Panel A. Founding-team size

	Mean	SD	P25	Median	P75	*N*
Founding directors	1.257	0.581	1	1	1	6,289,111
Human founding directors	1.253	0.576	1	1	1	6,289,111

Panel B. Founding-team composition and measurement

	Mean (%)	SD (pp)	*N*
Single founding director	79.37	40.47	6,289,111
Two or more founding directors	20.63	40.47	6,289,111
Three or more founding directors	3.61	18.65	6,289,111
Corporate founding director	0.33	5.76	6,289,111
Primary SIC changes across snapshots	5.08	21.97	6,289,111

Panel C. Sample composition over time

Period	Incorporation dates	Firms	Share (%)
Pre-2023	Jan. 2018–Dec. 2022	3,425,573	54.5
Transition	Jan. 2023–Dec. 2024	1,595,794	25.4
Post-2024	Jan. 2025–Aug. 2026	1,267,744	20.2

Notes: The sample contains private limited companies incorporated between January 2018 and August 2026 with a valid primary SIC and matched industry LLM exposure. The exposure measure is described in the text and displayed in Figure 1.
Founding directors are directors whose appointment date is on or before the company's incorporation date. Human founding directors exclude corporate directors. Binary variables are reported in percentage points.
The final sample period includes July and August 2026 only; it is therefore a partial third quarter.

Table 1 summarizes the analysis sample. The average new company begins with 1.257 directors, and 79.37 percent begin with a single director. Only 3.61 percent have three or more founding directors. The analogous average based only on natural-person directors is 1.253, confirming that corporate appointments contribute little to the distribution. The sample includes 3,425,573 companies incorporated during 2018–2022, 1,595,794 during 2023–2024, and 1,267,744 between January 2025 and August 2026. Because the final quarter contains only July and August 2026, specifications based on complete quarters are reported as a robustness check.

The treatment variable comes from the UK Department for Education's 2023 assessment of occupational and industry exposure to generative artificial intelligence (Department for Education, 2023). I use the published values in Annex 1, Table 9 [1], which measure the extent to which tasks performed in an industry could potentially be carried out or assisted by large language models. I map each company's earliest five-digit SIC code into the corresponding DfE industry and standardize the 17 industry scores to have an unweighted mean of zero and a standard deviation of one. The measure captures technological exposure rather than realized adoption: a highly exposed industry contains a larger share of tasks that LLMs may affect, but the score does not indicate whether any particular company uses these technologies.

[1]The DfE annex also reports a broader measure of AI occupational exposure. Across the 17 nonaggregate industries used in the analysis, it is almost perfectly correlated with LLM exposure (Pearson's $r = 0.985$; Spearman's $\rho = 0.949$). The implied variance-inflation factor from jointly including the two standardized measures is 33.7. The available industry variation therefore cannot meaningfully distinguish effects specific to LLM capabilities from effects associated with broader AI exposure.

Figure 1 displays the resulting variation. Finance and insurance, information and communication, professional and technical services, and education are among the most exposed industries. Agriculture, transport, motor trades, construction, and accommodation and food services are among the least exposed. The analysis therefore compares changes in the organization of newly registered companies across industries that differed in their ex ante scope for using LLMs.

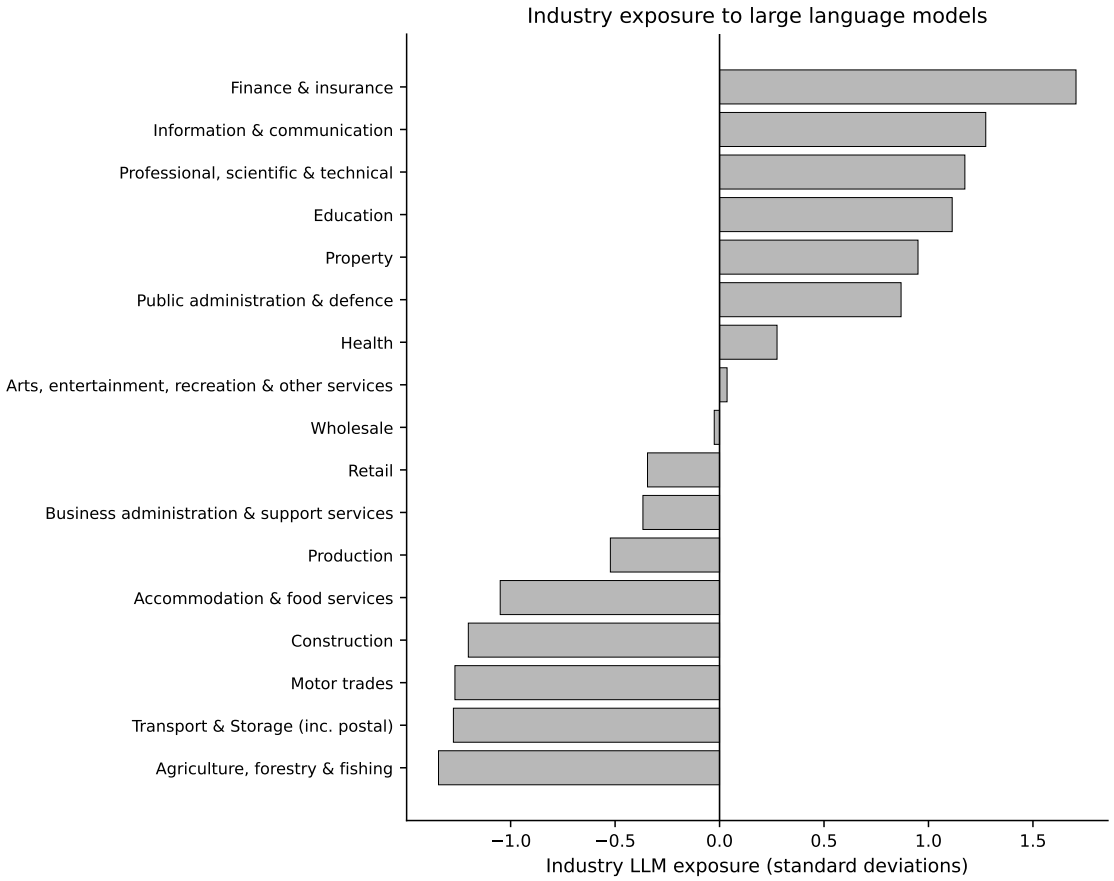


Figure 1. Industry Exposure to Large Language Models

Notes: Source: UK Department for Education (2023), Annex 1 to The Impact of AI on UK Jobs and Training, sheet 'LLM by Industry' (Table 9). The Annex values are used directly and standardized to have mean zero and unit standard deviation across the 17 nonaggregate exposure sectors. Higher values indicate greater exposure to tasks that can be performed or assisted by large language models.

