Social Spillovers and AI Adoption: Evidence from Denmark*

Philippe Aghion† Simon Bunel‡ Xavier Jaravel§ Thomas Mikaelsen¶
Alexandra Roulet‖ Jakob Søgaard**

March 2026

Preliminary and Incomplete. Do not cite or circulate.

Abstract

This paper estimates social spillovers in the adoption of artificial intelligence technologies. Combining Danish employer–employee records with population-wide data on different social networks, we measure each firm's exposure to early AI adoption through its employees' personal networks. In a shift-share design, we find that doubling average exposure to early AI adoption through the spousal network increases AI adoption by approximately 3 percentage points. This effect is driven by connections involving high-wage employees, and is confirmed using sibling ties as an alternative network. Using network exposure as an instrument for AI adoption, we estimate positive causal effects on firm employment, consistent with the productivity effect of AI dominating the displacement effect. The positive impacts on labor demand and value added at linked firms suggests that AI adoption generated positive externalities, raising the possibility that technology diffusion may be inefficiently slow, and that policy should consider supporting adoption alongside invention.

Keywords: Artificial intelligence, Technology adoption, Social spillovers, Networks.
JEL codes: O33, J12, D22, L23.

*For helpful comments and suggestions, we thank David Autor, Mireia Giné, Raffaela Sadun, Joonas Tuhkuri, and John Van Reenen, as well as participants at the Munich Workshop on Skills, Tasks and Technologies in the AI Era 2025 and at IESE's 2026 Economics of AI Conference.

†Collège de France & INSEAD. Email: philippe.aghion@insead.edu.
‡INSEAD. Email: simon.bunel@insead.edu.
§London School of Economics. Email: X.Jaravel@lse.ac.uk.
¶Stockholm University. Email: thomas.mikaelsen@su.se.
‖INSEAD. Email: alexandra.roulet@insead.edu.
**University of Copenhagen & CEBI. Email: jes@econ.ku.dk.

1 Introduction

A broad literature has established that the invention of new technologies generates positive spillovers across firms and industries. Theoretically, Romer (1990) and Aghion and Howitt (1992) showed that innovation produces positive externalities, as knowledge created by one firm can benefit others at little or no additional cost. Empirically, Bloom et al. (2013) provide direct evidence on the magnitude of these technology spillovers across firms, and show that they outweigh negative spillovers from business-stealing across firms. Knowledge spillovers from the invention of new technologies thus provide the theoretical foundation for public support of research and development, from R&D tax credits to DARPA-style direct investment programs. In contrast, much less is known about the magnitude of spillovers in the *diffusion* of technologies. While the social returns to inventing a new technology have been extensively studied, the question of whether its *adoption* by one firm generates positive externalities on other firms has received far less attention. This asymmetry in knowledge has a direct policy counterpart: most countries actively subsidize innovation, but very few subsidize the adoption and diffusion of existing technologies.

One potentially important channel through which technologies may diffuse is *social spillovers*: adoption decisions may propagate through the interpersonal links between individuals, via mechanisms such as peer influence, social learning, or information sharing. A growing theoretical literature has formalized how such interpersonal interactions can shape the pace and patterns of technology diffusion (Bikhchandani et al., 1992; Banerjee, 1992; Golub and Sadler, 2016). While the empirical literature has documented social spillovers in the diffusion of mid-tech agricultural technologies in developing countries (e.g., Foster and Rosenzweig (1995), Bandiera and Rasul (2006), Conley and Udry (2010)), little is known empirically about the magnitude of social spillovers for modern technologies adopted by firms in developed economies.

In this paper, we study social spillovers in the adoption of artificial intelligence technologies. Using comprehensive Danish administrative data, we ask whether a firm's likelihood of adopting AI is causally affected by the AI adoption decisions of other firms to which it is connected through its employees' personal networks – and, if so, what the downstream consequences of this network-induced adoption are for firm performance and labor demand. Our analysis focuses on spousal and sibling ties, not because we believe these are necessarily the most important channels for technology diffusion, but because they provide a credible source of identification. In this sense, our findings provide compelling evidence that social networks in general play a significant role in the diffusion of AI across firms.

Our identification strategy relies on a shift-share research design, exploiting Danish administrative data over the period 2014–2023. The *shocks* come from the early adoption of AI by some

firms, which we identify using mandatory firm surveys. The *exposure shares* are constructed from workers' personal networks: because employees' spouses and siblings may work at different firms, each focal firm's workforce is heterogeneously exposed to early AI adoption at other firms through these interpersonal ties.

A key identification challenge is the reflection problem highlighted by Manski (1993): firms connected through the same network may share unobserved characteristics or face correlated shocks that drive AI adoption independently of any social spillover. For instance, homophily could induce bias: if tech workers are more likely to be married to other tech workers, a firm operating in the tech industry will be more likely to be exposed to AI adoption via spousal links, while being also directly more likely to adopt AI given its prevalence in the tech sector.

To address this concern, we follow two strategies. First, we include a rich set of focal firm controls in our regression specification. Specifically, we control for time-varying industry and local labor market fixed effects, as well as for the predicted component of network exposure to AI adoption. The predicted network exposure is constructed using a machine learning algorithm trained on a rich set of linked firms' observable characteristics. Hence, our estimates rely only on the unpredicted (and thus plausibly exogenous) component of exposure. Intuitively, we compare firms that are very similar in terms of their activity, their workforce composition, and the characteristics of their linked firm network, but whose networks happen to experience different rates of early AI adoption.

Second, we conduct several falsification tests to assess whether our estimates reflect genuine information transmission rather than unobserved confounders. We first test for parallel trends, prior to differential exposure to early AI adoption. We then examining heterogeneity across worker wage groups. If the estimated effects were driven by unobserved confounders correlated across firms, rather than by information flowing from adopters to connected firms, we would not expect the effects to be concentrated among high-wage employees, who are the most plausible carriers of relevant information and the most likely to influence adoption decisions. We also test whether the research design is spuriously capturing correlated shocks across firms. If this were the case, changes in profits among firms in the linked network would co-move mechanically with profits of the focal firm, rather than varying specifically with technology adoption.

With this research design, we find significant social spillovers in AI adoption through spousal networks. Doubling average exposure to early AI adoption at other firms increases the probability that the focal firm adopts AI by approximately 3 percentage points, with no pre-trends. This effect is entirely driven by connections involving high-wage employees, consistent with the view that individuals in influential positions are the primary vectors through which information about new technologies is transmitted across firms. In contrast, there is no impact of profit shocks, which

bolsters the credibility of the research design.

Using sibling ties as an alternative network yields qualitatively similar results, confirming that close personal relationships are a relevant channel for technology diffusion. In contrast, weaker ties constructed from shared undergraduate cohorts produce no significant effect. While some links formed during university are undoubtedly strong, education networks also include many weak connections, reducing the average signal per link and limiting our ability to detect an effect. By contrast, spousal ties generate fewer connections per employee, but they are on average much stronger, yielding a clearer signal for identifying social spillovers.

Beyond its intrinsic interest for understanding how new technologies diffuse through social networks, the effect of network exposure on AI adoption can also serve as a first stage in an instrumental variable strategy. Using the spousal network as an instrument for AI adoption, we estimate the causal effect of adoption on firm-level outcomes. Our IV estimates indicate that AI adoption leads to an increase in total employment, suggesting that the productivity gains from AI dominate the displacement of workers at the firm level. We also find positive effects on sales, value added and on the employment of young workers, contributing to an ongoing debate on whether AI adoption harms or benefits entry-level positions. While the IV estimates are imprecise due to substantial heterogeneity in firm responses, they are broadly consistent with the significant estimates obtained from a standard event-study approach used, as for example in Aghion et al. (2025).

The positive impacts on labor demand and value added at linked firms suggests that adoption at focal firms generated positive externalities that were not internalized when deciding whether to adopt AI. These results thus raise the possibility that the current pace of AI diffusion, and perhaps of technology more broadly, may be inefficiently slow and that public policy should consider supporting not only the invention but also the adoption of new technologies.

