Board Characteristics and AI Adoption

Ruveyda Nur Gozen* Suleyman Faruk Gozen† Sezer Yasar‡

Preliminary Draft, July 2026

Abstract

We study the relationship between corporate board characteristics and firm-level artificial intelligence (AI) adoption. Linking BoardEx director data to Compustat for U.S. public firms, we find at the board level that firms whose directors have larger external networks, more qualifications, shorter tenure, and more diverse nationalities adopt substantially more AI than their industry and year peers, and that this relationship is strongest in small, young, high-growth, and AI-intensive firms. The same profile reappears at the CEO level: CEOs with larger networks, more qualifications, and shorter tenure run firms that adopt more AI. We then analyze why some boards adopt AI more. A board-network shift-share design shows that firms adopt more AI when their directors sit on other, more AI-intensive boards, which we read as evidence of a director-borne diffusion channel. Lastly, we use board-member deaths as a quasi-random shock to board composition and instrument the board index with a signed death shock; the component of board AI-orientation moved by these deaths maps into higher AI adoption.

Keywords: Corporate Boards, Board and CEO Characteristics, Artificial Intelligence, Technology Adoption, CEO and Director Networks.

JEL: D22, G34, J24, M51, O32, O33.

*Economics, Cardiff Business School & CEP LSE. Email: gozenr@cardiff.ac.uk.
†School of Economics, University of Bristol. Email: suleyman.gozen@bristol.ac.uk.
‡Economics, Cardiff Business School. Email: yasars3@cardiff.ac.uk.

1 Introduction

Artificial intelligence is spreading through the corporate sector rapidly, but also unevenly. The demand for AI-related labor in the United States grew roughly tenfold over the 2010s (Alekseeva et al., 2021), and AI now ranks among the small set of emerging technologies with the breadth of a general-purpose technology (Goldfarb et al., 2023). Adoption is nonetheless highly concentrated. Even within a narrow industry and a single year, some firms staff sizable AI workforces while otherwise comparable firms employ almost none. Because AI adoption is increasingly tied to firm growth, product innovation, and productivity (Babina et al., 2024; Brynjolfsson et al., 2021), the question of what determines which firms adopt is a first-order one for the firms themselves, for the labor market (Acemoglu et al., 2022), and for the distribution of the gains from the technology. The usual determinants, namely size, financial slack, and industry, explain only part of the variation, which directs attention to the organizational and governance margins of the adoption decision.

This paper examines a determinant that sits at the top of the firm and has received little attention in this setting: the board of directors. The board is the firm's apex decision-making body. It approves large and uncertain investments, hires and monitors the CEO, and channels outside information and connections into the firm (Adams et al., 2010). A large corporate-finance literature shows that board composition shapes consequential firm policies, including risk-taking and financial policy (Bernile et al., 2018; Coles et al., 2008), innovation (Balsmeier et al., 2017), and the quality of monitoring and advice (Fich and Shivdasani, 2006; Güner et al., 2008). Adopting a general-purpose, skill-intensive, and uncertain technology such as AI is exactly the kind of strategic, discretionary choice for which the vision, information, and risk tolerance embodied in the board should matter. We therefore ask two questions: how do the characteristics of the board relate to a firm's AI adoption, and through what channel does that relationship operate? To answer them we assemble a firm-year panel of U.S. public firms that links board characteristics from BoardEx to Compustat fundamentals and to a job-postings-based measure of AI adoption from Babina et al. (2024).

Our measurement strategy summarizes each board along seven dimensions: its external network, captured by two measures of the outside boards its directors serve on; its tenure; its education; its national diversity; its size; and its gender mix. Six of the seven move with AI adoption in a common direction, and we combine them into a single standardized AI-oriented board index. The index is formed in three steps: each trait is expressed in cross-sectional standard-deviation units, the six retained traits are

sign-aligned so that a higher value means a more networked, fresher, more credentialed, more diverse, and leaner board, and the sign-aligned traits are averaged and re-standardized.

The baseline result is that the board profile of AI adopters is distinctive and that it maps into adoption by an economically meaningful margin. Relative to their industry-year peers, the directors of AI adopters sit on more outside boards, have shorter tenure, hold more qualifications, and are more nationally diverse. A one-standard-deviation increase in the AI-oriented board index is associated with a 0.094 percentage point higher AI worker share, conditional on industry-year fixed effects and firm controls. The relationship is driven by the network and human-capital margins, namely outside directorships, qualifications, and national diversity, and not by board size or gender. The same outward-looking, credentialed, and short-tenured profile reappears at the level of the single most important director, the CEO: CEOs with larger networks, foreign experience, more qualifications, and shorter tenure run firms that adopt more AI, and a CEO index retains independent explanatory power alongside the board.

The relationship is heterogeneous in a sensible way. It is markedly stronger where adoption is most discretionary and most growth-driven, that is, in small, young, high-sales-growth, high-R&D, and AI-intensive firms. Estimated separately, the board-index slope is about 50 to 70 percent larger in small and young firms than in larger and more mature ones, and it is amplified by firm growth and R&D intensity.

The result is robust. It survives a narrower definition of AI-related postings, alternative board-index constructions, alternative regression specifications including different combinations of fixed effects, and excluding AI-intensive industries. We also apply propensity score matching that balances high- and low-index firms on observables, and the results still survive.

We then test why some boards adopt AI more. We hypothesize that if networked boards adopt more AI because their directors carry knowledge of the technology from board to board, then a firm's adoption could increase with the AI intensity of the other firms its directors are connected to. To test this, we use a shift-share design aimed at the diffusion mechanism. We find that a firm's board-network AI exposure is a robust and strong predictor of its own adoption, which we refer to as a director-borne diffusion channel.

Finally, we exploit the deaths of sitting directors as a quasi-random shock to board composition. Because a death removes a specific director, it shifts the board's measured AI-orientation in a direction that depends only on who died. Using this signed

shock to instrument the board index yields a first stage that is strong and a two-stage least squares estimate with the same sign as, and if anything larger than, the cross-sectional relationship.

The paper makes two main contributions. First, to the corporate-finance literature on boards (Adams et al., 2010; Coles et al., 2008; Balsmeier et al., 2017), we add a new and economically important outcome, AI adoption, and show that the board margins that matter are network reach and human capital rather than the board size or gender composition that much of the literature emphasizes. Second, to the economics of AI and technology diffusion (Babina et al., 2024; Acemoglu et al., 2022; Goldfarb et al., 2023), we provide one of the first systematic portraits of the governance of AI adopters, in contrast to an emerging literature on boards and AI that studies AI as a governance tool or the board's oversight of AI rather than board human capital as a determinant of adoption (Hilb, 2020; Ferreira and Li, 2025; Chatjuthamard et al., 2026), document that the governance correlate concentrates exactly where the adoption choice is most discretionary.

Our findings speak to the firm dynamics behind AI adoption and to the policies aimed at broadening it. Work on the uneven spread of new technology across firms has emphasized size, resources, and industry, and our results add the board as a distinct piece of a firm's organizational capital, since firms whose directors bring wide external networks and deep human capital move into AI faster than otherwise similar peers. Because this board capital is highly persistent, gaps in AI capacity across firms tend to be stable rather than transient, with well-connected, human-capital-rich firms pulling ahead while others lag. For policy, this places part of the AI divide in an organizational margin that input subsidies and infrastructure programs do not reach, and it suggests that broadening adoption may depend on lowering the information frictions that boards help firms overcome.

2 Related Literature

This paper contributes to several strands of the literature. A first literature establishes that AI is an economically consequential, broadly applicable technology whose adoption is uneven across firms. Brynjolfsson et al. (2021) argue that AI is a general-purpose technology whose productivity gains arrive with a lag as firms accumulate complementary intangible capital, and Goldfarb et al. (2023) show, using the breadth of skills demanded in job postings, that machine learning has an unusually general

footprint among emerging technologies. On the labor-market side, Acemoglu et al. (2022) document the firm-level penetration of AI through online vacancies, and Alekseeva et al. (2021) show that demand for AI skills grew roughly tenfold over the 2010s and carries a substantial wage premium. Most directly for our question, Babina et al. (2024) link firm-level AI investment to faster growth and product innovation, and find AI investment concentrated among larger, high-R&D firms, establishing both that adoption matters for firm outcomes and that it varies systematically with firm characteristics. We build directly on this work. We use a job-postings-based measure of firm AI adoption from Babina et al. (2024), and ask what governance characteristics, on top of the firm fundamentals this literature emphasizes, distinguish adopters.

A second, larger literature in corporate finance establishes that board characteristics shape the firm's most consequential choices. Adams et al. (2010) survey the dual board functions of monitoring and advising. On the advisory margin most relevant to a complex technology decision, Coles et al. (2008) show that board structure optimally varies with the firm's advising needs, and Güner et al. (2008) that directors' specific expertise shapes the financing and investment decisions they influence. Board diversity and composition affect risk-taking and corporate policy (Bernile et al., 2018; Adams and Ferreira, 2009); busy or over-committed boards affect monitoring intensity (Fich and Shivdasani, 2006); and, most directly related to our setting, Balsmeier et al. (2017) show that board independence shapes the innovation a firm pursues. A complementary strand emphasizes director networks as conduits of information and influence: board connectedness relates to firm performance (Larcker et al., 2013) and to corporate policies and acquisitions (Fracassi, 2017; Cai and Sevilir, 2012), and external networking shapes internal governance (Fracassi and Tate, 2012). We extend this literature to a new and economically important outcome, AI adoption, and show that the network and human-capital margins, rather than size or gender, are where boards line up with adoption.

A third literature connects individual leaders to firm policy. Bertrand and Schoar (2003) show that managers carry idiosyncratic styles that move firm policies; Kaplan et al. (2012) that measurable CEO characteristics predict performance; and Custódio et al. (2013) that general managerial human capital, the kind acquired across firms and roles, shapes firm choices and pay. Bennedsen et al. (2020), using CEO hospitalization events, show that CEOs causally affect investment and that these effects are larger in younger, growing, and human-capital-intensive firms, precisely the firms in which we find board characteristics to matter most. Our CEO-level results show that the same outward-looking, credentialed, short-tenured profile that characterizes

AI-adopting boards also characterizes their CEOs, and that the two carry distinct information. Methodologically, our identification designs draw on the director sudden-death literature (Nguyen and Nielsen, 2010) and on modern shift-share econometrics (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022).