I supplement the founding-team outcomes with several measures of other adjustment margins. To study entry, I aggregate eligible incorporations to the industry-month level and use the logarithm of new-company formations as the outcome. This sample does not require a measured founding-director count, although a live Companies House snapshot may omit some companies dissolved before they were first observed. For geography, I map each company's earliest registered-office postcode to official Office for National Statistics geography and construct an indicator for registration in Greater London; the measure describes the

registered office rather than necessarily the firm's operating location. Finally, I use directors' prior Companies House histories to identify companies with at least one founder who previously founded another incorporated company. This captures prior formal founding experience recorded in Companies House and does not include unincorporated businesses or histories that cannot be linked across records.

II. Empirical Strategy

The main analysis is conducted at the industry-quarter level. For each of the 17 exposure sectors and 35 incorporation quarters between 2018Q1 and 2026Q3, I calculate the mean of each founding-team outcome. This produces 595 sector-quarter cells, which I weight by the number of companies in the cell. The resulting estimates are equivalent to giving each company equal weight while allowing standard errors to be clustered at the level at which exposure varies. For binary outcomes, cell means and coefficients are expressed in percentage points.

A conventional difference-in-differences design would compare changes after the public release of ChatGPT across industries with different levels of LLM exposure. In this setting, however, the identifying assumption of parallel pre-treatment trends is not credible. The raw event study in Appendix Figure A1 shows that founding-team outcomes in industries with different exposure scores were already evolving differently before 2023. I therefore allow the outcome trend to vary linearly with industry exposure and estimate departures from the continuation of that pre-existing relationship.

Specifically, I estimate

$$Y_{sq} = \alpha_s + \lambda_q + \theta\,(E_s \times T_q) + \sum_{k \in \mathcal{P}} \beta_k\,(E_s \times \mathbf{1}(q = k)) + \varepsilon_{sq}, \tag{1}$$

where Y_{sq} is the mean outcome among companies incorporated in sector s and quarter q, E_s is the standardized LLM-exposure score, and T_q is a linear quarterly time index normalized to zero in 2022Q4. Sector fixed effects, α_s, absorb persistent differences in organizational structure across industries, while quarter fixed effects, λ_q, absorb shocks common to all industries. The interaction $E_s \times T_q$ allows industries with greater LLM exposure to follow a different linear trend before the diffusion of generative AI. The set $\mathcal{P}$ contains the quarters from 2023Q1 onward. Because each post-2022 quarter receives a separate exposure interaction, θ is identified by the relationship between exposure and outcomes during 2018Q1–2022Q4. Each β_k measures the subsequent deviation associated with a one-standard-deviation increase in exposure from the extrapolated pre-2023 trend.

The quarterly coefficients describe the timing of the organizational response but are imprecisely estimated individually. I therefore summarize them over two periods. The first is 2023–2024, which captures the initial diffusion of publicly available generative-AI tools. The second runs from 2025Q1 through August 2026 and captures the later period, when access to these tools was more widespread. For each period, I report the equal-weight average of its quarterly coefficients:

$$\overline{\beta}_P = \frac{1}{|\mathcal{Q}P|} \sum k \in \mathcal{Q}_P \beta_k. \tag{2}$$

Equal weighting prevents quarters with unusually large numbers of incorporations from disproportionately determining the period estimate. The final observation, 2026Q3, contains July and August only but enters the later-period average as one quarterly estimate. The appendix reports estimates that exclude this partial quarter.

Identification requires that, absent the diffusion of generative AI, the relationship between industry exposure and the outcomes would have continued along its pre-2023 linear path. This is weaker than assuming identical trends across industries, but it cannot be tested directly. As a diagnostic, I augment the pre-2023 specification with an exposure-by-quadratic-time term. I find no evidence of nonlinear differential pretrends for the principal founding-team outcomes, as reported in Table A2. I also examine alternative pre-period starting dates and placebo breaks within the pre-2023 period. These exercises support the linear approximation but do not rule out discrete post-2022 shocks correlated with exposure, including industry differences in product demand, financing conditions, remote-work practices, or skilled-labour availability. Sector and quarter fixed effects and exposure-specific pretrends do not eliminate such contemporaneous correlated shocks. The estimates should therefore be interpreted as reduced-form changes associated with the DfE LLM-exposure measure during the diffusion of generative AI, rather than as effects of observed firm-level adoption; the design also cannot establish that the changes are specific to LLM capabilities rather than broader AI exposure.

Exposure varies across only 17 industries, making conventional cluster-robust inference potentially unreliable. I report standard errors clustered by exposure sector, but base statistical conclusions on a restricted wild-cluster bootstrap using Webb six-point weights and 1,999 replications. The bootstrap is imposed under the null and applied directly to the relevant quarterly coefficient or period-average linear combination. Figure 2 reports conventional 95 percent clustered confidence intervals to display the quarterly dynamics, while the preferred Webb bootstrap p-values for the pooled periods appear in the tables.

I apply the same trend-adjusted design to the persistence, geography, and founder-experience outcomes. The firm-entry analysis is conducted at the industry-month level using log monthly incorporations, with month fixed effects and an exposure-specific linear trend. Across these exercises, the coefficients retain the same interpretation: the post-2022 deviation associated with a one-standard-deviation increase in industry LLM exposure relative to the continuation of the exposure-related pretrend.

III. Main Results

Figure 2 plots the quarterly deviations from the exposure-specific linear pretrends. The three panels show the effects on the number of founding directors, the probability of a single founding director, and the probability of three or more founding directors. The shaded area denotes the 2023–2024 transition period, the dashed vertical line marks the beginning of 2025, and the open final marker represents the partial 2026Q3 observation covering July and August. The estimates move in the direction of smaller founding teams during 2023–2024 and generally become larger thereafter. In more exposed industries, the number of founding directors and the share of companies with three or more directors decline relative to their pre-existing trends, while the share formed with a single director increases. The pattern develops over several quarters rather than appearing as a single-quarter discontinuity.