Related literature. Our paper relates to several strands of the literature. First, a theoretical literature has explored how information cascades and social learning can lead to inefficiently slow or fast technology diffusion (Bikhchandani et al., 1992; Banerjee, 1992; Golub and Sadler, 2016). Empirically, the most developed evidence on social spillovers in technology adoption comes from the agricultural economics literature. Foster and Rosenzweig (1995) show that farmers learn about new technologies from their neighbors' experiences. Besley and Case (1993) study how adoption decisions are influenced by the choices of nearby farmers. Munshi (2004), Bandiera and Rasul (2006), and Conley and Udry (2010) provide further evidence that social networks play a central role in the diffusion of agricultural innovations in developing countries. Our paper extends this line of research to a fundamentally different setting – the adoption of AI technologies by firms in a developed economy – and exploits a different type of network: the interpersonal ties between

employees at different firms.

A closely related paper is Bilgin et al. (2024), who study technology spillovers from diffusion through *firm* networks, using input-output linkages from Turkish VAT data. They find positive spillovers from robot adoption along the supply chain, but negative spillovers between horizontal competitors. Our approach differs in that we focus on technology diffusion through *personal* networks (specifically, spousal and family ties between workers) rather than through inter-firm trade linkages. This distinction is important, as the two channels operate through different mechanisms: firm networks transmit technology through production complementarities, while personal networks operate through information sharing and social learning.

Finally, our paper contributes to the rapidly growing literature on the economic effects of AI. At the industry or labor market level, the evidence is mixed: Acemoglu et al. (2022), Hampole et al. (2025), and Humlum and Vestergaard (2025) find limited or no aggregate labor market effects of AI adoption, while Babina et al. (2024) document positive effects on firm growth and product innovation. At the firm level, a growing body of evidence points toward positive effects of AI adoption on productivity, employment and sales (Babina et al., 2024; Aghion et al., 2025; Brynjolfsson et al., 2025b), consistent with a model where the productivity effect outweighs the displacement effect. A much broader literature has documented positive effects of AI on productivity at the worker level (e.g., Brynjolfsson et al. 2025b; Dell'Acqua et al. 2023; Choi et al. 2024; Noy and Zhang 2023). An ongoing debate concerns the effects of AI on young workers and entry-level positions (Brynjolfsson et al., 2025a; Humlum and Vestergaard, 2025; Hosseini and Lichtinger, 2025). We contribute to this growing literature on the firm-level effects of AI by using our shift-share research design to instrument AI adoption and estimate its effects on firm-level outcomes including employment, sales, value added, and the hiring of young workers.

The remainder of the paper is organized as follows. Section 2 describes the data sources, defines the sample and key variables, and presents summary statistics. Section 3 presents the identification strategy and the empirical specification. Section 4 reports the main results on social spillovers in AI adoption, robustness checks, and the instrumental variable estimates of the effects of AI adoption on firm-level outcomes. Section 5 concludes.

2 Data, Sample Definition and Summary Statistics

In this section, we describe the data sources, define the sample and key variables and present summary statistics.

2.1 Data

Danish administrative data provide a unique setting to study social spillovers in AI adoption. A key advantage of the data is that it combines two features that are rarely available jointly. First, matched employer–employee data cover the universe of firms and workers, providing detailed information on worker characteristics such as wages and occupation, and thereby offering a clear view of any firm's workforce. Second, and crucially for our purposes, Danish population records allow us to observe interpersonal links across individuals, notably spousal and sibling ties. This makes it possible to trace the network connections between employees at different firms, and thereby to measure the extent to which workers at a given firm are exposed to AI adoption decisions made at other firms through their personal relationships. To our knowledge, only a few countries outside Scandinavia offer this combination of comprehensive employer–employee data and population-wide information on family ties at the individual level, which is essential for credibly identifying social spillovers in technology adoption.

Beyond its exceptional data infrastructure, Denmark is also a particularly relevant country in which to study AI adoption and diffusion: throughout our sample period, it consistently exhibits the highest share of AI-adopting firms among all EU member states, well above the European average.[1] Overall, our analysis covers the period 2014–2023, spanning the early stages of modern AI diffusion in the Danish economy.

Measure of AI adoption. Our measure of firm-level AI adoption is drawn from the European ICT Enterprise Survey, administered by Statistics Denmark over the period 2017–2022. This survey is part of a harmonized European data collection effort and constitutes the most reliable source of information on actual AI usage by firms in Denmark. The survey covers enterprises with more than 10 employees operating in private, non-financial industries, corresponding to a target population of around 20,000 firms per year across our period of interest. Each year, a sample of approximately 4,000 enterprises is drawn, representative of both firm size and industry. The survey is mandatory, and the response rate is remarkably high, reaching 96%.[2] This yields a subsample covering roughly 25% of the target population in each survey wave. Importantly, the survey is exhaustive for firms with 100 or more employees, ensuring that larger firms are systematically represented.

A key strength of our approach is that the ICT Enterprise Survey provides very specific ques-

[1] According to Eurostat's ICT Enterprise Survey, the share of Danish firms using at least one AI technology reached 23.9% in 2021 and 27.6% in 2024, ranking first in the EU 27 in both years. By comparison, the EU 27 average stood at 7.7% in 2021 and 13.5% in 2024. The next highest adoption rates in 2021 were observed in Finland (15.8%), the Netherlands (13.1%), and Luxembourg (13.0%).

[2] Further details on the ICT Enterprise Survey are provided in Appendix A.1.

tions about the usage of AI technologies within firms, avoiding broad generalizations from respondents. Specifically, the survey defines AI adoption as the use of at least one of the following technologies: text mining, computer vision, speech recognition, natural language generation, machine learning, or deep learning. This technology-specific definition is important for two reasons. First, it anchors responses to concrete and well-identified technologies, thereby reducing the risk that firms overstate their AI adoption due to an overly broad interpretation of what constitutes artificial intelligence. Second, it ensures comparability across firms and over time, as the set of technologies listed in the survey provides a consistent benchmark for measuring adoption.

The ICT Enterprise Survey provides information on AI adoption over the period 2017–2022, which consistent questions on AI adoption in 2017–2019, and 2022. We use the early survey waves (2017–2019) to identify firms that had already adopted AI technologies by a given base year. These *early AI adopters* serve as the source of exogenous shocks in our shift-share design: it is their adoption status that generates variation in the network exposure of other firms, through the interpersonal connections of their employees. The later survey waves (the 2022 wave) provide an outcome measure, allowing us to assess whether focal firms that were more exposed to early AI adoption through their employees' spousal networks were themselves more likely to adopt AI subsequently.[3]

Employment data. We complement the survey evidence on AI adoption with administrative matched employer–employee data covering the universe of Danish firms and workers over the period 2014–2024. For each worker, we observe detailed characteristics including wages, hours, and occupation. At the firm level, this allows us to construct the total workforce in any given year, but also to distinguish subgroups of employees defined by their occupation or position in the within-firm wage distribution. This level of granularity is central to our empirical strategy: by decomposing each firm's workforce along these dimensions, we can construct heterogeneous measures of a focal firm's exposure to AI adoption through its employees' personal networks. In particular, we are able to test whether spillovers are primarily driven by high-wage employees (who are more likely to influence firm-level technology decisions) or whether they operate uniformly across the wage distribution.

Personal ties. Danish population records allow us to construct interpersonal links between individuals across the entire population. Depending on the administrative source used, we can define several types of personal ties for each individual: spousal links, identified from civil registration records; sibling ties, derived from parent–child relationships in the population register; and former

[3] See Section 3 for a complete discussion of the research design.

classmate networks, constructed from co-enrollment in the same undergraduate program at the same institution. In our main analysis, we focus on spousal links, which constitute the strongest form of social links. We also use sibling and education networks as alternative networks to better understand the channels through which social spillovers in adoption might operate.

These interpersonal ties are the key ingredient that allows us to connect the matched employer–employee data to the AI adoption survey. For each employee at a focal firm, we observe whether their spouse (or sibling, or former classmate) is employed at a firm that adopted AI early. Aggregating these individual-level connections at the firm level, we can then construct measures of each focal firm's exposure to early AI adoption through the personal networks of its workforce, as described in detail in Section 3.