A small but fast-growing literature connects boards specifically to artificial intelligence, though mostly from a different angle than ours. One strand studies AI as a tool that reshapes governance itself: Hilb (2020) asks how AI may reshape the future of corporate governance, and Ferreira and Li (2025) analyze the use of AI inside the boardroom and its consequences for board decision-making. A second strand studies the board's own role in overseeing the technology: Mulamula and Ruthkamp-Bleom (2026) examine the board's responsibility for governing AI technologies, and Agnese et al. (2025) discuss what the AI era implies for the composition and duties of the board. Closest to us in method, Chatjuthamard et al. (2026) relate firm-level AI to board size. We differ from this literature in both question and design. Rather than asking how AI changes governance, or how boards should oversee AI, we ask how a board's existing human and network capital predicts the firm's decision to adopt AI, and we move beyond board size to a multi-dimensional profile of the board, identified within industry-year cells and probed with a shift-share design and a board-death shock.

3 Data and Measurement

To study the dynamics between board characteristics and AI adoption, we utilize three main data sources.

Compustat. Firm fundamentals come from Compustat, the standard source for accounting and market data on U.S. public firms. For each firm-year we draw total assets, sales, capital expenditure, research-and-development expenditure, common equity, shares outstanding and price (to form Tobin's Q and the market-to-book ratio), and SIC and NAICS industry codes. From these we construct the firm controls used throughout, namely log assets, log sales, and the ratio of capital expenditure to assets, and the moderators (firm size, growth, and profitability) used in the heterogeneity analysis. Continuous variables are winsorized at the 1st and 99th percentiles to limit the influence of outliers. Industry is defined at the 2-digit NAICS level for the fixed effects and, separately, at the Fama and French 12 grouping of SIC codes for the

descriptive industry cuts.

AI Adoption. Our outcome is a measure of a firm-level share of AI workers from Babina et al. (2024). The AI worker share is the percentage of a firm's job postings in a year that are AI-related, that is, postings whose required skills include artificial-intelligence and machine-learning competencies, and it is expressed in percentage points. A narrower definition restricted to core AI skills is used for robustness. Job postings capture the demand for AI labor, and therefore a firm's intent to build AI capacity, particularly well: hiring is the margin on which a firm expands a skill-intensive technology, and postings are observed at high frequency and across the whole cross-section of firms. We note their limitation, namely that postings measure intended hiring rather than realized AI deployment or capital, and return to it in the conclusion.

BoardEx. Board characteristics come from BoardEx, a biographical database of directors and senior executives that records, for each individual, education and qualifications, nationality, gender, current and past board memberships, and tenure, and, for each firm-year, board-level aggregates of these characteristics. We use three layers of the data. The firm-year Overall Board Characteristics record supplies the board-level aggregates that define our seven board variables and the index built from them. The underlying director-level files supply the identity and characteristics of each firm's CEO, for the CEO analysis in Section 6; the full director-by-company membership history, from which we build the director-interlock network used in the shift-share design in Section 7; and director dates of death, used to construct the board-death shock in Section 8.

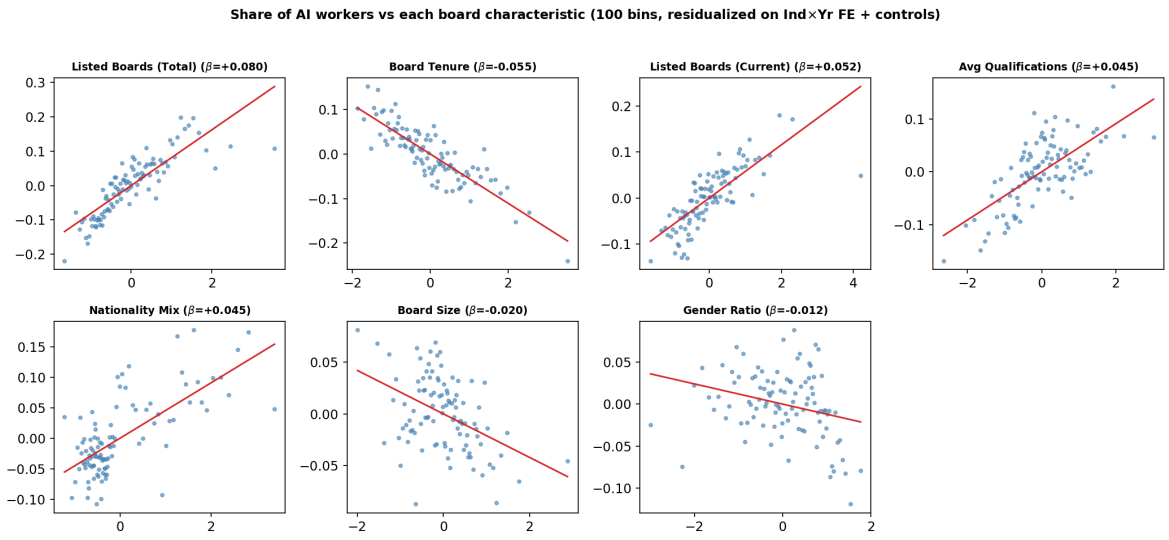
In the final data, we keep one board record per firm-year, require a non-missing AI worker share, a BoardEx board match, and a non-missing industry code and firm controls, and restrict to 2010 through 2018 due to AI worker share data availability.[1] The final data used in the analysis contains 19,363 firm-year observations on 3,149 unique firms.

[1]Table D.1 traces the resulting funnel.

4 Descriptive Evidence

The Board Profile of AI Adopters. We begin with simple descriptive raw correlations. Figure 1 plots the AI worker share against each board characteristic, with both variables residualized on industry-by-year fixed effects and the firm controls and grouped into 100 equal-sized bins. The relationships show clear indications: firms whose directors are more connected, more credentialed, and more internationally diverse post systematically more AI vacancies. Moreover, board tenure and board size are negatively correlated with AI adoption: younger (shorter in tenure) and smaller boards are more likely to adopt AI. Gender ratio is the one trait with no significant connection with AI adoption.

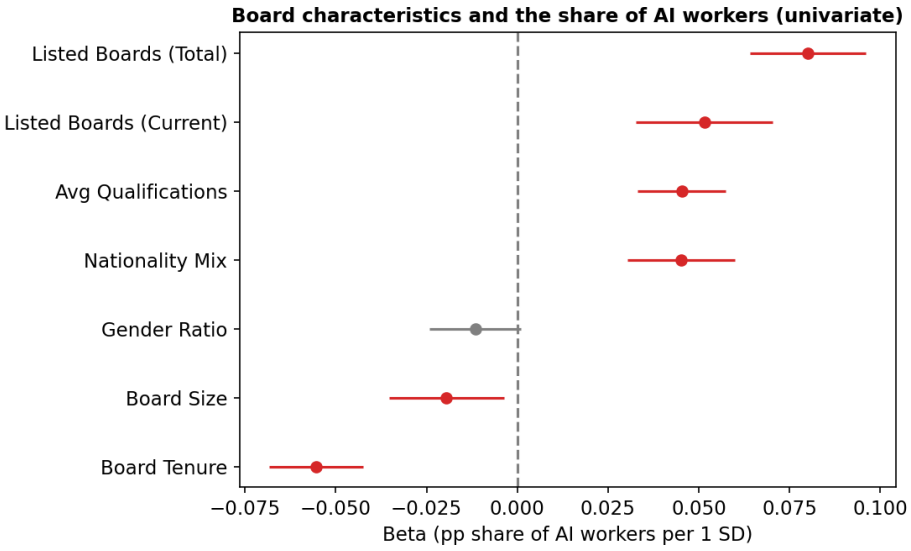
Figure 1: AI Worker Share and Board Characteristics



Notes. Binned scatters of the AI worker share against each board characteristic, both residualized on industry×year fixed effects and firm controls, in 100 equal-sized bins. The univariate slope is reported in each panel title.

We further show the relationship between board characteristics and AI adoption by conditioning on firm level controls (log assets, log sales, capital expenditure over assets), and industry-by-year fixed effects. Figure 2 plots the resulting univariate coefficients. All the traits carry the expected sign as shown with the raw correlation analysis. Better network reach, higher human capital, board members with shorter tenure, and smaller board sizes suggest higher AI adoption.

Figure 2: Board Characteristics and AI Adoption: Univariate Coefficients



Notes. Each point is the coefficient from a separate regression of the AI worker share on one standardized board characteristic (percentage points per standard deviation), with 95 percent confidence intervals, industry×year fixed effects, firm controls, and standard errors clustered by firm.

4.1 Constructing the AI-oriented Board Index

For further econometric analysis, we create an index to summarize the board profile in a single index rather than using each board characteristic independently. Doing so has three advantages: it captures the joint AI-orientation of the board in one interpretable number, it is statistically more powerful than any single trait, and it gives a natural scale, a one-standard-deviation move in overall board orientation, for the economic magnitudes.

The index is built in three steps. First, for each of the seven traits k we form the standardized value

$$z_{k,it} = \frac{x_{k,it} - \bar{x}_k}{\mathrm{sd}(x_k)}, \tag{1}$$

so each trait is expressed in cross-sectional standard-deviation units. Second, we use the six traits that load on AI adoption in a common direction, namely the two network measures, qualifications, nationality mix, and, with a sign flip, tenure and size, and assign each a sign $s_k \in \{+1, -1\}$ aligning "high" with more networked, shorter tenure, more credentialed, more diverse, and smaller: $s_k = +1$ for the two listed-board measures, qualifications, and nationality mix, and $s_k = -1$ for tenure

and board size. We average the sign-aligned standardized traits,

$$\widetilde{B}_{it} \;=\; \frac{1}{6}\sum_{k\in\mathcal{K}} s_k\, z_{k,it}, \qquad \mathcal{K} = \{\text{ListedTot, ListedCur, Quals, NatMix, } -\text{Tenure, } -\text{Size}\}, \tag{2}$$

and third, we re-standardize the average to mean zero and unit variance, $B_{it} = (\widetilde{B}_{it} - \overline{\widetilde{B}})/\text{sd}(\widetilde{B})$. We exclude gender, the one trait without a robust univariate association, and confirm in Section B that adding it, or dropping any single component, leaves the results unchanged.