Table 2 reports the corresponding period averages. During 2023–2024, a one-standard-

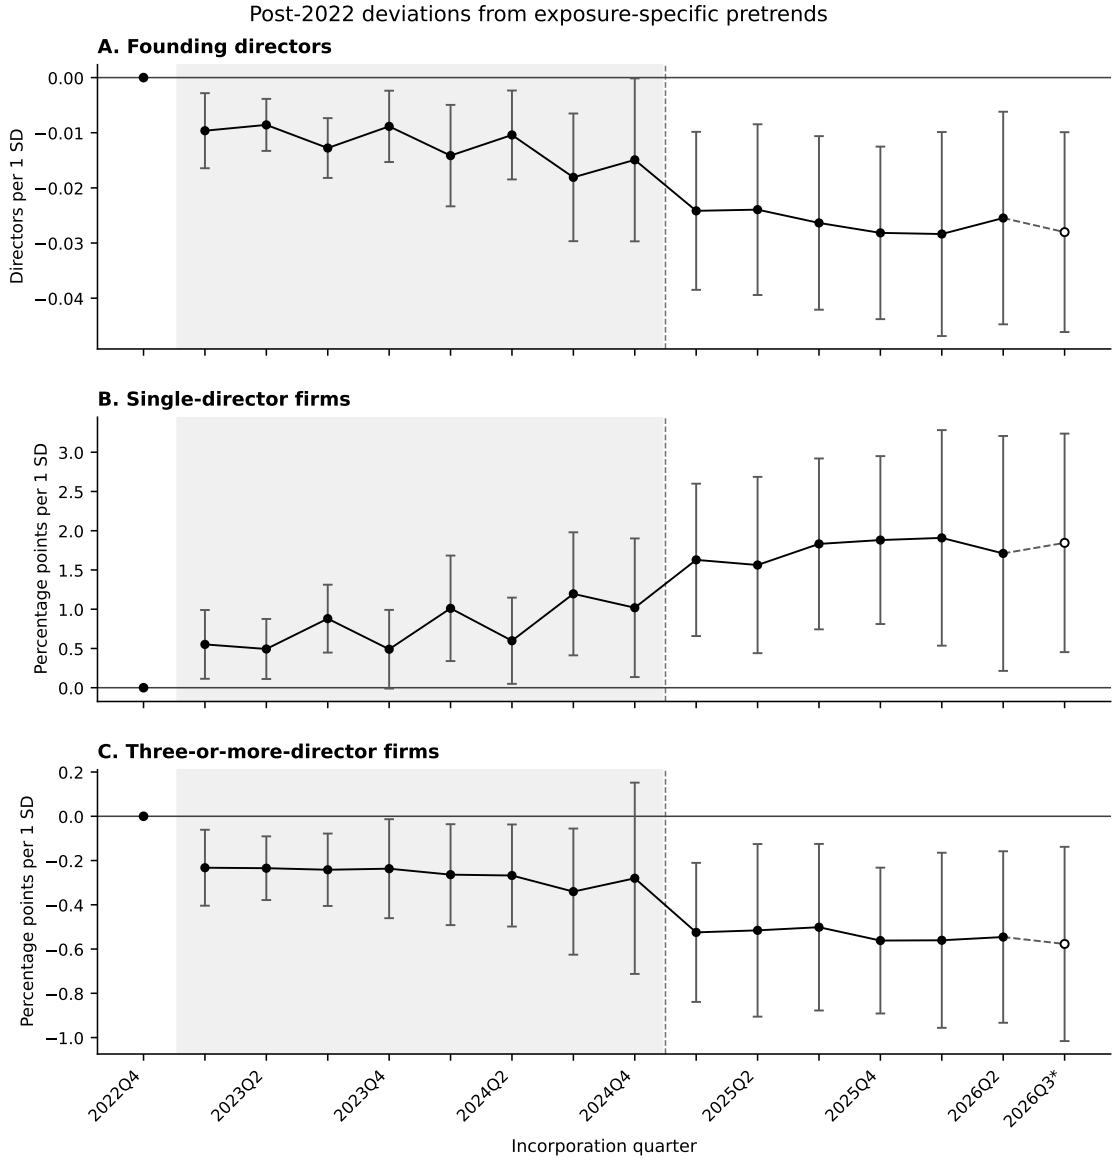


Figure 2. Post-2022 Deviations from Exposure-Specific Pretrends

Notes: The reference quarter is 2022Q4. Each post-2022 point is the estimated deviation associated with a one-standard-deviation increase in LLM exposure from an extrapolated linear exposure-specific pretrend estimated over 2018Q1–2022Q4. The shaded region denotes 2023–24; the dashed vertical line marks 2025Q1. The open 2026Q3 marker covers July and August only. Whiskers show 95 percent confidence intervals using conventional standard errors clustered by the 17 exposure sectors; preferred pooled-period inference is reported using Webb wild-cluster p-values in Table 2.

deviation increase in LLM exposure is associated with 0.012 fewer founding directors. The share of companies formed with a single director increases by 0.780 percentage points, while the share formed with three or more directors declines by 0.262 percentage points. All three estimates are statistically distinguishable from zero using the preferred Webb wild-cluster inference, with p-values of 0.006, 0.004, and 0.044, respectively. The emergence of these differences during the transition period is consistent with organizational adjustment beginning relatively soon after generative-AI tools became publicly available.

Table 2—Generative AI Exposure and Founding-Team Organization

Outcome	Pre-2023 mean	2023–24 effect	Webb p-value	2025–Aug. 2026 effect	Webb p-value
Founding directors	1.265	-0.012 (0.003)	0.006	-0.026 (0.008)	0.030
Relative to pre-2023 mean				-2.1%	
Single founding director	78.91%	+0.780 (0.218)	0.004	+1.767 (0.605)	0.023
Three or more founding directors	3.79%	-0.262 (0.103)	0.044	-0.541 (0.182)	0.048
Relative to pre-2023 mean				-14.3%	

Notes: Each row reports a separate weighted sector-by-quarter regression using 6,289,111 firms in 595 sector-quarter cells. Cells are weighted by their number of firms. The specification includes sector fixed effects, quarter fixed effects, an LLM-exposure-specific linear time trend, and quarter-specific post-2022Q4 deviations from that trend.
Effects are for a one-standard-deviation increase in industry LLM exposure. The founding-director outcome is reported in directors; binary outcomes are reported in percentage points. Cluster-robust standard errors are in parentheses. Preferred inference uses the restricted Webb six-point wild-cluster bootstrap over 17 industry clusters.
The 2023–24 and 2025–August 2026 coefficients are equal-weight averages of the corresponding post-quarter deviations. The 2026Q3 estimate covers July and August only and enters this average as a partial-quarter observation. Relative effects are shown for the two headline outcomes used in the abstract.