Firm-level outcomes. Finally, we have use balance sheet data covering 2014-2023 to construct a range of additional firm-level outcomes, such as sales and value added. These variables enable us to go beyond the question of whether social spillovers affect the *diffusion* of AI, and to examine whether AI adoption induced by network exposure also has real effects on firm performance and labor demand.

2.2 Sample construction

Early adopters. Before turning to the estimation sample of focal firms, we describe the population of firms surveyed in the ICT Enterprise Survey during the early waves (2017, 2018, and 2019), distinguishing between those that had already adopted AI by a given base year – the *early adopters* – and those that had not. Understanding the characteristics of early adopters is important, as they constitute the source of the shocks in our shift-share design: it is their adoption status that generates variation in the network exposure of focal firms through spousal ties, and thereby drives the diffusion channel we seek to identify.

Table 1 reports the number of firms surveyed in each ICT Enterprise Survey wave by AI adoption status. For each wave $Y \in \{2017, 2018, 2019\}$, we distinguish between firms that had already adopted at least one AI technology by year Y (early adopters) and those that had not (non early-adopters). Pooling across the three waves, the sample comprises 11,924 firm-by-wave observations, of which 1,012 (8.5%) are early adopters. The share of early adopters increases over time, from 7% in the 2017 wave to 10% in the 2019 wave, reflecting the progressive diffusion of AI technologies across Danish firms during this period.

Table 1. Firms in the ICT Enterprise Survey by Wave and AI Adoption

Survey wave	Non-adopters	Adopters	Total
2017	3,645	256	3,901
2018	3,607	342	3,949
2019	3,660	414	4,074
2022	3,126	1,064	4,190
Total	14,038	2,076	16,114

Notes: This table reports the number of firms surveyed in the ICT Enterprise Survey by wave and AI adoption status as of the survey year. Adopters are firms that had already adopted at least one AI technology by the survey year. Non-adopters are firms that had not adopted AI by that year.

Table 2 compares the base-year characteristics of early adopters and non early-adopters. Early adopters are, on average, substantially larger than non early-adopters: the mean employment of early adopters is 423 employees, roughly three times the mean of 126 for non early-adopters. Early adopters are also more productive, in terms of turnover per employee (849 vs. 416 k€ per FTE), gross value added per employee (133 vs. 114 k€ per FTE), and profits per employee (25 vs. 24 k€ per FTE). Average annual wages are higher among early adopters (74 vs. 63 k€), suggesting a wage premium in early-adopting firms, while the labor share is comparable across the two groups (78% vs. 77%). These patterns indicate that firms at the frontier of AI adoption tend to be larger and more productive, consistent with the idea that adopting a new general-purpose technology requires a certain scale and organizational capacity. The dispersion within each group is considerable, as indicated by the interquartile ranges: for instance, the employment of early adopters ranges from 30 at the 25th percentile to 294 at the 75th percentile.

Table 2. Descriptive Statistics: Surveyed Firms by Early AI Adoption Status (2017-2019)

	Variable	Mean	p25	p50	p75	N
Non early-adopters	Employment (FTE)	126	19	48	129	10,847
	Labor share (%)	77	69	82	92	10,847
	Average wage (k€)	63	52	60	70	10,847
	Turnover per employee (k€ per FTE)	416	143	222	385	10,847
	Gross value added per employee (k€ per FTE)	114	66	85	112	10,847
	Profits per employee (k€ per FTE)	24	2	9	22	10,847
Early adopters	Employment (FTE)	423	30	107	294	1,005
	Labor share (%)	78	62	80	92	1,005
	Average wage (k€)	74	58	70	83	1,005
	Turnover per employee (k€ per FTE)	849	151	261	503	1,005
	Gross value added per employee (k€ per FTE)	133	73	98	142	1,005
	Profits per employee (k€ per FTE)	25	1	10	30	1,005

Notes: This table reports descriptive statistics for firms surveyed in the ICT Enterprise Survey, by early AI adoption status as of the survey year. Early adopters are firms that had already adopted at least one AI technology by the survey year. Non early-adopters are firms that had not adopted AI by that year. All variables are measured in the base year. Average wage, turnover, gross value added, and profits are expressed in thousands of euros (k€) per FTE employee.

Estimation sample of focal firms. Our empirical analysis relies on a staggered cohort design. For each of the three base years – 2017, 2018, and 2019 – we restrict the sample to firms that answer the ICT Enterprise Survey in at least one base year, and in the later 2022 wave. Among these, we select the firms that had *not* adopted AI in the base year, and track whether these firms adopt AI by 2022.

On to this sample of firms we merge the matched employer-employee and balance sheet data to construct a full panel of firms between 2014–2024. Hence, the data window allows us to examine both pre-trends in the years leading up to the base year and the dynamic evolution of adoption in the years that follow. The pre-period is particularly important for our identification strategy, as it allows us to verify that firms with higher network exposure to early AI adopters did not already exhibit differential trends in adoption or other outcomes prior to the treatment period.

Table 3 reports the composition of this sample by base year and adoption status in 2022[4] Pooling across the three cohorts, the focal firm sample comprises 4,771 firm-by-cohort observations, of

[4]One might refer to these firms as "late adopters" to distinguish them from the early adopters described above. However, since AI adoption by 2022 remains confined to a relatively small number of firms, we prefer the simpler terminology of *adopters* versus *early adopters*, reserving the term "early" for the firms that had already adopted AI during the 2017–2019 period and whose adoption status serves as the shock in our shift-share design.

which 1,380 (29%) are adopters that transitioned to AI adoption between the base year and 2022. The share of adopters is relatively stable across cohorts, ranging from 26% for the 2019 cohort to 32% for the 2017 cohort. By exploiting variation in base years, the staggered design increases statistical power and allows us to assess the stability of our estimates across cohorts defined at different points in time during the early diffusion of AI technologies.

Table 3. Estimation Sample by Base Year and AI Adoption

Base year	Non-adopters	Adopters	Total
2017	1,003	471	1,474
2018	1,049	447	1,496
2019	1,339	462	1,801
Total	3,391	1,380	4,771

Notes: This table reports the number of focal firms by base year cohort and AI adoption status as of 2022. Focal firms are firms that were surveyed in the ICT Enterprise Survey both in the base year and in 2022, and that had not adopted AI by the base year. Non-adopters are firms that had still not adopted AI by 2022. Adopters are firms that adopted AI between the base year and 2022.

Table 4 compares the base-year characteristics of focal firms by their subsequent adoption status. Focal firms that adopt AI by 2022 are, on average, substantially larger than those that do not in the base year: the mean employment of adopters is 372 employees, roughly twice the mean of 180 for non-adopters. Adopters also display higher turnover per employee (444 vs. 350 k€ per FTE), higher gross value added per employee (130 vs. 104 k€ per FTE), and higher profits per employee (33 vs. 20 k€ per FTE). Average wages and the labor share are higher among adopters (respectively 68 vs. 63 k€ and 68% vs. 57%). Importantly, the focal firms in our stacked sample are substantially larger than the average Danish firm, reflecting both the threshold of 10 employees for being potentially sampled in the ICT survey and the full inclusion of larger firms in the survey.

Table 4. Descriptive Statistics: Estimation Sample by AI adoption

	Variable	Mean	p25	p50	p75	N
Non-adopters	Employment (FTE)	180	63	120	195	3,383
	Labor share (%)	57	69	81	91	3,383
	Average wage (k€)	63	53	61	70	3,383
	Turnover per employee (k€ per FTE)	350	149	233	380	3,383
	Gross value added per employee (k€ per FTE)	104	71	87	111	3,383
	Profits per employee (k€ per FTE)	20	2	10	23	3,383
Adopters	Employment (FTE)	372	106	186	365	1,378
	Labor share (%)	68	64	78	91	1,378
	Average wage (k€)	68	57	65	76	1,378
	Turnover per employee (k€ per FTE)	444	175	292	496	1,378
	Gross value added per employee (k€ per FTE)	130	77	98	131	1,378
	Profits per employee (k€ per FTE)	33	3	12	28	1,378

Notes: This table reports descriptive statistics for focal firms, by AI adoption status as of 2022. Non-adopters are focal firms that had still not adopted AI by 2022. Adopters are focal firms that adopted AI between the base year and 2022. All variables are measured in the base year. Average wage, turnover, gross value added, and profits are expressed in thousands of euros (k€) per FTE employee.