The resulting index B_{it} is a single number, in standard-deviation units, that is high for a board that has better network, shorter tenure, higher human capital, and are more internationally diverse, and smaller in size. Economically, a high-index board would represent a board positioned to recognize the value of innovative general-purpose technology and support the investment to adopt it. Technically, a one-unit increase in B_{it} is interpreted as moving a firm one standard deviation up this composite orientation, holding its industry and year fixed.[2]

5 Reduced-Form Analysis

Our baseline regression is

$$\text{AI share}_{it} \;=\; \beta\, B_{it} + \gamma' X_{it} + \delta_{j(i)t} + \varepsilon_{it}, \tag{3}$$

where B_{it} is the AI-oriented board index (or a single standardized trait), X_{it} are firm controls (log assets, log sales, capital expenditure over assets), and $\delta_{j(i)t}$ are industry-by-year fixed effects (2-digit NAICS interacted with year). β is identified only from differences across firms in the same industry and the same year.

Table 1 estimates equation (3) with the board index, adding fixed effects step by step. The coefficient is stable and highly significant across specifications. In our preferred column (4), with industry-by-year fixed effects and firm controls, a one-standard-deviation increase in the index is associated with a 0.094 percentage point higher AI worker share. It is worth being concrete about what such an increase is, since the index bundles several traits: moving a firm one standard deviation up the index

[2]To see how tightly the index lines up with AI adoption, Figure E.2 plots the binned scatter of the AI worker share against the index, both residualized on industry×year fixed effects and the firm controls.

means giving it a board that is meaningfully more outward-looking (its directors hold more outside directorships), younger in the boardroom (shorter average tenure), more highly educated (more qualifications per director), and more internationally diverse, all at once. In short, it is a board richer in external information and human capital.

Table 1: Board Index and AI adoption

	(1)	(2)	(3)	(4)
AI-oriented board index	0.116***	0.091***	0.091***	0.094***
	(0.007)	(0.007)	(0.007)	(0.008)
Controls	No	No	No	Yes
Fixed effects	Year	Ind+Yr	Ind×Yr	Ind×Yr
Observations	19,363	19,363	19,360	19,299
Within R^2	0.056	0.039	0.040	0.042

Notes. Dependent variable: firm AI worker share (pp). The AI-oriented board index (6 components, incl. Board Size) is standardized, so coefficients are pp change in AI share per 1 SD. Controls: log assets, log sales, CapEx/assets. SE clustered by firm. *** p<0.01, ** p<0.05, * p<0.10.

Against a sample mean of 0.49 percentage points the estimated effect is economically large, about 19 percent of the mean. In other words, two firms in the same industry and year whose boards differ by one standard deviation along this combined profile, one with the more networked, shorter tenure, more credentialed, and more diverse board, differ on average by about *a fifth* of the typical firm's entire AI workforce share.

5.1 Heterogeneity

We hypothesize that if boards matter because they shape discretionary choices, their association with adoption should be largest where the firm's choices are most malleable and most growth-driven. To do this test, we study three dimensions: firm age and size, growth orientation, and industrial heterogeneity.

We first start with firm size and age heterogeneity. Table 2 interacts the index with indicators for small and young firms. In all specifications, we see that being a small and/or young firm generate a strong positive signal for further AI adoption.

Table 2: Board Index and AI: Size and Age Heterogeneity

	(1)	(2)	(3)	(4)
AI-oriented board index	0.094***	0.079***	0.077***	0.068***
	(0.008)	(0.009)	(0.009)	(0.010)
× Small firm		0.038**		0.029**
		(0.015)		(0.014)
× Young firm			0.039***	0.030**
			(0.013)	(0.013)
Small firm		0.113***		0.110***
		(0.023)		(0.023)
Young firm			0.040**	0.039**
			(0.016)	(0.016)
Controls	Yes	Yes	Yes	Yes
Fixed effects	Ind×Yr	Ind×Yr	Ind×Yr	Ind×Yr
Observations	19,299	19,299	19,299	19,299

Notes. Dependent variable: AI worker share (pp). Small / Young = bottom tercile of assets / firm age within year. Ind×Yr FE, controls, SE clustered by firm.

Table 3: Board Characteristics and AI: Growth Heterogeneity

	(1)	(2)	(3)	(4)
AI-oriented board index	0.083***	0.088***	0.090***	0.047***
	(0.008)	(0.008)	(0.007)	(0.007)
× High sales growth	0.036***			
	(0.009)			
× High asset growth		0.023***		
		(0.008)		
× High market-to-book			-0.006	
			(0.015)	
× High R&D				0.044***
				(0.016)
Fixed effects	Ind×Yr	Ind×Yr	Ind×Yr	Ind×Yr
Observations	19,299	19,299	19,299	19,299

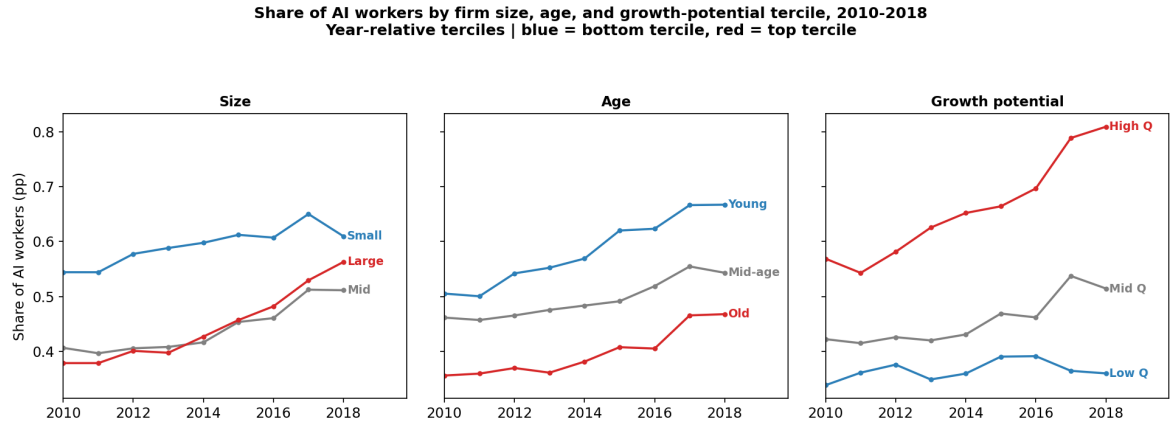
Notes. Dependent variable: AI worker share (pp). High sales / asset growth = top tercile of one-year log growth within year; High market-to-book / R&D = top tercile within year. Each column interacts the board index with one growth proxy. Ind×Yr FE, controls, SE clustered by firm.

We then categorize whether the *growth orientation* of a firm matters. In Table 3, we

interact the index with proxies for growth potential from the data. We find that firms with high sales growth, high asset growth, and high R&D intensity are more likely to adopt AI.

Combining both the age, size, and growth orientation aspects, we further trace the AI worker share by firm size, age, and growth-potential (Tobin's Q) tercile over 2010 to 2018 in Figure 3. Similar to findings above, we find that small, young, and high-growth firms lead throughout, and the gap widens, most sharply for growth potential, where high-Q firms rise from 0.57 to 0.81 percentage points while low-Q firms stay flat near 0.35. The firms in which the board index matters most are the same firms driving aggregate adoption.

Figure 3: AI Worker Share by Firm Size, Age, and Growth Potential



Notes. Average AI worker share by terciles of firm size, firm age, and growth potential (Tobin's Q), 2010 to 2018.

Finally, we study industrial heterogeneity because the relationship could be larger where AI itself is more relevant. We find that there is significant heterogeneity across industries. Telecom, Finance, Business Equipment, Manufacturing, and Healthcare are strongly related with higher AI adoption whereas it is near zero in Utilities, or Energy industry. Therefore, we categorize industries into AI-intensity categorization if the industry average has above-median mean AI worker share.

In Table 4, we show whether it is the board characteristics and/or industry level AI intensity that derives higher AI adoption. We find that board characteristics overall remain highly significant. In Column (2), we interact the index with an AI-intensive-industry indicator and we find that the interaction is positive and significant, suggesting that having a more AI-oriented board index in a more AI-intense industry raises the likelihood of more AI adoption. Column (3) interacts the index instead with a high-technology indicator (Business Equipment, Telecommunications, or Software).

We examine this because AI hiring is concentrated in these sectors, where the technology is most readily deployed and where the complementary engineering talent is found, so it is natural to ask whether the board relationship simply reflects being in a high-tech sector. We find that it does not. The interaction confirms that the link is stronger in high-tech, but the relationship between board characteristics index and AI adoption remains positive and significant regardless of being in a high-tech sector.

Table 4: Board Characteristics and AI: Industrial Heterogeneity

	(1)	(2)	(3)
AI-oriented board index	0.094***	0.073***	0.077***
	(0.008)	(0.008)	(0.006)
× AI-intensive industry		0.041***	
		(0.015)	
× High-tech			0.046*
			(0.027)
Fixed effects	Ind×Yr	Ind×Yr	Ind×Yr
Observations	19,299	19,299	19,299

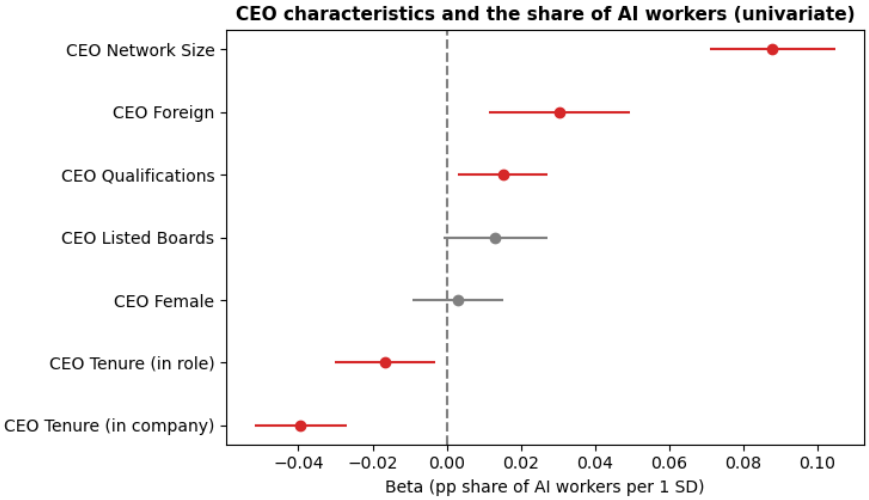
Notes. The dependent variable is the firm's AI worker share, in percentage points. An AI-intensive industry is a 2-digit NAICS industry with an above-median mean AI worker share. Column (1) is the baseline; column (2) interacts the board index with an AI-intensive-industry indicator; column (3) interacts it instead with a high-technology indicator (Business Equipment, Telecommunications, or software firms), the broad sectors in which AI hiring is most common, as a complementary cut. All columns include industry×year fixed effects and firm controls, with standard errors clustered by firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6 CEOs and AI Adoption

The board is a collective body, but one director, the CEO, dominates execution, and a literature on managerial style and CEO effects shows that the individual at the top moves firm policy (Bertrand and Schoar, 2003; Kaplan et al., 2012; Bennedsen et al., 2020). Using BoardEx director records we identify the CEO of each firm-year (a 97 percent match, 18,776 firm-years) and build CEO analogues of the board traits: the CEO's professional network and outside directorships, qualifications, company tenure, and indicators for foreign nationality and female.