The differences are larger from 2025 onward. A one-standard-deviation increase in exposure is associated with 0.026 fewer founding directors, with a Webb p-value of 0.030. Relative to the pre-2023 mean of 1.265 directors, this represents a decline of 2.1 percent. The probability that a company is formed with a single director increases by 1.767 percentage points ($p = 0.023$), equivalent to approximately 2.2 percent of its pre-2023 mean. At the same time, the probability of beginning with three or more directors falls by 0.541 percentage points ($p = 0.048$). Because only 3.79 percent of companies began with three or more directors before 2023, this estimate corresponds to a 14.3 percent reduction in the prevalence of larger founding teams. Put differently, for every 1,000 new companies, a one-standard-deviation increase in industry exposure is associated with approximately 18 additional single-director companies and five fewer companies with three or more founding directors.

The absolute change in the average number of directors is modest because the distribution is already concentrated at one director. The changes in the distribution are more economically revealing: greater LLM exposure is associated both with movement from multi-director to single-director companies and with a particularly large proportional decline in the upper tail of founding-team size. The results therefore indicate organizational compression rather than a uniform reduction of equal importance across all team sizes.

These estimates do not show that any particular company adopted an LLM or replaced a specific prospective director with technology. They instead show that, after 2022, new companies in industries whose tasks are more exposed to LLM capabilities were increasingly formed with fewer formal directors than would be predicted by the continuation of their pre-

2023 trends. This pattern is consistent with generative AI reducing the need to assemble some complementary capabilities within the formal founding team, although other post-2022 shocks correlated with industry exposure cannot be ruled out. The next section examines whether the organizational differences remain after incorporation and whether greater exposure is also associated with changes in firm entry or the observable composition of new firms.

IV. Persistence and Other Margins of Adjustment

The smaller founding teams observed in more LLM-exposed industries could be temporary if companies add directors soon after incorporation. Table 3 examines this possibility by measuring active directors six and twelve months after formation. Each analysis is restricted to incorporation cohorts with complete follow-up in the September 2026 officer records. To ensure comparability, I re-estimate the incorporation outcome within the same fixed-horizon sample before examining the corresponding number of active directors.

Table 3—Persistence of Founding-Team Differences

Outcome	N	Pre-2023 mean	Effect	Relative effect	Webb p-value
Panel A. Number of directors					
At incorporation (6-month cohort)	5,790,823	1.265	-0.026 (0.008)	-2.0%	0.028
Active directors at 6 months	5,790,823	1.283	-0.028 (0.007)	-2.2%	0.028
At incorporation (12-month cohort)	5,410,319	1.265	-0.024 (0.007)	-1.9%	0.030
Active directors at 12 months	5,410,319	1.286	-0.027 (0.007)	-2.1%	0.028
Panel B. Three or more directors					
At incorporation (6-month cohort)	5,790,823	3.79%	-0.525 (0.173)	-13.8%	0.046
Three or more active directors at 6 months	5,790,823	4.20%	-0.540 (0.152)	-12.9%	0.030
At incorporation (12-month cohort)	5,410,319	3.79%	-0.518 (0.171)	-13.7%	0.040
Three or more active directors at 12 months	5,410,319	4.28%	-0.593 (0.159)	-13.9%	0.029

Notes: The six- and twelve-month samples include only incorporation quarters with complete follow-up in the September 2026 officer snapshot. The corresponding birth outcomes are re-estimated within each fixed-horizon sample so that the birth and follow-up estimates use the same firms.
Effects are deviations from the same exposure-specific linear pretrend specification used in the main analysis. Count effects are in directors per one-standard-deviation increase in LLM exposure; binary effects are in percentage points. Cluster-robust standard errors are in parentheses and preferred p-values use the restricted Webb wild-cluster bootstrap.

The organizational differences remain after companies begin operating. In the six-month sample, a one-standard-deviation increase in exposure is associated with 0.026 fewer directors at incorporation and 0.028 fewer active directors six months later. These estimates represent reductions of 2.0 and 2.2 percent relative to their respective pre-2023 means, and both have Webb p-values of 0.028. In the twelve-month sample, the estimated difference is 0.024 fewer directors at incorporation and 0.027 fewer active directors after twelve months. The latter represents a 2.1 percent reduction relative to the pre-2023 mean and has a Webb p-value of 0.028.

The persistence is also visible in the upper tail of the distribution. Among companies with six months of follow-up, greater exposure is associated with a 0.525 percentage-point reduction in the probability of beginning with at least three directors and a 0.540 percentage-point reduction in the probability of having at least three active directors after six months. In the twelve-month sample, the corresponding estimates are reductions of 0.518 and 0.593

percentage points. Relative to their pre-2023 means, these effects range from 12.9 to 13.9 percent, and all are statistically distinguishable from zero using the preferred bootstrap inference.

The close correspondence between the incorporation and follow-up estimates indicates that the main result is not rapidly reversed through subsequent director appointments. Companies in more exposed industries are not only born with smaller formal teams; on average, they remain smaller along this dimension for at least one year. The supplementary decompositions in Appendix Table A9 do not provide statistically precise evidence on whether this persistence reflects differences in the addition of new directors, the departure of founding directors, or both. The data therefore establish persistence in the net organizational difference but do not isolate the particular appointment and exit decisions generating it.

I next examine whether greater exposure is associated with adjustment along other margins. Table 4 reports estimates for firm entry, registered-office location, and founders' prior entrepreneurial experience. These outcomes help distinguish a change in how firms are organized from changes in how many firms are created or which firms enter.

Table 4—Additional Outcomes

Outcome	2023–24 effect	Webb p-value	2025–Aug. 2026 effect	Webb p-value	Quadratic pretrend p
Firm entry	+5.02%	0.214	+10.20%	0.172	0.173
Share registered in Greater London	-1.535 pp	0.079	-2.436 pp	0.134	0.520
Firm has an experienced founder	+0.633 pp	0.660	+0.242 pp	0.907	0.296

Notes: Each row reports a separate trend-adjusted exercise using the same industry LLM exposure. Firm entry is estimated from log monthly formations and is reported as an approximate percentage effect. Greater London and founder-experience outcomes are reported in percentage points. The final period includes the partial 2026Q3 observation covering July and August.

Preferred inference uses restricted Webb wild-cluster bootstrap p-values over the 17 exposure sectors. The final column reports the conventional cluster-robust p-value on an exposure-by-quadratic-time term estimated in the pre-2023 period as a diagnostic for nonlinear pretrends.

None of the reported effects is statistically distinguishable from zero under the preferred Webb inference. The entry and Greater London point estimates are economically meaningful but imprecisely estimated and should not be interpreted as evidence of zero effects.