3 Research Design

In this section, we present our identification strategy, formalize the exposure measure, and describe the empirical specification used to estimate social spillovers in AI adoption.

3.1 Adoption through personal networks

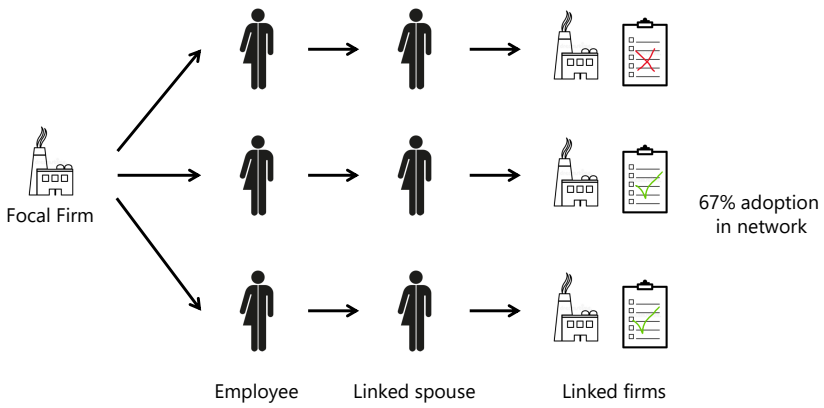
Our goal is to study the diffusion of modern AI technologies through social networks. The core idea is that technology adoption decisions at the firm level may be influenced by the personal connections of a firm's employees to workers at other firms that have already adopted AI.

To fix ideas, consider the case of spousal ties. Spouses are likely to share information about their respective work environments on a regular basis. Suppose that a member of a couple is employed in a firm that adopted AI early. If this individual is convinced of the value of the technology (because it allows either themselves or their colleagues to be more productive) they may relay this information to their partner and suggest that their partner's firm consider adopting it as well. The more a firm's workforce is exposed to early AI adoption through employees' partners, the greater the firm's exposure to AI adoption through the channel of social spillovers. Figure 1 illustrates this logic with a stylized example of two focal firms. Firm A has a higher share of employees whose

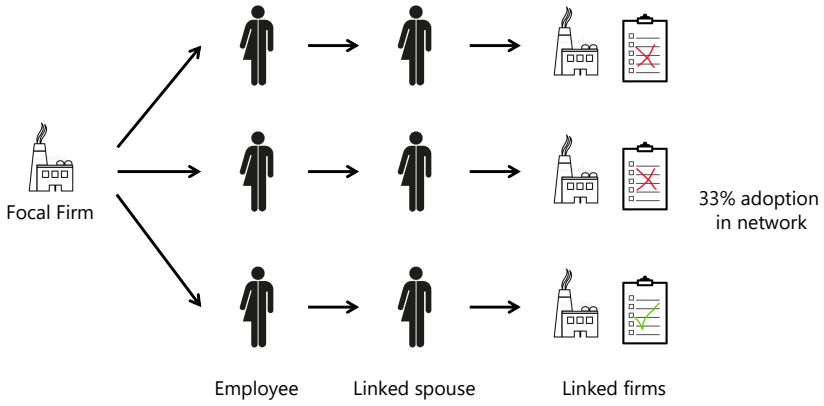
spouses work at early AI-adopting firms (67% of connected firms in the network have adopted AI) than Firm B (33%), and is therefore more exposed to potential social spillovers.

Figure 1. Exposure Shares Through Spousal Networks: Stylized Example

(a) Firm A: high exposure (67% adoption in network)



(b) Firm B: low exposure (33% adoption in network)



Notes: Each panel depicts a focal firm whose employees are linked to other firms through spousal ties. In this illustrative example, for simplicity all linked firms have answered the ICT survey, making it possible to measure AI adoption for all of them. Firm A's employees are connected to a network in which 67% of linked firms are early AI adopters, while Firm B's employees are connected to a network in which only 33% have adopted.

We formalize this intuition using a shift-share research design. The design relies on two distinct sources of variation. The first source – the *shocks* – comes from the early adoption of AI by

some firms. Using the ICT Enterprise Survey, we identify firms that had adopted at least one AI technology in a given base year Y. By 2019, approximately 6% of all surveyed firms and 23% of large firms (those with 250 or more employees) had adopted AI, providing meaningful variation in the shock component. The second source of variation – the *exposure shares* – comes from the spousal network connections of each focal firm's employees to workers at these early-adopter firms. Specifically, each focal firm's exposure to early AI adoption depends on the share of its employees whose spouse is employed at an early AI-adopting firm, as measured in the base year Y (either 2017, 2018 or 2019, as discussed in Section 2.2). Firms whose workforce happens to include more employees with spousal ties to early adopters are thus more exposed to the potential diffusion of AI through social spillovers.

Formally, we define the exposure of a focal firm i to early AI adoption as:

$$z_i = \frac{1}{N_i} \sum_e AI_e \tag{1}$$

where N_i denotes the number of employees at firm i whose partner is employed at another firm in the base year Y, and AI_e is a binary indicator equal to one if employee e's partner works at a firm that had adopted AI by the base year (and answered the ICT survey).

We can re-write the exposure share as a shift-share: Let j index firms in the network in the base year Y and define $L_{ij} = \#\{\text{employees in focal firm } i \text{ whose partner works at firm } j\}$ so that $N_i = \sum_j L_{ij}$. Further define the firm-level shock $g_j = \mathbb{1}\{\text{network firm } j \text{ had adopted AI in } Y\}$. Then the exposure of focal firm i can be written as the shift-share

$$z_i = \sum_j s_{ij} g_j \qquad \text{with} \qquad s_{ij} = \frac{L_{ij}}{N_i}. \tag{2}$$

3.2 Empirical specification

Rather than estimating a single event study pooling all years, we adopt a stacked difference-in-differences design (de Chaisemartin and D'Haultfœuille, 2020; Borusyak et al., 2024; Wing et al., 2024). This approach avoids the well-documented biases that can arise in standard two-way fixed effects specifications with staggered treatment adoption and heterogeneous treatment effects: by constructing separate cohort-specific comparisons, the stacked design ensures that already-treated firms are never used as controls for later-treated firms.

For AI adoption, we estimate event-time-specific regressions of the level of the adoption indicator on base-year network exposure. For firm outcomes such as employment, sales, and value added, we instead work with changes relative to the base year (or the pre-period average), so that

the dynamic coefficients can be interpreted as event-time treatment effects relative to baseline.

Indexing firms by i, base-year cohorts by c, and event time by t (relative to the base year Y), we estimate the following specification separately for each event time t:

$$AI_{ict} = \alpha_t + \delta_t z_i + X_i'\beta_t + \gamma_t \hat{z}_i + \lambda_{ct} + \varepsilon_{ict} \tag{3}$$

$$\Delta Y_{ict} = \tilde{\alpha}_t + \tilde{\delta}_t\, z_i + X_i'\tilde{\beta}_t + \tilde{\gamma}_t \hat{z}_i + \tilde{\lambda}_{ct} + \varepsilon_{ict}, \tag{4}$$

where AI_{ict} is an indicator for AI adoption, and ΔY_{ict} is the change in the outcome of interest between year $Y + t$ and the base year Y. λ_{ct} denotes cohort fixed effects, which are estimated separately for each event time horizon and absorb any differences across the three cohorts defined by base years 2017, 2018, and 2019. X_i and $\hat{z}_i$ denote controls described in detail below.