In Figure 4, we show the different profiles of CEOs and AI Adoption at the individual level. CEOs with larger networks, foreign experience, more qualifications, and early in tenure run firms with higher AI adoption whereas CEO gender seems to be unrelated.

Figure 4: CEO Characteristics and AI Adoption: Univariate Coefficients



Notes. Each point is the coefficient from a separate regression of the AI worker share on one standardized CEO characteristic (percentage points per standard deviation), with 95 percent confidence intervals, industry×year fixed effects, firm controls, and standard errors clustered by firm.

Table 5: CEO Characteristics and AI: Heterogeneity

	(1)	(2)	(3)	(4)
CEO AI-orientation index	0.071***	0.070***	0.062***	0.065***
	(0.012)	(0.012)	(0.014)	(0.010)
× Small firm	0.013			
	(0.019)			
× Young firm		0.006		
		(0.021)		
× AI-intensive industry			0.027	
			(0.020)	
× High sales growth				0.045***
				(0.014)
Fixed effects	Ind×Yr	Ind×Yr	Ind×Yr	Ind×Yr
Observations	9,973	9,973	9,973	9,973

Notes. Dependent variable: AI worker share (pp). Each column interacts the CEO AI-orientation index with one moderator (defined as in the board-level analysis). Ind×Yr FE, controls, SE clustered by firm. *** p<0.01, ** p<0.05, * p<0.10.

We then perform a similar heterogeneity analysis with firm size, age, and growth orientation. As shown in Table 5, we find that the CEO link is clearly amplified by firm growth, being largest in high-sales-growth and high-R&D firms, but, unlike the

board characteristics index, it does not vary significantly with firm size or age or AI-intensity of the industry.

A CEO AI-orientation Index. Mirroring the board index, we combine the CEO traits, namely network, outside directorships, qualifications, foreign nationality, and, with a sign flip, company tenure, into a single standardized CEO AI-orientation index, built exactly as in equations (1) and (2). We also ask whether board characteristics and CEO characteristics matter differently for AI adoption. Table 6, column (1), shows that a one-standard-deviation increase in the CEO index is strongly and positively associated with higher AI share, conditional on industry-by-year fixed effects and controls. Column (2) includes the board and CEO indices together: we find they both matter for AI adoption in a given firm, with the board index coefficient being significantly higher than the CEO index coefficient.

Table 6: CEO Characteristics Index and AI

	CEO	Both
CEO AI-orientation index	0.076***	0.044***
	(0.010)	(0.012)
Board AI-orientation index		0.066***
		(0.012)
Controls	Yes	Yes
Fixed effects	Ind×Yr	Ind×Yr
Observations	9,973	9,973
Within R^2	0.031	0.046

Notes. The dependent variable is the firm's AI worker share, in percentage points. The sample is firm-years with a CEO identified in BoardEx. The CEO AI-orientation index is the standardized mean of the sign-aligned standardized values of CEO network size, outside listed directorships, qualifications, foreign nationality, and minus company tenure, built exactly as the board index. Column (1) enters the CEO index alone; column (2) is a horse race that adds the board index, in which both indices retain independent and highly significant explanatory power. Both columns include industry×year fixed effects and firm controls, with standard errors clustered by firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7 Why Do Some Boards Adopt AI More?

So far, we studied how board and CEO characteristics are associated with AI adoption. However, this correlation is silent on why such boards adopt more.

One candidate explanation is that networked boards function as conduits of information: a director who observes AI in use on one board may carry that knowledge to another or a director who sits on the board of an AI-adopting firm observes the technology at work, including its uses, its returns, and the talent it requires, and carries that knowledge to the firm's other board members (Fracassi and Tate, 2012; Larcker et al., 2013). This explanation yields a direct empirical implication, distinct from the board's own composition. If diffusion through shared directors is the operative channel, a firm's adoption should respond not to the characteristics of its own board but to the AI intensity of the other firms its directors are connected to.

To test this hypothesis, we first use a board-network shift-share design that isolates the AI a firm is exposed to through its directors. We use the BoardEx record of which directors sit on which boards. For each firm-year, let P_{it} denote the firm's set of interlock peers, the other sample firms on whose boards firm i's directors also sit. Board-network AI exposure is the leave-out mean AI worker share of those peers,

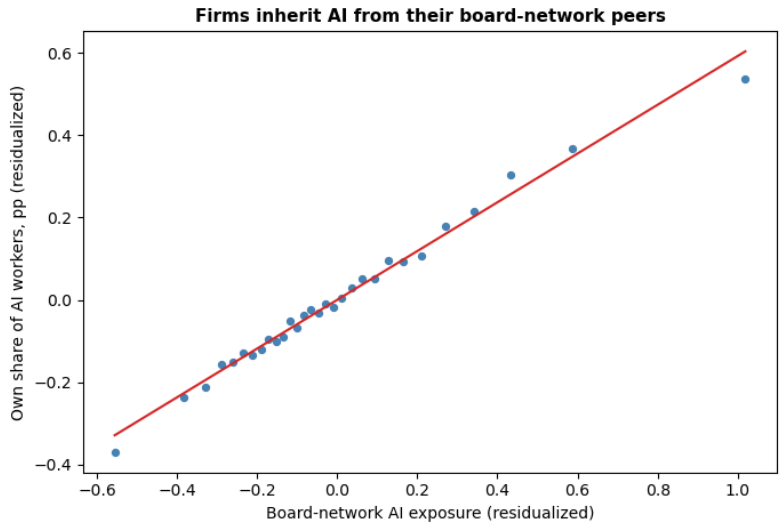
$$\text{Exposure}_{it} = \frac{1}{|P_{it}|} \sum_{j \in P_{it}, j \neq i} \text{AI share}_{jt}, \tag{4}$$

the AI intensity of the firm's director-transmitted information environment. Eighty-eight percent of firm-years have at least one such peer. We then estimate equation (3) with standardized exposure in place of the board index.[3]

In Figure 5, we show that a one-standard-deviation increase in board-network AI exposure is associated with a 0.199 percentage point higher own AI share (column 1), more than twice the board-index coefficient, and the scatter is tight and close to linear. Exposure is largely distinct from the board's own composition: its correlation with the board index is only 0.22, and including the board index (column 5) leaves exposure dominant at 0.120 while the board index falls to 0.066. Much of what board composition proxies may be the information the board's network carries.

[3]Validity tests of the instrument can be found in the Appendix.

Figure 5: AI Worker Share and Board-Network AI Exposure



Notes. Binned scatter of a firm's own AI worker share against its board-network AI exposure (the leave-out mean AI share of the interlock peers reached through its directors), both residualized on industry×year fixed effects and firm controls.

A raw peer correlation is not yet evidence of diffusion. Two threats stand out. The first is reverse causality and common shocks: connected firms are similar, may have selected one another, and may adopt AI together in response to the same forces, so that exposure and own adoption move together without any causal flow from peer to firm. The second is endogenous network formation: a firm expanding into AI may recruit directors precisely from AI-adopting firms, so that the network is chosen in response to the adoption we are trying to explain. Following the shift-share logic of Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022), in which identification can rest either on exogenous shares or on exogenous shocks, the columns of Table 7 address these threats one at a time, progressively isolating variation that the threats cannot explain.

Column (2) uses a predetermined network. We fix the set of interlock peers at its 2010 composition, so that the network (the shares) is determined before the adoption episode we study and cannot be chosen in response to it, which neutralizes endogenous network formation over the window. Column (3) adds a lagged peer signal. On top of the fixed network, we measure peers' AI intensity lagged one year, so that a firm's current adoption cannot mechanically move its own contemporaneous exposure and the identifying variation comes from predetermined peer shocks. Column (4), the key check, retains only peers in a different 2-digit industry from the focal firm. Exposure then comes solely from firms that share no product market with the firm, so it cannot reflect shocks common to the firm's own industry, the most natural remain-

ing confound. That a firm's adoption tracks the AI intensity of out-of-industry firms reached only through its directors is very hard to explain except as director-borne diffusion. Column (5) adds the board index, which leaves exposure dominant, as noted above, so exposure is not merely a restatement of the firm's own board composition.

Overall, we find that exposure to more AI is a robust driver of a firm's own AI adoption.

Table 7: Board-network AI Exposure: Shift-Share Design

	(1)	(2)	(3)	(4)	(5)
Board-network AI exposure	0.199***	0.141***	0.145***	0.127***	0.120***
	(0.011)	(0.011)	(0.012)	(0.012)	(0.012)
Board AI-orientation index					0.066***
					(0.009)
Network	Contemp.	Base-2010	Base-2010	Base-2010	Base-2010
Peer AI timing	t	t	$t-1$	$t-1$	$t-1$
Peers	all	all	all	diff. ind.	diff. ind.
Observations	16,954	13,913	12,276	11,055	11,055
Within R^2	0.172	0.109	0.112	0.096	0.114

Notes. Dependent variable: firm AI worker share (pp). Board-network AI exposure is the leave-out mean AI worker share of the other sample firms on whose boards the firm's directors also sit (director interlocks), standardized. (1) contemporaneous network and peer AI; (2) network fixed at its 2010 composition (predetermined shares); (3) predetermined network with peer AI lagged one year (predetermined shifts), addressing simultaneity; (4) further restricts to peers in a *different* 2-digit industry, so exposure cannot reflect shared own-industry shocks; (5) adds the board index. Ind×Yr FE, controls, SE clustered by firm. *** p<0.01, ** p<0.05, * p<0.10.