The firm-entry point estimates are positive but imprecise. A one-standard-deviation increase in exposure is associated with approximately 5.0 percent greater entry during 2023–2024 and 10.2 percent greater entry from 2025 onward. The corresponding Webb p-values are 0.214 and 0.172. These estimates therefore do not provide statistically persuasive evidence that LLM exposure increased firm formation, although their magnitude and uncertainty also preclude interpreting the results as evidence of no entry response.

The geographic estimates are negative, indicating fewer registrations in Greater London in more exposed industries. The estimated difference is 1.535 percentage points during 2023–2024 and 2.436 percentage points from 2025 onward, with Webb p-values of 0.079 and 0.134. The results are suggestive of geographic dispersion but are not sufficiently precise to support a firm conclusion under the preferred inference. Moreover, the outcome records a company's registered office and need not identify where its economic activity is performed.

Finally, I find little indication that exposure changed the prior entrepreneurial experience of founders. The probability that a company has at least one experienced founder increases

by an estimated 0.633 percentage points during 2023–2024 and 0.242 percentage points from 2025 onward. The Webb p-values are 0.660 and 0.907. Alternative measures based on the share of experienced founders and their number of prior foundings yield the same conclusion in Appendix Table A10.

Taken together, the results are concentrated in the internal organization of newly incorporated companies. Greater LLM exposure is associated with smaller formal teams at formation, and those differences remain visible six and twelve months later. By contrast, the available evidence does not establish corresponding changes in the number of companies formed, their geographic concentration, or founders' prior experience. This pattern is consistent with generative AI changing how entrepreneurial activity is organized conditional on entry rather than primarily changing who enters or whether entry occurs.

V. Robustness and Falsification

Table 5 examines whether the main organizational results are sensitive to the measurement of founding teams, the construction of the sample, or the assumed form of the pretrend. The table focuses on the 2025–August 2026 estimates for the number of founding directors and the probability of beginning with three or more directors.

Table 5—Robustness of the 2025–August 2026 Organizational Effects

Specification	N	Founding directors effect	Webb p-value	Three-plus effect (pp)	Webb p-value
Panel A. Measurement and sample definitions					
Baseline	6,289,111	-0.026 (0.008)	0.030	-0.541 (0.182)	0.048
Unique-person founders	6,289,111	-0.026 (0.008)	0.035	-0.532 (0.180)	0.048
Origin-3 founders	6,288,766	-0.025 (0.008)	0.034	-0.500 (0.166)	0.044
Stable SIC only	5,969,355	-0.025 (0.010)	0.042	-0.532 (0.198)	0.052
No corporate founder	6,268,189	-0.026 (0.008)	0.032	-0.526 (0.170)	0.040
Coverage $\geq$ 97%	5,407,562	-0.026 (0.009)	0.033	-0.525 (0.180)	0.046
Coverage $\geq$ 98%	4,675,147	-0.027 (0.009)	0.028	-0.530 (0.184)	0.052
Panel B. Alternative trend specification					
Sector-specific linear trends	6,289,111	-0.027 (0.009)	0.051	-0.501 (0.190)	0.061

Notes: Panel A retains the preferred trend-adjusted specification while varying the founder definition or sample. The baseline uses the validated date-based director count. The unique-person and origin-3 rows use alternative founder measures; the remaining rows impose sample restrictions.

Panel B allows every exposure sector to have its own linear time trend. Effects are equal-weight averages over 2025Q1–2026Q3 for a one-standard-deviation increase in LLM exposure; 2026Q3 covers July and August only. Count effects are in directors; three-plus effects are in percentage points. Cluster-robust standard errors are in parentheses; p-values use the restricted Webb wild-cluster bootstrap.

Under sector-specific trends, the point estimates remain close to the baseline but the Webb p-values are marginal (0.051 and 0.061), so this specification supports magnitude stability more clearly than conventional five-percent significance.

Leave-one-sector-out estimates, placebo breaks, and pre-period-start sensitivity are reported in the appendix rather than the main robustness table.

The measurement and sample checks produce remarkably similar point estimates. Counting unique individuals rather than appointment records, using an alternative officer-origin definition, excluding companies with corporate founders, and restricting the sample to companies with stable SIC codes all yield reductions of between 0.025 and 0.026 founding directors per standard deviation of exposure. The estimated decline in the share with three

or more directors ranges from 0.500 to 0.532 percentage points across these specifications. Restricting the analysis to registration cohorts with at least 97 or 98 percent observed coverage also leaves the estimates essentially unchanged. Some Webb p-values move just above 0.05 as the sample is restricted, but the stability of the point estimates indicates that the results are not driven by a particular founding-director definition, corporate appointments, changing industry codes, or uneven historical coverage.

The preferred specification assumes that pre-2023 trends differ linearly in proportion to the continuous exposure score. A more flexible specification instead allows each of the 17 sectors to have its own linear time trend. Under this specification, the estimated effect on the number of founding directors is -0.027, compared with -0.026 in the baseline, while the effect on the probability of three or more directors is -0.501 percentage points, compared with -0.541 in the baseline. The Webb p-values increase to 0.051 and 0.061, respectively. Thus, the point estimates are robust to substantially more flexible trends, although statistical significance becomes marginal when identification relies on only 17 industry clusters.

The estimates are also not driven by any single industry. Appendix Table A5 excludes each sector in turn. Across all 17 exclusions, the estimated effect on founding directors remains negative and ranges from -0.032 to -0.022. The single-director effect remains positive, ranging from 1.352 to 2.100 percentage points, while the three-or-more effect remains negative, ranging from -0.680 to -0.447 percentage points. Every leave-one-sector-out estimate therefore retains the sign of the baseline result. This is important because the exposure distribution contains several relatively high- and low-exposure sectors that might otherwise exert disproportionate influence.

Several exercises assess the linear-pretrend assumption more directly. In the pre-2023 sample, adding an exposure-by-quadratic-time term yields p-values of 0.774 for the number of founding directors, 0.862 for the single-director indicator, and 0.581 for the three-or-more indicator. These tests do not indicate meaningful nonlinear departures from the exposure-specific linear trend, although their failure to reject should not be interpreted as proof of the identifying assumption. Appendix Table A3 also moves the event date to 2019Q4, 2020Q4, and 2021Q4 while using only pre-2023 observations. None of the nine placebo estimates is statistically distinguishable from zero; the Webb p-values range from 0.406 to 0.856 and do not reproduce the post-2022 pattern of fewer directors, more single-director companies, and fewer companies with three or more directors.