The main coefficient of interest is δ_t, which captures the effect of a one-unit increase in network exposure z_i (as defined in equation (1)) on the outcome at event time t. Estimating the specification separately for each event time $t \in \{-3, -2, -1, 0, +1, +2, +3, +4\}$ allows us to trace the dynamic path of this effect. In particular, the pre-period coefficients ($t < 0$) serve as a test of the identifying assumption: if firms with higher spousal network exposure to early AI adoption were already on differential trajectories before the treatment period, the coefficients δ_t for $t < 0$ would be significantly different from zero.

Controls. The vector X_i includes a set of firm-level controls measured in the base year: one-digit industry fixed effects, local labor market fixed effects, and log baseline employment. We also control for the number of links from the focal firm to other firms, and the share of these links that work in firms, who answered the ICT survey in the base year. These controls absorb differences across firms that may correlate with both network exposure and AI adoption propensity.

Controlling for Manski's reflection problem. As previously mentioned, a key challenge in identifying social spillovers is the reflection problem emphasized by Manski (1993): firms within the same network may share similar characteristics or face correlated shocks that affect AI adoption independently of any genuine social spillover.

To address this concern, we include in the specification the *predicted* exposure measure $\hat{z}_i$, constructed as follows. For each firm in the spousal network of a focal firm's employees, we predict the probability of early AI adoption using an eXtreme Gradient Boosting (XGBoost) algorithm trained on a rich set of observable characteristics: year, industry, and region fixed effects; balance sheet variables (investments, assets, profits, exports, imports, wage bill, value added, and the labor share); labor force composition (full-time employment, youth share, low-skill share, and female

share); and within-firm wage inequality (the p90/p10 and p75/p25 ratios). The predicted exposure of the focal firm is then defined as:

$$\hat{z}_i = \frac{1}{N_i} \sum_e \widehat{AI}_e \tag{5}$$

where $\widehat{AI}_e$ denotes the predicted probability that the firm employing the spouse of employee e is an early AI adopter.

By controlling for $\hat{z}_i$, our specification absorbs variation in the outcome that is driven by observable characteristics of linked firms. The identifying variation for the main coefficient of interest, δ_t, in equation (4) thus comes from the residual variation in actual exposure z_i, i.e. conditional on what would have been predicted based on linked firms' observables – that is, from the "unexpected" component of early AI adoption in the spousal network. Taken together, our specification controls for observable characteristics of both the focal firm (through X_i) and the linked firms in its spousal network (through $\hat{z}_i$), thereby eliminating many potential confounders. However, we cannot rule out that unobservable factors correlated with both network exposure and AI adoption remain. This motivates a series of falsification tests, which we present in Section 4.2.

3.3 Description of social networks

Before turning to the results, we describe the characteristics of the spousal and sibling networks that underpin our identification strategy. Tables 5 and 6 report descriptive statistics on the network structure of focal firms. For each network, we report the number of links in the base year, the share of linked firms that responded to the ICT Enterprise Survey (which determines the extent to which the exposure measure z_i is observed), the share of links connecting to early AI adopters (the actual exposure z_i), and the predicted exposure $\hat{z}_i$ as defined in equation (5).

The spousal network provides the primary source of variation in our shift-share design. On average, focal firms have 142 spousal links in the base year, with considerable dispersion: the interquartile range spans from 44 to 147. Among these links, approximately 27% connect to firms that responded to the ICT Enterprise Survey, and about 6% connect to early AI-adopting firms. The predicted exposure $\hat{z}_i$ closely tracks the actual exposure, with a mean value of 5%, confirming that a substantial share of the variation in z_i is accounted for by linked firms' observable characteristics.

Table 5. Descriptive Statistics of the Spousal Network

Variable	Mean	p25	p50	p75	N
Number of spousal links	142	44	82	147	4,702
Share of linked firms surveyed (%)	27	21	26	32	4,702
Exposure to early AI adopters, z_i (%)	6	3	6	9	4,702
Predicted exposure, $\hat{z}_i$ (%)	5	3	5	7	4,702

Notes: This table reports descriptive statistics on the spousal network of focal firms. The number of spousal links is the number of employees at the focal firm whose spouse is employed at another firm in the base year. The share of linked firms surveyed is the fraction of these links for which the spouse's employer responded to the ICT Enterprise Survey. Exposure to early AI adopters (z_i) is the share of links connecting to a firm that had adopted AI by the base year, as defined in equation (1). Predicted exposure ($\hat{z}_i$) is the predicted probability of early AI adoption aggregated across linked firms, as defined in equation (5). All variables are measured in the base year.

The sibling network is substantially more extensive than the spousal network, as each employee may have several siblings working at different firms but at most one spouse. On average, there are 361 sibling links per focal firm, roughly two and a half times the corresponding figure for the spousal network, with an interquartile range spanning from 95 to 346. The share of linked firms surveyed in the ICT Enterprise Survey is comparable to that of the spousal network (about 28%), as is the share of links to early AI adopters (about 6%). Predicted exposure closely tracks actual exposure, with a mean value of 6%.

Overall, the two networks display similar patterns in terms of survey coverage, exposure to early AI adoption, and predicted exposure, but differ markedly in size. These network characteristics set the stage for the estimation results presented in Section 4.

Table 6. Descriptive Statistics of the Sibling Network

Variable	Mean	p25	p50	p75	N
Number of links	361	95	186	346	4,754
Share of linked firms surveyed (%)	28	25	28	31	4,754
Exposure to early AI adopters, z_i (%)	6	4	6	8	4,754
Predicted exposure, $\hat{z}_i$ (%)	6	4	6	7	4,754

Notes: This table reports descriptive statistics on the sibling network of focal firms. The number of sibling links is the number of employees at the focal firm who have a sibling employed at another firm in the base year. The share of linked firms surveyed, exposure to early AI adopters (z_i), and predicted exposure ($\hat{z}_i$) are defined analogously to Table 5. All variables are measured in the base year.

4 Results

4.1 The social spillovers of adoption

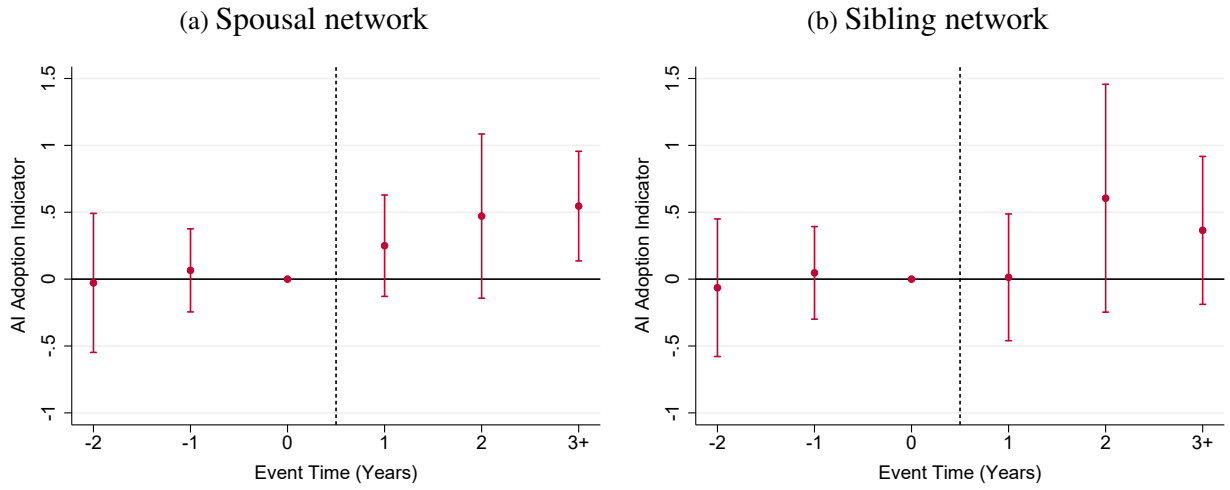
We find that firms whose employees are more exposed to early AI adoption through their spousal network are themselves significantly more likely to adopt AI. Panel (a) of Figure 2 implements our main specification using spousal ties to construct the exposure measure. For the pre-period ($t \leq 0$), the estimated coefficients $\hat{\delta}_t$ are close to zero and statistically insignificant, i.e. there are no differential pre-trends between more and less exposed firms.