8 A Board-member Death Shock

A natural way to ask whether the board's AI-orientation causes adoption is to find variation in board composition that the firm did not choose in response to its own AI plans. The death of a sitting director provides one. It removes a board member on a schedule the firm does not control, and, this is the key to the design, it shifts the board's measured AI-orientation mechanically, in a direction and magnitude that depend only on who died. This is the logic of the director sudden-death literature, which uses unexpected director deaths as exogenous shocks to board composition and value (Nguyen and Nielsen, 2010).

We identify deaths from the date-of-death field in the BoardEx individual-profile file, which records a death date for roughly nineteen thousand directors worldwide.

Matching these records to the directors who sit on our sample firms' boards during the window, using the BoardEx director-by-company membership history that also underlies the network design, yields 836 on-board director deaths at 602 firms. For each death we measure the deceased director's own AI-orientation, o_d, exactly as we measure the board's: the sign-aligned mean of the standardized values of the same traits that define the board index (outside listed directorships, qualifications, network size, foreign nationality, and minus board tenure), taken from that director's last BoardEx record before death.

The economic content of the design is that a death is a removal, so its mechanical effect on the board's average orientation is the negative of the deceased director's orientation, scaled by how much one seat moves a board of a given size:

$$Z_{it} = \sum_{\tau \le t} \sum_{d \in D_{i\tau}} \frac{-o_d}{\text{size}_{i\tau}}, \tag{5}$$

where $D_{i\tau}$ is the set of firm i's directors who die in year τ and $\text{size}_{i\tau}$ is board size. The intuition is concrete. When a long-tenured, locally rooted director (a low-orientation director, $o_d < 0$) dies, the term $-o_d$ is positive: the board mechanically becomes younger (shorter in tenure), and more outward-looking, so its index rises. When a highly networked, credentialed, or foreign director (a high-orientation director, $o_d > 0$) dies, the index falls. Dividing by board size captures that a single departure reshapes a small board more than a large one, and the cumulative sum reflects that the change in composition persists after the death. Z_{it} is thus a signed, predetermined shifter of the firm's board orientation, driven by an event, the timing of a death, that the firm does not choose. We standardize it and use it to instrument the board index in equation (3). Read this way, the exclusion restriction is that a director's death affects the firm's AI hiring only through its effect on the composition of the board.

Table 8 reports ordinary least squares in column (1) alongside the two-stage least squares estimate in column (2), both with industry-by-year fixed effects, firm controls, and standard errors clustered by firm. Three points describe what we find.

Table 8: Board-member Deaths as an Instrument for Board AI-orientation

	OLS	2SLS
Board AI-orientation index	0.094***	0.325***
	(0.008)	(0.089)
First stage: death shock Z		-0.094***
		(0.016)
First-stage F		36
Fixed effects	Ind×Yr	Ind×Yr
Controls	Yes	Yes
Observations	19,299	19,299

Notes. The dependent variable is the firm's AI worker share, in percentage points. The instrument Z_{it} is the firm's cumulative, size-scaled board-death shock, defined in equation (5): each on-board director death enters as minus the deceased director's own AI-orientation divided by board size, summed over the firm's directors and accumulated over time, then standardized. A director's own AI-orientation is the sign-aligned mean of the standardized values of the same traits that define the board index (outside listed directorships, qualifications, network size, foreign nationality, and minus board tenure), measured from that director's last BoardEx record before death; death dates are taken from the BoardEx individual-profile file. Column (1) reports ordinary least squares; column (2) instruments the board index with Z_{it} by two-stage least squares. Both columns use industry×year fixed effects. Standard errors are clustered by firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

First, the instrument is strong. The first-stage F statistic is about 36, comfortably above conventional weak-instrument thresholds, so the accumulated death shock moves the board index with enough force to identify the second stage.

Second, the two-stage least squares coefficient on the board index is $+0.325$. It has the same sign as the ordinary least squares estimate of $+0.094$ and is, if anything, several times larger. The component of board AI-orientation that is moved by the quasi-random removal of directors maps into AI adoption at least as strongly as the raw cross-sectional relationship does. The instrumental-variables evidence therefore reinforces rather than overturns the reduced-form pattern: when a death makes a board more AI-oriented, that firm hires more AI labor.

Yet, the first stage carries a sign that deserves comment. The first-stage coefficient on Z_{it} is -0.094, that is, negative, even though the shock is built so that, mechanically within a firm, an orientation-reducing death lowers the index and an orientation-raising death lifts it. The negative sign is a between-firm phenomenon: firms that accumulate orientation-raising death shocks are disproportionately the low-orientation firms to begin with, and firms that lose high-orientation directors are disproportionately the high-orientation firms, so across firms the cumulative shock and the level of the index line up with opposite signs. Because the specification identifies the co-

efficient from variation across firms rather than from following a single firm through a death, the two-stage least squares estimate should be read as corroborating the cross-sectional relationship under an instrument the firm does not choose, not as a within-firm treatment effect.

Sharpening the design toward a within-firm causal effect would require isolating plausibly sudden, unanticipated deaths, which we discuss as a next step in the conclusion.

9 Conclusion

We study whether the composition of a firm's board bears on its adoption of artificial intelligence. Linking board characteristics from BoardEx to Compustat and to a job-postings measure of AI hiring, we summarize each board by its external network, tenure, education, national diversity, size, and gender, and combine the traits that move together into a single AI-oriented board index. Firms whose directors are more externally connected, shorter in tenure, more credentialed, and more internationally diverse hire substantially more AI talent than their industry and year peers. The gap is economically large. The association is largest in the small, young, high-growth, and AI-intensive firms where the choice to build an AI capability is most discretionary.

The same profile reappears in the person of the CEO. CEOs with larger networks, more qualifications, and shorter tenure run firms that adopt more AI, and the CEO's orientation carries information that the board's does not.

We find that a firm hires more AI when its directors sit on other, more AI-intensive boards, a board-network exposure that holds up when we fix the interlock network at its start-of-sample composition, lag the peer signal, and keep only peers in other industries, which we read as a director-borne diffusion channel. We further test the board index with an instrument built from board-member deaths, a change in board composition the firm does not choose, and find a result that corroborates the cross-sectional relationship.

In future work we aim to sharpen causal identification. First, we plan to verify that the director deaths behind our board-death shock are sudden. Hand-collecting the cause of death, in the manner of the sudden-death literature we build on, would let us restrict attention to plausibly unanticipated transitions and separate them from deaths that the firm could foresee and plan around.

Second, we plan to turn the shift-share logic into a design that holds the focal firm's own board fixed and uses a change on a connected firm instead. Consider a director who sits on both a focal firm and a connected firm. When the connected firm changes its board through a hire or a departure, while the focal firm's board does not change, the connected firm's own AI orientation moves for a reason that is external to the focal firm. In a first stage, this board change at the connected firm predicts that firm's AI adoption; in a second stage, the focal firm's board-network AI exposure, and its own adoption, respond to the predicted AI of the connected firm. Because the focal firm's board is unchanged, this isolates diffusion that runs from the peer to the firm rather than any contemporaneous change in the firm's own governance.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, "Artificial intelligence and jobs: Evidence from online vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

Adams, Renée B and Daniel Ferreira, "Women in the boardroom and their impact on governance and performance," *Journal of Financial Economics*, 2009, *94* (2), 291–309.

_ , Benjamin E Hermalin, and Michael S Weisbach, "The role of boards of directors in corporate governance: A conceptual framework and survey," *Journal of Economic Literature*, 2010, *48* (1), 58–107.

Agnese, Paolo, Francesca Romana Arduino, and Domenico Di Prisco, "The era of artificial intelligence: what implications for the board of directors?," *Corporate Governance: The International Journal of Business in Society*, 2025, *25* (2), 272–287.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska, "The demand for AI skills in the labor market," *Labour Economics*, 2021, *71*, 102002.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial intelligence, firm growth, and product innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Balsmeier, Benjamin, Lee Fleming, and Gustavo Manso, "Independent boards and innovation," *Journal of Financial Economics*, 2017, *123* (3), 536–557.

Bennedsen, Morten, Francisco Pérez-González, and Daniel Wolfenzon, "Do CEOs matter? Evidence from hospitalization events," *The Journal of Finance*, 2020, *75* (4), 1877–1911.

Bernile, Gennaro, Vineet Bhagwat, and Scott Yonker, "Board diversity, firm risk, and corporate policies," *Journal of Financial Economics*, 2018, *127* (3), 588–612.

Bertrand, Marianne and Antoinette Schoar, "Managing with style: The effect of managers on firm policies," *The Quarterly Journal of Economics*, 2003, *118* (4), 1169–1208.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," *The Review of Economic Studies*, 2022, *89* (1), 181–213.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The productivity J-curve: How intangibles complement general purpose technologies," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 333–372.

Cai, Ye and Merih Sevilir, "Board connections and M&A transactions," *Journal of Financial Economics*, 2012, *103* (2), 327–349.

Chatjuthamard, Pattanaporn, Pornsit Jiraporn, and Sang Mook Lee, "Artificial intelligence (AI) and corporate governance: Evidence from board size," *Journal of Behavioral and Experimental Finance*, 2026, p. 101144.

Coles, Jeffrey L, Naveen D Daniel, and Lalitha Naveen, "Boards: does one size fit all?," *Journal of Financial Economics*, 2008, *87* (2), 329–356.

Custódio, Cláudia, Miguel A Ferreira, and Pedro Matos, "Generalists versus specialists: Lifetime work experience and chief executive officer pay," *Journal of Financial Economics*, 2013, *108* (2), 471–492.

Ferreira, Daniel and Jin Li, "Artificial Intelligence in the Boardroom," ECGI Finance Working Paper 1087/2025, European Corporate Governance Institute 8 2025. HKU Jockey Club Enterprise Sustainability Global Research Institute Paper No. 2025/132. Winner of the 2026 John L. Weinberg Center for Corporate Governance / IRRCi Research Paper Competition. Available at SSRN: https://ssrn.com/abstract=5409542.

Fich, Eliezer M and Anil Shivdasani, "Are busy boards effective monitors?," *The Journal of Finance*, 2006, *61* (2), 689–724.

Fracassi, Cesare, "Corporate finance policies and social networks," *Management Science*, 2017, *63* (8), 2420–2438.

_ **and Geoffrey Tate**, "External networking and internal firm governance," *The Journal of Finance*, 2012, *67* (1), 153–194.

Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis, "Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings," *Research Policy*, 2023, *52* (1), 104653.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

Güner, A Burak, Ulrike Malmendier, and Geoffrey Tate, "Financial expertise of directors," *Journal of Financial Economics*, 2008, *88* (2), 323–354.