Changing the beginning of the estimation period from 2018Q1 to 2019Q1 or 2020Q1 produces the same qualitative results. The estimated director-count effect changes from -0.026 in the baseline to -0.024 and -0.022, while the single-director effect remains between 1.60 and 1.77 percentage points. The estimate for three or more directors remains negative, although it attenuates from -0.541 to -0.329 percentage points when the pre-period begins in 2020. The results therefore do not depend on the earliest years of the sample, but the upper-tail magnitude is less stable when the available pre-period is shortened substantially.

Finally, Appendix Table A4 excludes July and August 2026 so that the later period contains only complete quarters. The resulting estimates are -0.026 founding directors, a 1.754 percentage-point increase in single-director companies, and a 0.535 percentage-point decline in companies with three or more directors. These are virtually identical to the baseline estimates of -0.026, 1.767, and -0.541, respectively, and remain statistically

distinguishable from zero using the preferred bootstrap inference.

Taken together, these exercises show that the organizational results are stable across alternative definitions, sample restrictions, trend specifications, estimation windows, and sector exclusions, and that comparable effects do not appear at placebo dates before the diffusion of generative AI. The checks strengthen the interpretation of the estimates as post-2022 departures from exposure-related pretrends. They do not, however, eliminate the possibility of another post-2022 industry shock correlated with LLM exposure, which remains the principal limitation of the research design.

VI. Conclusion

This paper examines whether generative AI changes how new firms are organized. I combine Companies House records for approximately 6.3 million UK companies with industry-level exposure to large language models. Because organizational outcomes were already evolving differently across exposure levels before 2023, the analysis estimates post-2022 deviations from exposure-specific linear pretrends.

New companies in more exposed industries increasingly begin with smaller formal founding teams. From 2025 onward, a one-standard-deviation increase in LLM exposure is associated with 2.1 percent fewer founding directors and a 14.3 percent decline in the share of companies formed with three or more directors. These differences remain after six and twelve months and are robust to alternative variable definitions, samples, trends, and sector exclusions. By contrast, I find no statistically persuasive evidence of changes in firm entry, geographic concentration, or founders' prior experience.

The results are consistent with generative AI allowing entrepreneurs to access some capabilities without adding another person to the formal founding team. They do not establish firm-level adoption or effects on employment, and other post-2022 industry shocks correlated with LLM exposure cannot be ruled out. Nevertheless, the findings identify organizational structure as an important margin through which generative AI may reshape entrepreneurship: more exposed firms are born smaller and remain smaller during their first year. Future research should determine whether leaner initial governance lowers coordination costs and facilitates experimentation, how it changes demand for complementary labour, and whether reduced capability diversity affects subsequent scaling.

Appendix

Table A1—Analysis Sample and Variable Definitions

Period	Incorporation dates	Firms	Share (%)
Pre-2023	Jan. 2018–Dec. 2022	3,425,573	54.5
Transition	Jan. 2023–Dec. 2024	1,595,794	25.4
Post-2024	Jan. 2025–Aug. 2026	1,267,744	20.2

The final analysis sample contains private limited companies with a valid primary SIC and matched industry LLM exposure.
Founding directors are director appointments dated on or before incorporation. The baseline includes both natural-person and corporate directors; excluding corporate directors is reported as a robustness check. The final period ends in August 2026, so 2026Q3 is partial.

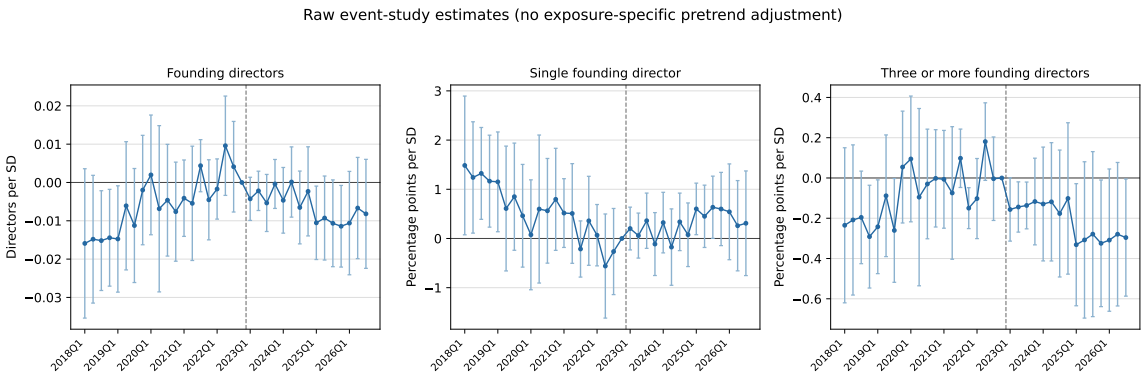


Figure A1. Raw Event-Study Estimates

Notes: Each point is the unadjusted exposure-by-quarter coefficient associated with a one-standard-deviation increase in industry LLM exposure. The specification includes sector and quarter fixed effects but does not include an exposure-specific linear pretrend. The omitted reference quarter is 2022Q4, and the dashed vertical line separates the pre-2023 and post-2022 periods. Whiskers show 95 percent confidence intervals based on standard errors clustered by the 17 exposure sectors.

Table A2—Exposure-Specific Pretrend Diagnostics

Outcome family	Outcome	Start	End	Quadratic p	Pass at 10%
Founding team	Founding directors	2018Q1	2022Q4	0.774	Yes
Founding team	Single founding director	2018Q1	2022Q4	0.862	Yes
Founding team	Two or more founding directors	2018Q1	2022Q4	0.862	Yes
Founding team	Three or more founding directors	2018Q1	2022Q4	0.581	Yes
Firm entry	Firm Entry	2018Q1	2022Q4	0.173	Yes
Official geography	Registered in Greater London	2018Q1	2022Q4	0.520	Yes
Official geography	Registered in a major-city core LAD	2018Q1	2022Q4	0.003	No
Founder experience	Firm has an experienced founder	2018Q1	2022Q4	0.296	Yes
Founder experience	Share Experienced Founders	2018Q1	2022Q4	0.280	Yes
Founder experience	Mean Prior Foundings Per Founder	2018Q1	2022Q4	0.596	Yes

The diagnostic augments the exposure-specific linear pretrend with an exposure-by-time-squared term. 'Pass' means the quadratic term has $p \geq 0.10$; it is a diagnostic, not proof of parallel trends.
The following exploratory outcome fails this diagnostic and should be interpreted cautiously: Registered in a major-city core LAD.