After the base year ($t = 0$), the effect of network exposure on AI adoption increases steadily, reaching approximately 0.5 by event time $t = 3+$, which pools years 3, 4, and 5 (the 2022 survey wave for the three base year cohorts). To interpret this magnitude, a 10 percentage point increase in exposure raises the probability of AI adoption by 5 percentage points. Given that average exposure is 6.3%, another way to assess the economic significance of this estimate is that doubling average exposure to early AI adoption increases AI adoption by approximately 3 percentage points.

Panel (b) of Figure 2 presents the results using sibling ties as an alternative network. The pattern is qualitatively similar to the spousal network – coefficients are near zero in the pre-period and turn positive after the base year with similar magnitude as for the spousal network – but the estimates are a bit less precise, reflecting somewhat lower statistical power with this network. This may be expected, as there might be more heterogeneity in the extent of interactions and information sharing between siblings than between spouses.

Appendix Figure A5 reports the results using education network ties, constructed from co-enrollment in the same undergraduate cohort. In contrast with the spousal network, we find no clear effect of education-based network exposure on AI adoption in the long run. Indeed, while the coefficient at $t = 2$ is positive, the estimate for $t = 3+$, which pools years 3, 4, and 5, is close to zero, suggesting that these ties have no lasting impact on adoption, unlike spousal or sibling networks. It suggests that the social spillovers operate through close personal ties that involve regular and sustained information exchange (such as those between spouses or siblings) rather than through weaker or more distant connections (in contrast with some theories highlighting the importance of weak ties, e.g. Granovetter 1973).

Figure 2. Effect of network exposure on AI adoption



Notes: Each panel plots the estimated coefficients $\hat{\delta}_t$ from equation (3), where the outcome is an indicator for AI adoption. Panel (a) uses spousal ties to construct the exposure measure and Panel (b) uses sibling ties. The dashed vertical line marks the base year. Confidence intervals are at the 95% level.

4.2 Falsification tests

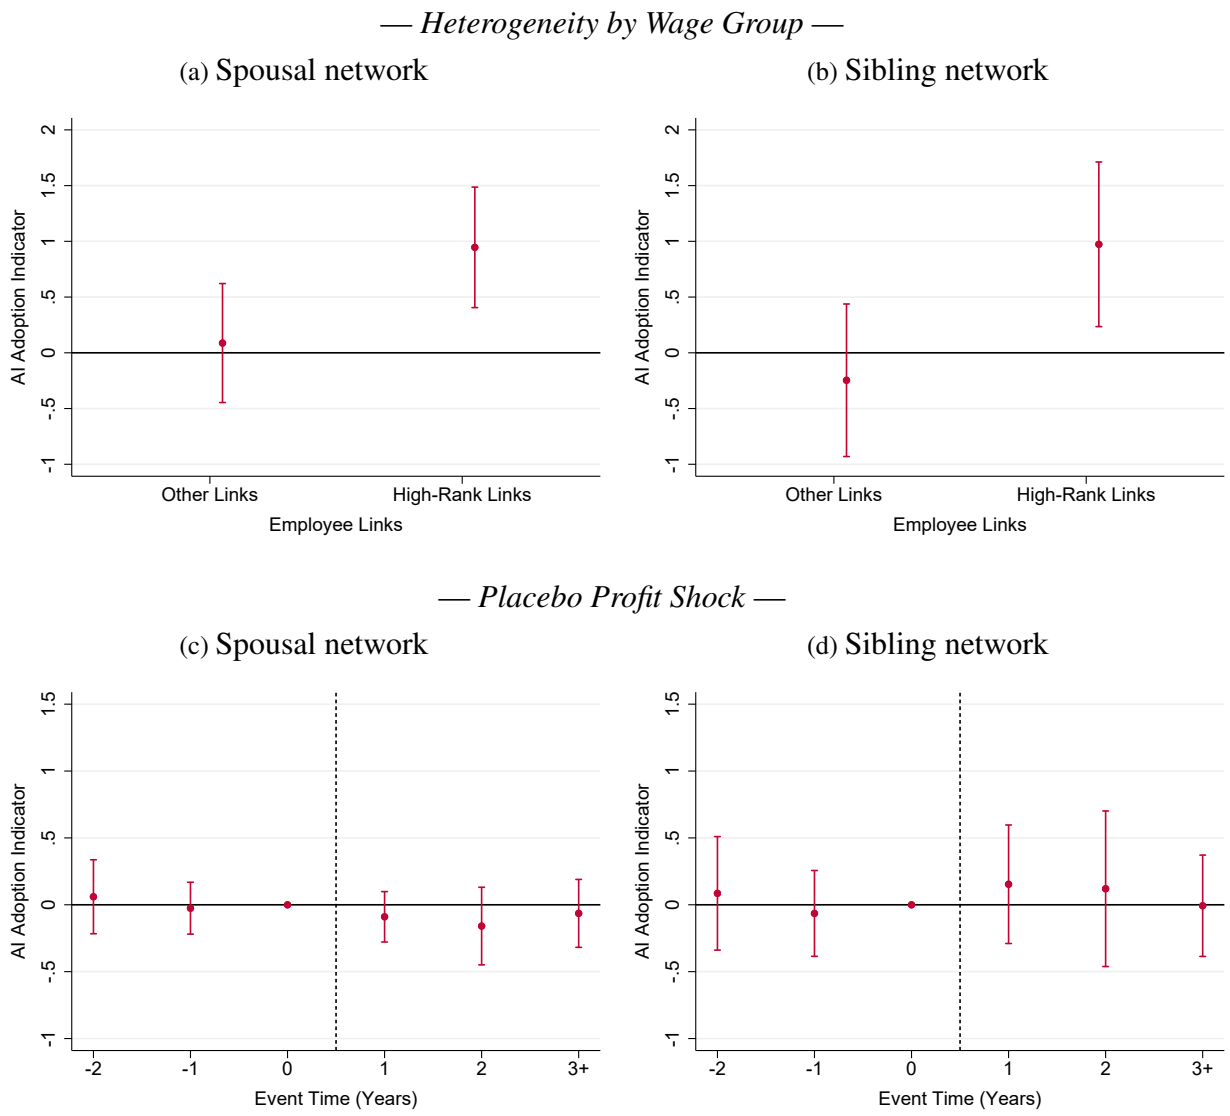
As discussed in Section 3.2, our specification controls for a wide range of firm and network characteristics, as well as for the predicted component of exposure. Nevertheless, unobservable confounders may remain. We now present a series of falsification tests designed to assess the plausibility of a causal interpretation.

Heterogeneity by wage group. Our first falsification test exploits the fact that not all personal connections within a firm's spousal network (resp. sibling network) are equally relevant for driving AI adoption. On the one hand, not all spouses (resp. siblings) linked to the focal firm's workforce are necessarily in positions where they are exposed to AI in their own workplace. On the other hand, not all employees of the focal firm are in positions where they would either be directly exposed to AI technologies or be close enough to the firm's decision-makers to influence the adoption decision. If our results were driven by unobserved confounders correlated with network structure rather than by genuine social spillovers, we would expect the effect to be similar regardless of the position of the linked individuals in the wage distribution.

To test this, we construct two alternative exposure measures by restricting the set of spousal links (resp. sibling links) used to compute z_i. The first—labeled "High-Rank Links"—retains only links in which both of the two connected individuals (the employee at the focal firm and the

spouse/sibling at the linked firm) is a high-wage worker, defined as being in the top half of the Danish wage distribution. The second—labeled "Other Links"—retains all other links. We then estimate the average post-base-year effect of each type of exposure on AI adoption.

Figure 3. Robustness and Falsification Tests



Notes: Panels (a) and (b) report the effect of network exposure on AI adoption separately by wage group of the linked employees: connections involving at least one high-wage worker versus connections between low-wage workers only, for the spousal and sibling networks respectively. Panels (c) and (d) report a placebo test in which the exposure measure is constructed using positive profit shocks in network firms (defined as a change in $\operatorname{asinh}(\text{net profit})$ falling in the top 5% of the industry-by-year distribution) instead of early AI adoption, again for the spousal and sibling networks respectively. Confidence intervals are at the 95% level.