Hilb, Michael, "Toward artificial governance? The role of artificial intelligence in shaping the future of corporate governance: M. Hilb," *Journal of Management and Governance*, 2020, *24* (4), 851–870.

Kaplan, Steven N, Mark M Klebanov, and Morten Sorensen, "Which CEO characteristics and abilities matter?," *The Journal of Finance*, 2012, *67* (3), 973–1007.

Larcker, David F, Eric C So, and Charles CY Wang, "Boardroom centrality and firm performance," *Journal of Accounting and Economics*, 2013, *55* (2-3), 225–250.

Mulamula, Xitshembhiso Russel and Emma Ruthkamp-Bleom, "The role of the board in artificial intelligence technologies governance," *AI and Ethics*, 2026, *6* (1), 23.

Nguyen, Bang Dang and Kasper Meisner Nielsen, "The value of independent directors: Evidence from sudden deaths," *Journal of Financial Economics*, 2010, *98* (3), 550–567.

Appendices

A Data Structure and Sample Statistics

Table A.1 describes the composition of the sample. The panel is balanced across years, with about 2,000 to 2,260 firm-years per year, each 10 to 12 percent of the sample, over which the mean AI worker share rises steadily from 0.44 percentage points in 2010 to 0.56 percentage points in 2018, a near-30 percent increase over the window (Panel A). Across the twelve Fama and French industries (Panel B), Business Equipment (computers, software, electronics) is both the largest cell (20 percent of observations) and by far the most AI-intensive (mean 0.99 percentage points), followed by Telecommunications (0.57). Wholesale and Retail, Utilities, and Energy adopt least. This concentration motivates our within-industry-year identification and our attention to AI-intensive industries in the heterogeneity analysis.

We use seven board variables. Two capture the board's external network, we use the listed boards Listed Boards (Total) and Listed Boards (Current), the average number of other listed-company boards a firm's directors have ever, and currently, served on. We retain both deliberately, even though they are strongly correlated. The total count captures the cumulative breadth of outside experience and connections a board has accumulated over its directors' careers, while the current count captures the board's live, contemporaneous interlocks, the active conduits through which information flows in the present. The first is a stock of accumulated network capital; the second is the channel through which the diffusion mechanism we study Section 7 would currently operate. One variable captures freshness: Board Tenure, average director tenure in years, where shorter tenure means a more recently refreshed board. One captures human capital: Avg Qualifications, the average number of formal credentials per director. One captures diversity: Nationality Mix. The remaining two are Board Size and Gender Ratio (the fraction of directors who are male). Each variable is standardized to mean zero and unit standard deviation, so that a regression coefficient is the change in AI worker share (in percentage points) associated with a one-standard-deviation change in the characteristic. We construct the composite AI-oriented board index from these traits in Section 5, where it is the central object of the reduced-form analysis.

Table A.1: Sample Composition

Panel A. By year			
Year	Firm-years	% of sample	Mean AI share (pp)
2010	2,000	10.3	0.443
2011	2,053	10.6	0.440
2012	2,175	11.2	0.461
2013	2,262	11.7	0.465
2014	2,240	11.6	0.480
2015	2,205	11.4	0.507
2016	2,142	11.1	0.517
2017	2,138	11.0	0.564
2018	2,148	11.1	0.561
All years	19,363	100.0	0.494
Panel B. By Fama–French 12 industry			
Industry	Firm-years	% of sample	Mean AI share (pp)
Consumer Nondurables	792	4.1	0.310
Consumer Durables	509	2.6	0.480
Manufacturing	1,850	9.6	0.373
Energy (Oil, Gas & Coal)	633	3.3	0.332
Chemicals	478	2.5	0.333
Business Equipment	3,824	19.7	0.993
Telecommunications	436	2.3	0.575
Utilities	559	2.9	0.303
Wholesale & Retail	1,985	10.3	0.252
Healthcare & Pharma	2,361	12.2	0.457
Finance	3,599	18.6	0.359
Other	2,337	12.1	0.394

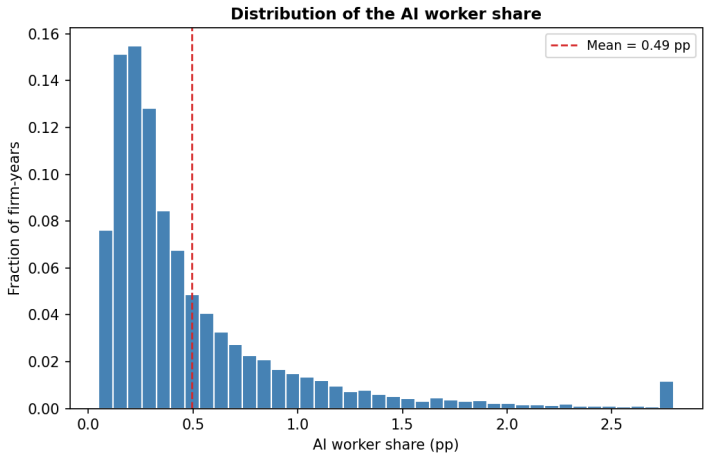
Notes. The estimation sample consists of 19,363 firm-year observations on 3,149 unique firms (3,158 unique BoardEx boards) over 2010–2018. Panel A reports the number of firm-year observations in each year, their share of the full sample, and the mean AI worker share (the percentage of a firm's job postings that are AI-related). Panel B reports the same information by Fama–French 12 industry, with industries assigned from firms' SIC codes. Business Equipment is both the largest and the most AI-intensive industry in the sample.

Table A.2: Summary Statistics

	Mean	SD	P25	Median	P75	N
AI worker share (pp)	0.49	0.49	0.20	0.32	0.59	19363
Narrow-AI share (pp)	0.08	0.22	0.00	0.01	0.05	19363
Listed Boards (Total)	3.06	1.37	2.00	2.90	3.90	19354
Listed Boards (Current)	1.71	0.56	1.30	1.60	2.00	19327
Board Tenure (yrs)	8.11	4.55	4.90	7.90	10.70	19363
Avg Qualifications	2.14	0.51	1.80	2.10	2.40	19363
Nationality Mix	0.10	0.18	0.00	0.00	0.20	18597
Board Size	8.87	2.39	7.00	9.00	10.00	19363
Gender Ratio (frac. male)	0.87	0.11	0.80	0.88	1.00	19363
AI-oriented board index	0.00	1.00	-0.63	-0.02	0.63	19363
Firm age (yrs)	36.41	24.75	19.00	31.00	48.00	19363
Log(Assets)	7.34	1.96	5.93	7.33	8.66	19356
Log(Sales)	6.72	2.03	5.47	6.78	8.07	19353
CapEx/Assets	0.04	0.04	0.01	0.02	0.05	19302
Tobin's Q	2.06	1.48	1.13	1.53	2.36	19310
Market-to-book	3.75	4.85	1.37	2.24	4.00	18529
R&D/Assets	0.05	0.10	0.00	0.00	0.05	19356

Notes. Firm-year observations, 2010–2018, BoardEx boards matched to Compustat with AI worker shares from job-posting data. AI worker share = share_allai $\times$ 100. The AI-oriented board index is the standardized mean of sign-aligned z-scores of Listed Boards (Total & Current), Avg Qualifications, Nationality Mix, $-$Board Tenure, and $-$Board Size. Continuous variables winsorized at 1/99%.

Figure A.1: Distribution of the AI Worker Share



Notes. Distribution of the AI worker share (the percentage of a firm's job postings that are AI-related) across firm-years, 2010 to 2018. The vertical axis is the fraction of firm-years in each bin, and the dashed line marks the sample mean.

Table A.2 reports summary statistics for the outcome, the board characteristics, and the firm controls, and Figure A.1 plots the distribution of the outcome. The mean AI worker share is 0.49 percentage points, but the distribution is strongly right-skewed. A large mass of firm-years post few or no AI-related vacancies, while a long upper tail of intensive adopters pulls up the mean.

B Robustness Checks

Alternative Measures and Specifications. We probe the robustness of the main result along two lines: the specification, and the construction of the index. Table B.1 re-estimates the baseline (Table 1, column 4) under a series of perturbations, taken one at a time.

Table B.1: Robustness Checks: Alternative Specifications

	Narrow AI	Ind+Yr	NAICS3×Yr	ex-Finance	ex-BusEq	Clu. ind
AI-oriented board index	0.025***	0.093***	0.070***	0.080***	0.078***	0.094***
	(0.003)	(0.007)	(0.008)	(0.009)	(0.006)	(0.012)
Observations	19,299	19,302	19,256	15,745	15,481	19,299

Notes. Each column re-estimates the baseline (Table 1, col. 4). (1) narrow-AI outcome; (2) separate industry and year FE; (3) 3-digit NAICS×year FE; (4) excluding Finance; (5) excluding Business Equipment; (6) SE clustered by industry. Controls throughout.

We first ask whether the result depends on how broadly AI-related postings are defined. Replacing the outcome with the narrow-AI share, which counts only postings requiring core AI and machine-learning skills and excludes adjacent data and analytics roles, leaves the index positive and significant at 0.026, so the relationship is not an artifact of a loose AI definition. Second, our baseline interacts industry and year; to check that the result does not hinge on this demanding specification, we instead include industry and year fixed effects separately, which absorbs less variation but is more transparent, and the coefficient is unchanged at 0.094. Third, a concern is that 2-digit industries are too coarse, so that the estimate compares firms in genuinely different product markets; tightening the fixed effects to 3-digit NAICS×year, and so comparing firms within much narrower industry-years, gives 0.070, somewhat smaller but still highly significant. Fourth, financial firms have distinctive boards (regulatory seats, different disclosure) and distinctive technology needs; dropping them entirely gives 0.079, so the relationship is not driven by the financial sector. Fifth,

because adoption is concentrated in Business Equipment, one might worry the result is a within-tech phenomenon, but dropping Business Equipment, the highest-AI sector, gives 0.082, confirming that the relationship holds among non-technology firms. Finally, we guard against correlated errors within industries by clustering standard errors by 2-digit industry rather than by firm, and significance is unchanged at 0.094.

A separate concern is that the result depends on the particular way we combine the traits into the index. Table B.2 re-estimates the baseline with six alternative constructions: the equal-weighted six-trait index; dropping board size; dropping the count of currently held listed directorships; adding gender; weighting the traits by their first principal component rather than equally; and weighting them by their fitted relationship with AI. The six versions correlate 0.87 to 0.96 with one another, and the estimated coefficient is essentially unchanged, ranging from 0.090 to 0.099 and always highly significant. The finding is a property of the underlying traits, not of a particular recipe.