Table A3—Preperiod-Start Sensitivity and Preperiod Placebo Breaks

Test	Start/reference quarter	Estimate	Webb p
Panel: Founding directors			
Preperiod starts	2018Q1	-0.026	0.030
Preperiod starts	2019Q1	-0.024	–
Preperiod starts	2020Q1	-0.022	–
Placebo break	2019Q4	-0.002	0.779
Placebo break	2020Q4	-0.003	0.701
Placebo break	2021Q4	+0.001	0.856
Panel: Single founding director			
Preperiod starts	2018Q1	+1.767	0.023
Preperiod starts	2019Q1	+1.604	–
Preperiod starts	2020Q1	+1.645	–
Placebo break	2019Q4	+0.192	0.663
Placebo break	2020Q4	+0.122	0.816
Placebo break	2021Q4	-0.140	0.713
Panel: Three or more founding directors			
Preperiod starts	2018Q1	-0.541	0.048
Preperiod starts	2019Q1	-0.482	–
Preperiod starts	2020Q1	-0.329	–
Placebo break	2019Q4	+0.036	0.761
Placebo break	2020Q4	-0.112	0.406
Placebo break	2021Q4	-0.039	0.744

Binary-outcome estimates are percentage points per one-standard-deviation higher industry exposure; the director-count estimate is directors per standard deviation.
Preperiod-start rows report the preferred 2025Q1–August 2026 period under alternative estimation starts. Placebo rows move the event into the pre-2023 period.

Table A4—Sensitivity to Excluding Partial 2026Q3

Outcome	Through Aug. 2026	Webb p	Through 2026Q2	Webb p	Difference
Founding directors	-0.026	0.030	-0.026	0.042	+0.000
Single founding director	+1.767	0.023	+1.754	0.032	-0.013
Three or more founding directors	-0.541	0.048	-0.535	0.045	+0.006

The complete-quarter column excludes July–August 2026 and therefore ends in 2026Q2. Both columns use sector and quarter fixed effects, a 2018Q1–2022Q4 exposure-specific linear pretrend, sector-clustered standard errors, and a restricted Webb wild cluster bootstrap.
The complete-quarter bootstrap uses 1,999 draws with seed 20260921.

Table A5—Leave-One-Sector-Out Sensitivity

Outcome	Baseline	Sector trends	Webb p	LOSO min.	LOSO max.	Same sign
Founding directors	-0.026	-0.027	0.051	-0.032	-0.022	17/17
Single founding director	+1.767	+1.879	0.048	+1.352	+2.100	17/17
Three or more founding directors	-0.541	-0.501	0.061	-0.680	-0.447	17/17

The sector-trends column adds separate linear time trends for all industry sectors. Each leave-one-sector-out (LOSO) estimate drops one of the 17 industry clusters. Binary effects are in percentage points; the count effect is in directors.

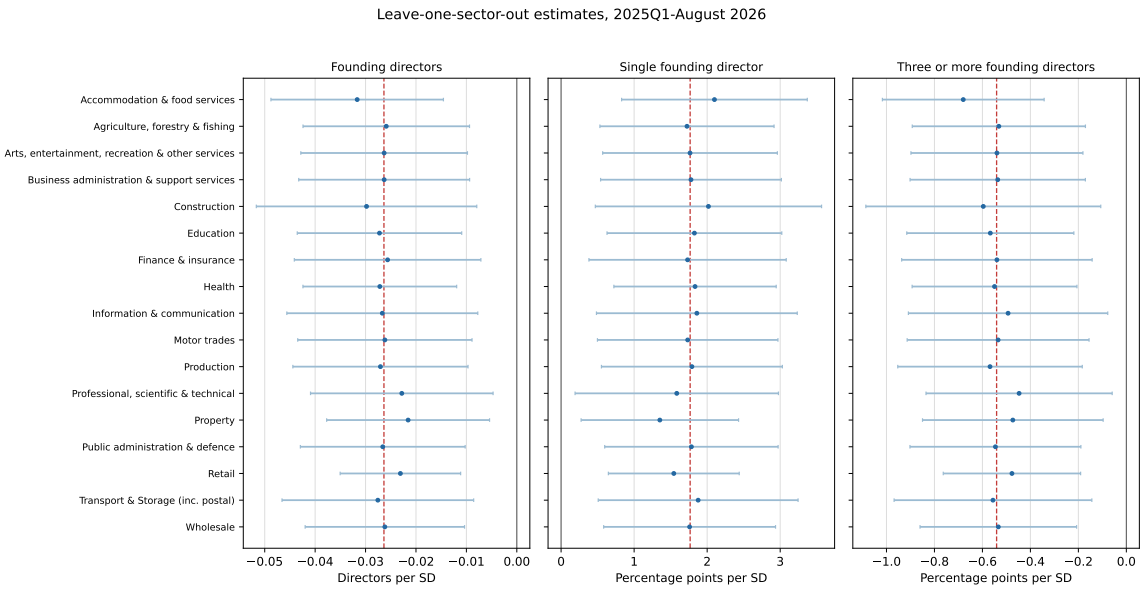


Figure A2. Leave-One-Sector-Out Estimates

Notes: Each point reports the preferred 2025Q1–August 2026 period-average estimate after excluding the named industry sector. The red dashed lines show the corresponding full-sample estimates. Whiskers show 95 percent confidence intervals based on sector-clustered standard errors. The founding-director outcome is measured in directors per one-standard-deviation increase in LLM exposure; binary outcomes are measured in percentage points per standard deviation.

Table A6—Full Measurement and Sample Robustness

Specification	N	Directors	Single	Two or more	Three or more
Baseline	6,289,111	-0.026 [0.030]	+1.767 [0.023]	-1.767 [0.023]	-0.541 [0.048]
Unique-person founders	6,289,111	-0.026 [0.035]	+1.760 [0.023]	-1.760 [0.026]	-0.532 [0.048]
Origin=3 founders	6,288,766	-0.025 [0.034]	+1.724 [0.032]	-1.724 [0.026]	-0.500 [0.044]
Stable SIC	5,969,355	-0.025 [0.042]	+1.625 [0.050]	-1.625 [0.046]	-0.532 [0.052]
No corporate founder	6,268,189	-0.026 [0.032]	+1.738 [0.026]	-1.738 [0.025]	-0.526 [0.040]
Coverage at least 97%	5,407,562	-0.026 [0.033]	+1.793 [0.027]	-1.793 [0.029]	-0.525 [0.046]
Coverage at least 98%	4,675,147	-0.027 [0.028]	+1.862 [0.028]	-1.862 [0.022]	-0.530 [0.052]

Entries are estimates followed by Webb wild-cluster p-values in brackets. The director-count effect is in directors; binary effects are percentage points. The table reports the complete set of measurement and sample variants.