The results are presented in Panels (a) and (b) of Figure 3 for the spousal and sibling networks respectively. In both cases, the effect of network exposure on AI adoption is driven by connections involving at least one high-wage worker, while the estimated effect through other links is close to zero and statistically insignificant. This pattern is consistent with the social spillovers channel: high-wage employees are more likely to hold positions where they are exposed to AI technologies and to have influence over (or access to) the firm's technology adoption decisions. Thus, our findings do not appear to be driven by unobserved shocks affecting all employees within a network uniformly, but rather by information transmission through individuals who are plausibly in a position to both acquire knowledge about AI and act on it.

Placebo: profit shocks. An alternative concern is that positive performance shocks at network firms, transmitted through the spousal network (resp. sibling network), could independently encourage the focal firm to adopt new technologies or be correlated across similar firms, regardless of whether the shock is specifically related to AI. If this were the case, our results could reflect a general tendency for firms exposed to successful peers to invest in new technologies, rather than a specific social spillover channel operating through information about AI.

To test this, we construct a placebo exposure measure by replacing each linked firm's early AI adoption status with an indicator for whether it experienced a large positive profit shock. Specifically, we define a network firm as experiencing a positive shock when its change in $\operatorname{asinh}(\text{net profit})$ falls in the top 10% of the year-specific distribution. We then average these firm-level profit shocks through the spousal network (resp. sibling network), exactly as we do for early AI adoption, to compute a focal firm's exposure to positive performance shocks in its personal network. We estimate the same specification as in equation (4), using this placebo exposure in place of z_i.

The results are presented in Panels (c) and (d) of Figure 3 for the spousal and sibling networks respectively. In both cases, the estimated coefficients are close to zero and statistically insignificant throughout the event window, with no discernible pattern in either the pre- or post-period. This finding is reassuring: it indicates that exposure to generally successful firms through personal networks does not, by itself, drive AI adoption at the focal firm. The social spillovers we identify thus appear to be specific to the diffusion of AI technologies through personal ties, rather than reflecting a broader channel through which positive economic conditions propagate across connected firms.

4.3 Using spousal network as an instrument

Beyond its intrinsic interest for understanding how a new technology diffuses through social spillovers, the effect of network exposure on AI adoption documented in Section 4.1 can also serve as a first stage in an instrumental variable strategy. The idea is straightforward: since we

have established that spousal network exposure causally affects a firm's probability of adopting AI, we can use this variation to instrument the endogenous adoption decision and estimate the causal effects of AI adoption on a range of firm-level outcomes. This approach addresses the endogeneity concern that firms choosing to adopt AI may differ systematically from non-adopters in ways that also affect their performance and labor demand.

We use base-year network exposure to AI adoption as an instrument for whether a firm transitions from non-adoption in the base year to adoption by 2022. The resulting IV estimates should therefore be interpreted as the causal effect of being induced to become an adopter by 2022, computing the average response in the four years after adoption by the early adopter.

Figure 4 presents the 2SLS estimates of the effect of AI adoption on two firm-level outcomes, using the spousal network exposure measure z_i as an instrument. In each panel, we plot the dynamic IV estimates across event times, allowing us to assess both the presence of pre-trends and the evolution of the treatment effect over time.

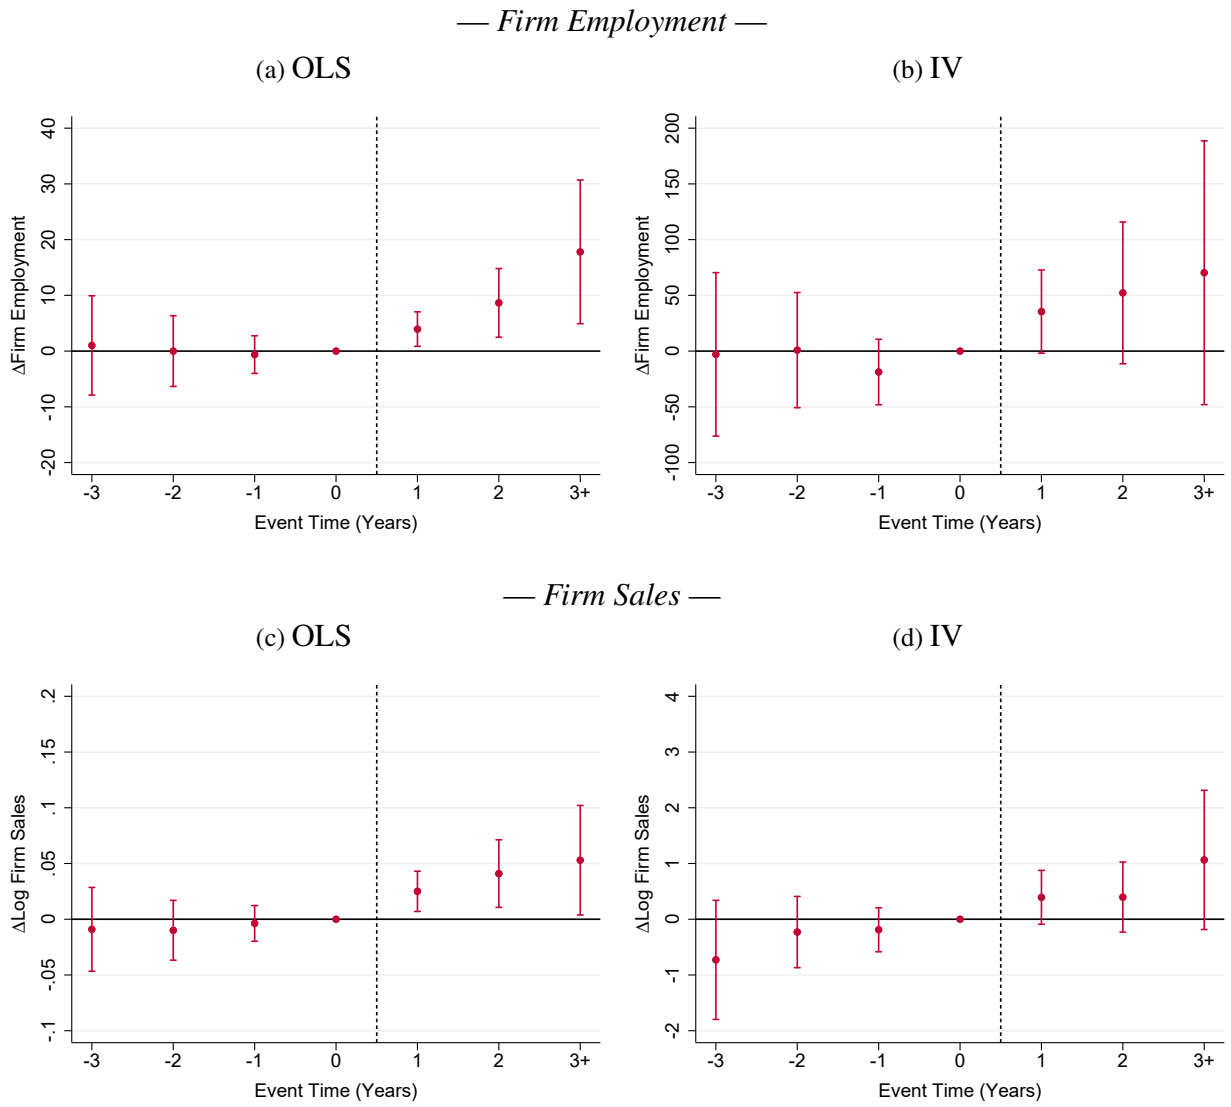
Panels (a) and (b) report the results for total firm employment. Panel (a) presents the OLS estimates: the pre-period coefficients are close to zero and the effect of AI adoption on employment is positive in the post-period. Panel (b) presents the corresponding IV estimates, using spousal network exposure as an instrument for AI adoption. The pre-period coefficients are again close to zero, i.e. there are no differential trends between more and less exposed firms prior to the base year. After the base year, AI adoption leads to a significant increase in employment, with point estimates growing steadily over the event window. A comparison of OLS and IV magnitudes is discussed in the paragraph below.