Table B.2: Robustness Checks: Alternative Board-index Construction

	(1)	(2)	(3)	(4)	(5)	(6)
Board index (standardized)	0.094***	0.094***	0.095***	0.091***	0.090***	0.099***
	(0.008)	(0.008)	(0.007)	(0.007)	(0.009)	(0.007)
Within R^2	0.042	0.040	0.044	0.039	0.035	0.048
Observations	19,299	19,299	19,299	19,299	19,299	18,506

Notes. Dependent variable: AI worker share (pp). Each column uses a different construction of the board index, each standardized to unit SD. (1) baseline equal-weighted 6 components; (2) excluding Board Size; (3) excluding Listed Boards (Current); (4) adding Gender Ratio; (5) first principal component of the 6 components; (6) weights from an FE regression of AI on the components. Ind×Yr FE, controls, SE clustered by firm.

Propensity Score Matching. Our regressions condition on industry×year fixed effects and a set of firm controls, but one might still worry that firms with AI-oriented boards differ systematically from other firms on observable dimensions, in that they may be larger, faster-growing, or more research-intensive, and that it is these differences, rather than board composition itself, that explain their AI adoption. A linear control for such covariates assumes a functional form and extrapolates across regions of the covariate space where the two groups barely overlap. Propensity-score methods address both concerns. They reweight or match observations so that treated and control firms have the same distribution of observables, and they restrict attention to the region of common support where comparable firms of both types actually exist,

making no functional-form assumption about how the covariates enter.

We implement this in three steps. First, because the board index is continuous, we define a binary treatment equal to one for firms with an above-within-year-median AI-oriented board index, so that treated firms are those with relatively AI-oriented boards in a given year. Second, we estimate each firm-year's propensity score, its predicted probability of treatment, from a logit of the treatment indicator on the firm's log assets, log sales, capital intensity, age, Tobin's Q, R&D intensity, and year indicators, and trim the sample to the common-support region in which both treated and control observations are present. Third, we use the propensity score in two complementary ways. We form stabilized inverse-propensity weights, which upweight firms that are under-represented in their treatment group so that the reweighted treated and control samples are balanced on the covariates, and re-estimate our baseline specification with the continuous index on this reweighted sample; and we separately compute a nearest-neighbour matched estimate, pairing each treated firm-year with the control firm-year closest in propensity score and averaging the within-pair difference in AI adoption to obtain the average treatment effect on the treated.

Table B.3: Board Characteristics and AI: Propensity Score Matching

	Board index (continuous)	PSM ATT (high vs. low)
Baseline (unweighted)	0.093***	
	(0.008)	
Inverse-propensity reweighted	0.062***	
	(0.008)	
Nearest-neighbour PSM ATT		0.048***
		(0.007)
Observations (common support)	18,932	9,381 pairs

Notes. "Treatment" is an indicator for an above-within-year-median AI-oriented board index. The propensity score is the fitted probability of treatment from a logit of the treatment indicator on log assets, log sales, CapEx/assets, firm age, Tobin's Q, R&D/assets, and year indicators; observations are trimmed to the common-support region $[0.05, 0.95]$. Column 1 reports the coefficient on the *continuous* board index, unweighted and then weighted by stabilized inverse-propensity weights, in the baseline specification (equation 3); column 2 reports the average treatment effect on the treated from nearest-neighbour matching on the propensity score, in percentage points. Standard errors clustered by firm (columns) and from matched-pair variation (ATT). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.3 reports the results. On the inverse-propensity-reweighted sample the coefficient on the continuous board index is 0.062 ($p < 0.01$), modestly below the unweighted 0.093, as one would expect once the groups are balanced on growth and

size, but firmly positive and highly significant. The nearest-neighbour matched comparison tells the same story: firms with AI-oriented boards adopt 0.048 percentage points more AI than otherwise-similar firms with less AI-oriented boards ($p < 0.01$). Balancing the two groups on observables thus attenuates the relationship only slightly and leaves it strongly significant. The board-AI link is not an artifact of which firms happen to have AI-oriented boards; it survives a design that compares only firms made comparable on their observable characteristics.

Firm Fixed Effects. A natural concern is timing and within-firm variation such that a firm that is already hiring AI talent might subsequently reshape its board, so that the contemporaneous correlation reflects boards responding to AI rather than the reverse.

To address this, and to ask whether the relationship survives a within-firm comparison, we re-estimate the baseline using the board index measured in the prior year ($t-1$) and progressively tighten the fixed effects, culminating in a specification that adds firm fixed effects so that the coefficient is identified purely from within-firm variation over time. The regressor is thus predetermined relative to the AI outcome, and all columns condition on contemporaneous firm controls.

Table B.4: Lagged Board Index: Between-firm vs. Within-firm

	(1)	(2)	(3)	(4)	(5)
Lagged board index ($t-1$)	0.123***	0.121***	0.095***	0.094***	0.014*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Controls (t)	Yes	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	–	–
Industry FE	No	No	Yes	–	–
Industry × Year FE	No	No	No	Yes	Yes
Firm FE	No	No	No	No	Yes
Within R^2	-	0.075	0.044	0.042	0.002
Observations	15,744	15,744	15,744	15,740	15,464

Notes. Dependent variable: firm AI worker share (pp). The regressor is the six-component AI-oriented board index measured in the prior year ($t-1$), standardized to unit SD, so the coefficient is the pp change in AI share per 1 SD. All columns include contemporaneous (t) firm controls (log assets, log sales, capx/assets). Columns (1)–(4) progressively add year, industry, and industry×year fixed effects but identify the coefficient across firms; column (5) adds firm fixed effects, identifying only within firms over time. Standard errors clustered by firm in parentheses. The predetermined board index strongly predicts AI hiring across firms (≈ 0.09) but not within firms, consistent with the between-firm (firm-type) interpretation in the main text. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.4 reports the results. Across firms, the lagged board index is a strong predictor

of AI hiring: with controls and no fixed effects the coefficient is 0.123, and it remains 0.094 once we absorb industry×year fixed effects (columns 1 through 4, all $p < 0.01$), essentially identical to the contemporaneous estimate, so the relationship is not an artifact of boards reacting to AI within the same year. Column (5) then adds firm fixed effects, identifying the coefficient purely from within-firm variation over time. The estimate falls to 0.014. We therefore read the reduced-form relationship as a between-firm, firm-type pattern rather than as a within-firm dynamic.

C Validity Test of the Shift-Share Instrument

A shift-share design is only as credible as the diagnostics it survives, and the recent econometric literature (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022) is explicit about what those diagnostics should be: the identifying variation should be unrelated to predetermined characteristics (balance), it should not be an artifact of a few influential peers or of measurement noise, and it should be robust to controlling for the dimensions on which balance is imperfect. We therefore subject the predetermined, out-of-industry exposure to these checks in Table C.1, motivated in each case by a specific threat to validity.

Table C.1: Shift-Share Validity Checks for Board-network AI Exposure

	(1)	(2)	(3)	(4)	(5)	(6)
Board-network AI exposure	0.127***	0.252***	0.290***	0.127***	0.087***	0.303***
	(0.012)	(0.024)	(0.035)	(0.012)	(0.011)	(0.026)
Restriction	–	≥3 peers	≥5 peers	+peer size	+Q, R&D	2SLS
First-stage F						417
Observations	11,055	6,499	3,746	11,055	11,032	10,749

Notes. Dependent variable: AI worker share (pp). All use the predetermined (2010), one-year-lagged, out-of-industry exposure (Table 7 col. 4) unless noted. (2)–(3) require a minimum number of interlock peers; (4) controls for the mean size of peer firms; (5) controls for the firm's own Tobin's Q and R&D intensity (growth); (6) is 2SLS instrumenting contemporaneous exposure with the predetermined out-of-industry network (first-stage F reported). Ind×Yr FE, controls, SE clustered by firm. *** p<0.01, ** p<0.05.

We begin with balance on size. If exposure were as good as randomly assigned, it should not line up with predetermined firm characteristics; regressing it on firm size yields 0.005 ($p = 0.72$), so exposure is unrelated to size. It does tilt toward higher-Tobin's-Q firms, which we address directly below. We next ask whether firms have enough peers. Where a firm has few interlock peers, the leave-out mean is

measured with noise, attenuating the estimate; forty-one percent of firm-years have fewer than three peers, and requiring at least three, then five, strengthens the estimate to 0.25 and 0.29, exactly the pattern expected if classical measurement error in low-peer firms attenuates the baseline. We then check that the result is not just large peers: a firm connected to large, visible firms might adopt for unrelated reasons, but controlling for the average size of the peer firms leaves the estimate unchanged at 0.127. Because exposure was not perfectly balanced on Tobin's Q and R&D, we also confirm that the result is not just growth firms by controlling for the firm's own values of both; the estimate falls modestly to 0.087 but remains positive and significant, so the channel is not merely that growth firms cluster together. Finally, we instrument contemporaneous exposure with the predetermined, out-of-industry network in two-stage least squares. The instrument is very strong (first-stage $F = 417$, far above any weak-instrument threshold) and the estimate is 0.30.

The exposure relationship thus passes the checks that ordinarily overturn spurious peer effects.

D Additional Tables

Table D.1: Sample Construction

Step	Firm-years	Unique firms
Compustat firm-years, 2010–2018	83,719	14,311
with a non-missing AI worker share	50,172	6,997
matched to a BoardEx board record	19,822	3,218
Final regression sample	19,363	3,149

Notes. The table traces the construction of the estimation sample. The first row reports the universe of Compustat firm-years over 2010–2018 (one observation per gvkey-fiscal year). We then keep firm-years for which an AI worker share (the AI-related share of job postings) can be computed and is non-missing, require an inner match to a BoardEx "Overall Board Characteristics" record. The final regression sample additionally requires a non-missing industry code and the firm controls used in equation (3).