Table A7—Firm-Entry Robustness

Specification	Effect (%)	Log estimate	SE	Webb p
Preferred	+10.20	+0.097	0.067	0.172
Coverage at least 97%	+10.42	+0.099	0.071	0.198
Coverage at least 98%	+10.31	+0.098	0.075	0.229
Stable SIC	+8.61	+0.083	0.068	0.249

The leave-one-sector-out percentage-effect range is +5.88 to +14.72 percent. The dependent variable is log firm formations.

Table A8—Official-Geography Robustness

Sample	Outcome	Estimate (pp)	SE (pp)	Webb p
Preferred Official Match	Registered in Greater London	-2.436	1.375	0.134
Preferred Official Match	Registered in a major-city core LAD	+0.395	0.497	0.532
Stable Postcode	Registered in Greater London	-2.362	1.327	0.144
Stable Postcode	Registered in a major-city core LAD	+0.635	0.462	0.270

Geography uses the official postcode lookup. The Greater London leave-one-sector-out range is -3.383 to -1.727 percentage points. Greater London is the preferred outcome; the major-city core LAD outcome fails the quadratic pretrend diagnostic and is exploratory.

Table A9—Organizational Dynamics at Six and Twelve Months

Months	Outcome	N	Pre-mean	Estimate	Relative (%)	Webb p
6	Founding directors	5,790,823	1.265	-0.026	-2.0	0.028
6	Three or more founding directors	5,790,823	3.792	-0.525	-13.8	0.046
6	Active directors	5,790,823	1.283	-0.028	-2.2	0.028
6	Three or more active directors	5,790,823	4.195	-0.540	-12.9	0.030
6	Directors added	5,790,823	0.087	+0.025	+28.3	0.204
6	Any director added	5,790,823	6.883	+2.647	+38.4	0.135
6	Founder retention share	5,790,823	0.953	-0.028	-2.9	0.127
6	Any founder exit	5,790,823	5.771	+2.696	+46.7	0.159
12	Founding directors	5,410,319	1.265	-0.024	-1.9	0.030
12	Three or more founding directors	5,410,319	3.792	-0.518	-13.7	0.040
12	Active directors	5,410,319	1.286	-0.027	-2.1	0.028
12	Three or more active directors	5,410,319	4.282	-0.593	-13.9	0.029
12	Directors added	5,410,319	0.122	+0.018	+14.9	0.428
12	Any director added	5,410,319	9.206	+2.248	+24.4	0.209
12	Founder retention share	5,410,319	0.937	-0.025	-2.6	0.174
12	Any founder exit	5,410,319	7.867	+2.318	+29.5	0.214

Rows use the six- or twelve-month fixed-horizon sample indicated in the first column. Each sample includes all eligible incorporation quarters with complete follow-up in the September 2026 officer snapshot, and the corresponding incorporation outcomes are re-estimated using the same firms. Reported effects are equal-weight averages of the available post-2024 quarterly coefficients within each horizon.

Count outcomes are reported in directors, indicator outcomes in percentage points, and the founder-retention share as a proportion between zero and one. The relative-effect column expresses each estimate as a percentage of its pre-2023 mean. Webb p-values use the restricted wild-cluster bootstrap over the 17 exposure sectors.

TABLE A10—FOUNDER-EXPERIENCE ROBUSTNESS

Outcome	Preperiod start	Estimate	SE	Webb p	Quadratic p
Firm has an experienced founder	2018Q1	+0.242	1.845	0.907	0.296
Share of founders with prior experience	2018Q1	+0.357	1.779	0.881	0.280
Mean prior companies per founder	2018Q1	-1.753	2.322	0.424	0.596

Rows report equal-weight averages of the quarterly coefficients from 2025Q1 through August 2026 under the preferred trend-adjusted specification. Effects correspond to a one-standard-deviation increase in industry LLM exposure.

The first two outcomes are reported in percentage points. Mean prior companies per founder is reported in its natural unit: the number of previously founded incorporated companies per founder. Standard errors are clustered by the 17 exposure sectors, and preferred inference uses the restricted Webb wild-cluster bootstrap. The final column reports the conventional clustered p-value on the exposure-by-quadratic-time term estimated over the pre-2023 period.

*

Declaration of Competing Interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

*

Data Availability

The underlying company-registration and officer records are publicly available from Companies House, and the industry-exposure measures are publicly available from the UK Department for Education. If the paper is accepted, the analysis data used to produce the reported results, replication programs, a full list of sources, and documentation needed to reproduce the tables and figures will be supplied as supplementary material for publication on the Oxford Economic Papers website. The full underlying Companies House records can be obtained from their original source.

*

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this work, the author used OpenAI ChatGPT to assist with manuscript organization, language refinement, and critical review of the exposition. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

*

REFERENCES

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics*, 151: 103745.

Bena, Jan, Bo Bian, and Mariassunta Giannetti. 2026. "Prompted to Start: How Generative AI Is Transforming Entrepreneurship." SSRN Working Paper.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics*, 117(1): 339–376.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics*, 140(2): 889–942.

Cai, Junhui Jeff, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao, and Wu Zhu. 2025. "AI as "Co-founder": GenAI for Entrepreneurship." *arXiv preprint arXiv:2512.06506*.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern. 2018. "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis." National Bureau of Economic Research NBER Working Paper 24449.

Department for Education. 2023. "The Impact of AI on UK Jobs and Training." UK Department for Education, London.

Felten, Edward W., Manav Raj, and Robert Seamans. 2023. "How Will Language Modelers Like ChatGPT Affect Occupations and Industries?" *SSRN Electronic Journal*.

Forbes, Daniel P., Patricia S. Borchert, Mary E. Zellmer-Bruhn, and Harry J. Sapienza. 2006. "Entrepreneurial Team Formation: An Exploration of New Member Addition." *Entrepreneurship Theory and Practice*, 30(2): 225–248.

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy*, 108(5): 874–904.

Garicano, Luis, and Esteban Rossi-Hansberg. 2006. "Organization and Inequality in a Knowledge Economy." *The Quarterly Journal of Economics*, 121(4): 1383–1435.

Lazar, Moran, Ella Miron-Spektor, Rajshree Agarwal, Miriam Erez, Brent Goldfarb, and Gilad Chen. 2020. "Entrepreneurial Team Formation." *Academy of Management Annals*, 14(1): 29–59.

Marchesi, Edoardo, and Sixun Tang. 2025. "When AI Meets Entrepreneurship: Evidence from the Commercialization of ChatGPT." SSRN Working Paper.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, 381(6654): 187–192.