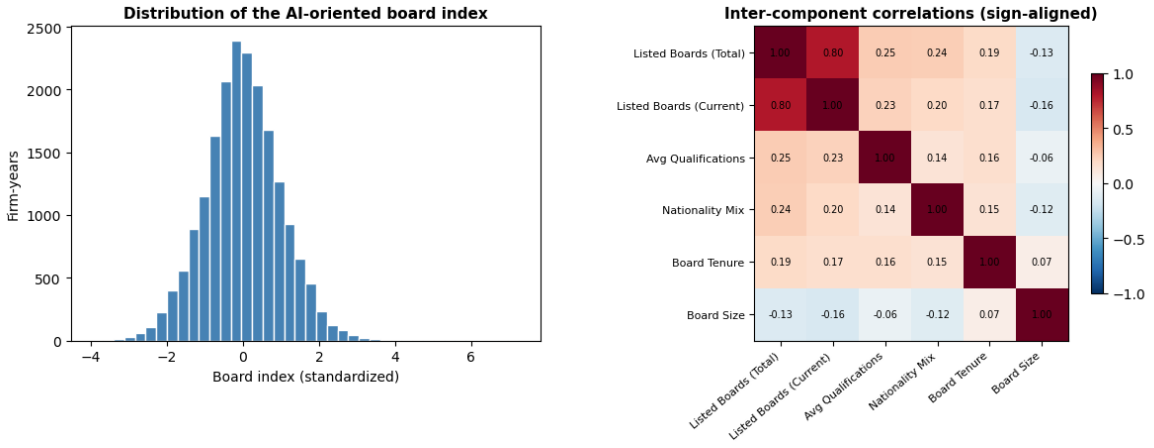
Table D.2: The AI-oriented Board Index and Its Components

Component (sign-aligned)	Mean	SD	Corr. with index	Corr. with other comps.
Listed Boards (Total)	3.06	1.37	0.74	0.53
Listed Boards (Current)	1.71	0.56	0.70	0.48
Avg Qualifications	2.14	0.51	0.53	0.26
Nationality Mix	0.10	0.18	0.50	0.21
Board Tenure	8.11	4.55	0.54	0.27
Board Size	8.87	2.39	0.18	-0.13

Notes. Each component is a standardized z-score, sign-aligned to the AI-favorable direction (−Board Tenure, −Board Size, so a higher value means shorter tenure / smaller board). The index is the standardized average of the six. "Corr. with other comps." is the correlation of each component with the average of the other five. The components are positively but loosely related (mean pairwise correlation 0.10, excluding the two near-identical listed-board measures), so the index summarizes several distinct traits rather than one underlying quality.

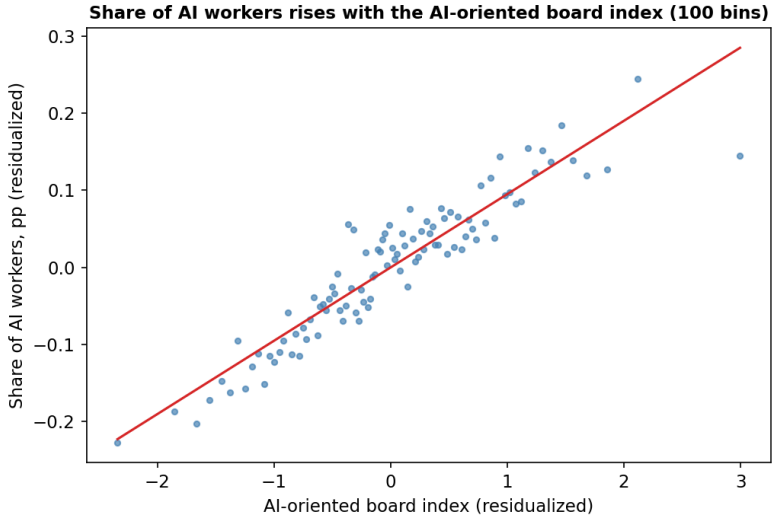
E Additional Figures

Figure E.1: AI-Oriented Board Index: Distribution and Components



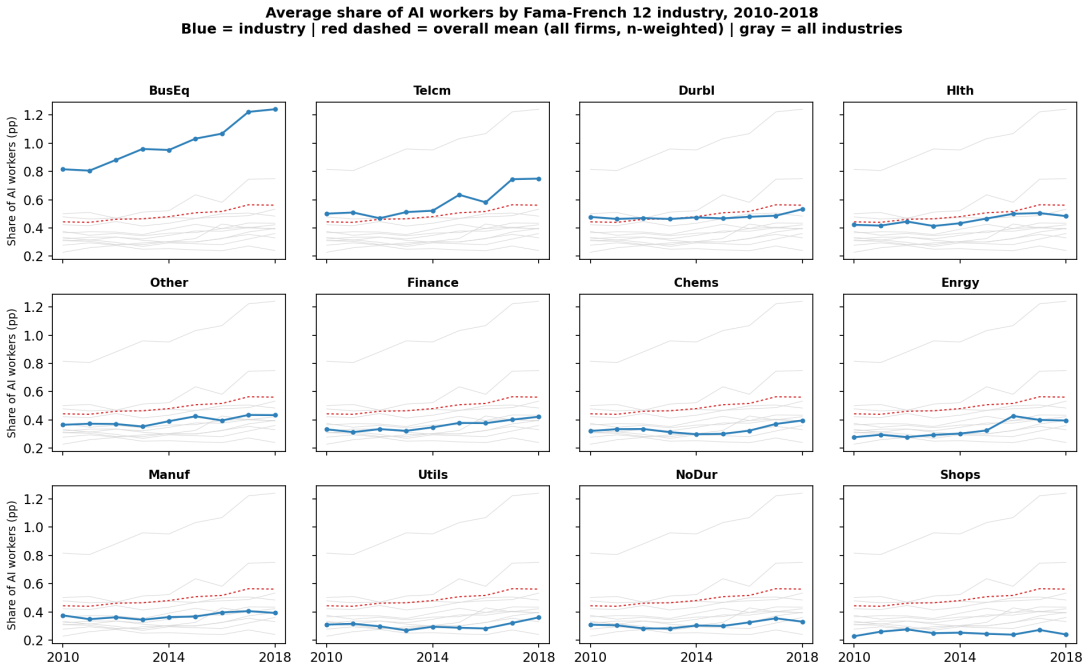
Notes. Left panel: distribution of the standardized AI-oriented board index. Right panel: pairwise correlations among the index's six sign-aligned components.

Figure E.2: AI Worker Share and the AI-Oriented Board Index



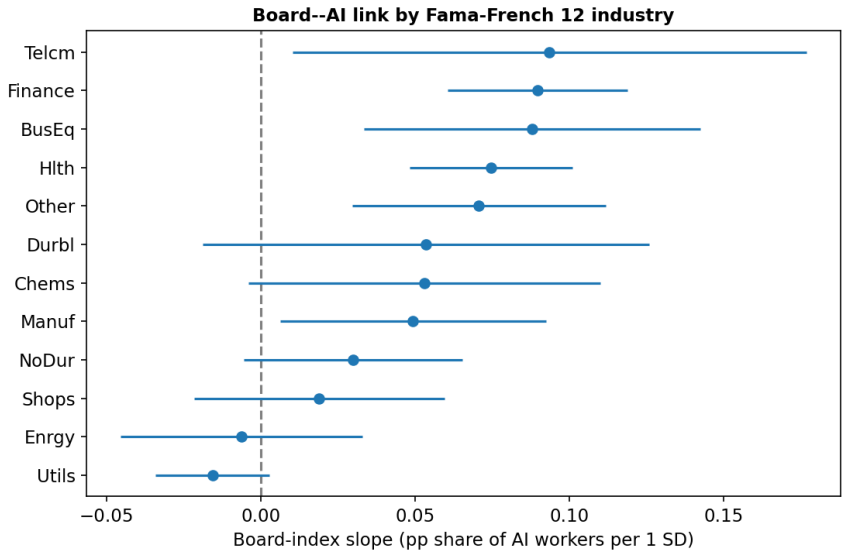
Notes. Binned scatter of the AI worker share against the AI-oriented board index, both residualized on industry×year fixed effects and firm controls, in 100 equal-sized bins.

Figure E.3: AI Worker Share by Industry



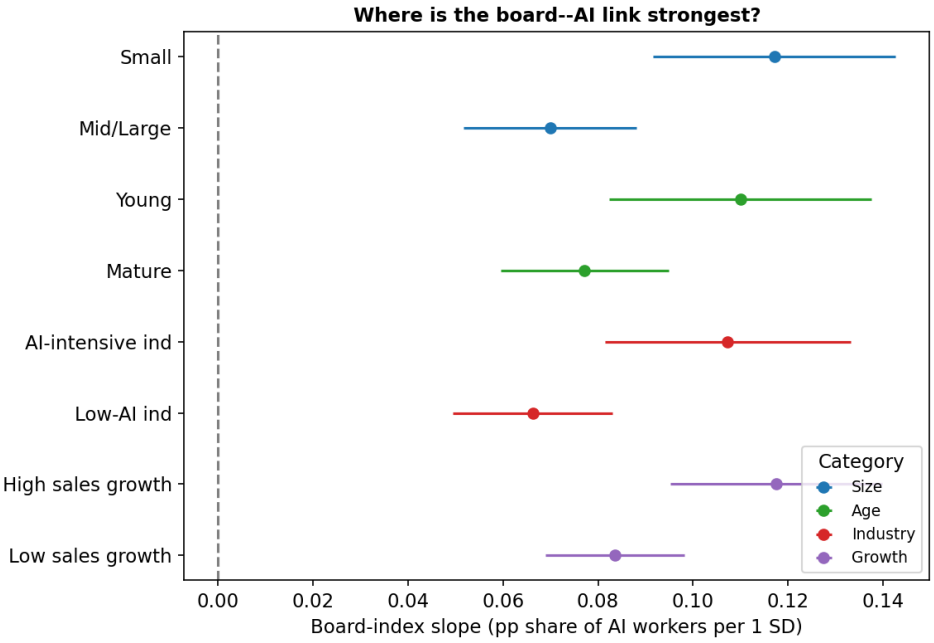
Notes. Average AI worker share by Fama and French 12 industry, 2010 to 2018. The solid line is the industry mean, the dashed line is the overall size-weighted mean, and the gray lines show all industries for comparison.

Figure E.4: Board Index Slope by Industry



Notes. The board-index slope (percentage points per standard deviation), estimated separately within each Fama and French 12 industry, with 95 percent confidence intervals.

Figure E.5: Board Index Slope across Subgroups



Notes. The board-index slope (percentage points per standard deviation), estimated on subgroups defined by firm size, firm age, industry AI-intensity, and sales growth, colored by category, with 95 percent confidence intervals.