Panels (c) and (d) present the OLS and IV estimates respectively for log firm sales. The OLS estimates in Panel (c) show small positive effects in the post-period, while the IV estimates in Panel (d) display a similar pattern but with larger magnitudes. In both cases, the pre-period coefficients are close to zero, supporting the validity of the research design. After the base year, AI adoption leads to a significant increase in sales, consistent with the positive employment response documented above.

Overall, these patterns are consistent with a model where the productivity effects of AI outweigh the direct displacement effects. By making firms more productive, AI adoption leads to an increase in firm sales – through lower prices or higher quality, which we do not directly observe –, which ultimately raises labor demand.[5]

[5]These findings are consistent with the literature on the firm-level effects of automation or robots, where a majority of studies documents a positive response (e.g., Acemoglu et al. 2020; Dixon et al. 2021; Domini et al. 2021; Humlum 2021; Koch et al. 2021; Curtis et al. 2021; Acemoglu et al. 2023; Bilgin et al. 2024; Aghion et al. 2023), with a few studies estimating a negative response (Bessen et al. 2020; Bonfiglioli et al. 2025).

Figure 4. Effect of AI Adoption on Firm-Level Outcomes: OLS and IV Estimates



Notes: Panels (a) and (b) report OLS and 2SLS estimates respectively of the effect of AI adoption on total firm employment (FTE), using the spousal network exposure measure z_i as an instrument for AI adoption in the IV specification. Panels (c) and (d) report OLS and 2SLS estimates respectively of the effect of AI adoption on log firm sales. The dashed vertical line marks the base year. Confidence intervals are at the 95% level.

Comparing OLS and IV magnitudes. Table 7 complements the analysis by reporting average post-base-year estimates of the effect of AI adoption on a range of firm-level outcomes. Panel A presents the OLS estimates and Panel B the corresponding IV estimates using spousal network

exposure as an instrument.

The OLS estimates in Panel A indicate that AI adoption is associated with a significant increase in total firm employment of approximately 17.8 employees, relative to a mean of 372.2. Given this baseline mean, the OLS point estimate implies an increase of roughly 4.8% in firm employment. This effect is partly driven by an increase in young workers' employment ($+11.7$ employees) and in low-skilled employment ($+14.4$ employees). As seen with the event studies, AI adoption is also associated with an increase in log sales ($+5.3\%$) and log value added ($+5.5\%$). The IV estimates in Panel B point in the same direction but are substantially larger in magnitude, with a total employment effect of 70.3 employees – implying an increase of roughly 18.9% relative to the baseline mean – although the estimates are less precisely estimated and not statistically significant. A similar pattern emerges for log sales, with an IV point estimate of 1.065 compared to 0.053 in OLS, and for log value added, with an IV point estimate of 0.170 compared to 0.055 in OLS.

Table 7. Firm-Level Effects of AI Adoption: OLS and IV Estimates

	Firm Employment			Log	Log
	Total	Young	Low Skilled	Sales	Value Added
Panel A: OLS					
AI adoption	17.80**	11.70***	14.42*	0.053*	0.055*
	(6.58)	(3.16)	(5.73)	(0.025)	(0.024)
Panel B: IV					
AI adoption	70.30	32.02	44.24	1.065	0.170
	(60.37)	(27.87)	(50.47)	(0.637)	(0.406)
Mean Outcome	372.2	129.4	129.4	6.088	5.002
Observations	13,712	13,712	13,712	9,074	8,959
Clusters	2,091	2,091	2,091	2,089	2,075

Notes: This table reports estimates of the average post-base-year effect of AI adoption on firm-level outcomes. Panel A presents OLS estimates and Panel B presents IV estimates using the spousal network exposure measure z_i as an instrument for AI adoption. The dependent variables are total firm employment (in FTE), young workers' employment (defined as workers aged below 35), low-skilled workers' employment, log sales, and log value added. Standard errors (in parentheses) are clustered at the firm level. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Several considerations may help reconcile the gap between the OLS and IV estimates. First, the standard errors of the IV estimates are large: for instance, the 95% confidence interval around the IV employment estimate (70.3 ± 118.3) includes the OLS point estimate of 17.8, so that the two estimates are not statistically distinguishable from each other. The difference in magnitudes

should therefore be interpreted with caution. Second, the compliers in the IV strategy (that is, the firms whose adoption decision is influenced by spousal network exposure) may not be representative of the average adopter, and their characteristics warrant further investigation. Third, firms may be more likely to endogenously adopt AI during economic downturns or periods of restructuring, when they are not experiencing large positive demand shocks. If adoption tends to coincide with unfavorable conditions that independently depress employment, the OLS estimate will be negatively biased, understating the true causal effect.[6]

4.4 A simple benchmark optimal subsidy for AI adoption

A simple stylized symmetric diffusion model delivers transparent policy implications from our previous estimates. Firms indexed by i optimally decide whether to pay a (firm-specific) fixed cost C_i to acquire information about how to adopt AI. Conditional on paying that cost, firms can invest in AI optimally, until the marginal cost of investment equates the marginal benefit. These investments generate a private value P of adoption, assumed to be identical across all firms for simplicity. Thus, firm i's optimal choice is to adopt AI whenever $P > C_i$.

We now introduce an *information externality*. For simplicity, we consider a stark case, assuming that the fixed cost is reduced to zero for firms exposed to AI adoption at other firms via social spillovers. All exposed firms can thus derive private value P from adoption. The externality drives a wedge between the privately optimal and socially optimal adoption decisions: adoption at any given firm expands the set of firms for which adopting becomes privately worthwhile.

In this highly stylized setting, one can very simply relate the information externality to our empirical estimates and obtain an optimal Pigouvian subsidy. Let FS denote the first-stage effect of exposure to adopters on a firm's own adoption decision. If one firm's adoption induces FS additional adoptions in the next round on average, and if the same propagation repeats in subsequent rounds, then one initial adoption generates $1 + FS + FS^2 + \cdots = 1/(1 - FS)$ total adoptions. The total surplus from adoption (net of investment costs but not accounting for the fixed cost C_i) is therefore $1/(1 - FS)$ times the private value P. Therefore it is socially optimal for firm i to adopt whenever $P/(1 - FS) > C_i$. The optimal Pigouvian subsidy t closing the wedge between the private and social optimums is therefore such that $(1+t)P = P/(1-FS)$. With our first-stage estimate of approximately $FS = 0.5$, this yields $t = FS/(1 - FS) = 100\%$.

Thus, in this stylized benchmark, the external effect of one adoption is about as large as its direct private effect. Accordingly, the optimal subsidy rate is roughly one-for-one. This argument

[6]This endogeneity concern is related to the "opportunity cost" mechanism formalized by Aghion and Saint-Paul (1998), who show that firms have stronger incentives to invest in productivity-improving activities during downturns, when the opportunity cost of disrupting current production is lower.

only requires that adoption has a positive direct social value; it does not require a precise estimate of the level effect of adoption on firm outcomes.[7]

This exercise should be interpreted as a partial-equilibrium benchmark rather than as a full welfare analysis. In particular, it abstracts from fiscal costs, general equilibrium effects, and any negative business-stealing externalities across firms, as in Bloom et al. (2013).

5 Conclusion

This paper provides evidence that modern AI technologies diffuse across firms through social spillovers. Using Danish administrative data that uniquely combine matched employer–employee records with population-wide information on interpersonal ties, we show that a firm's probability of adopting AI is significantly influenced by the AI adoption decisions of other firms to which it is connected through its employees' spousal networks. Doubling average network exposure to early AI adoption raises the focal firm's probability of adopting AI by approximately 3 percentage points. This effect is driven by connections involving high-wage employees, consistent with the idea that individuals in senior or influential positions are the primary vectors of technology diffusion through personal networks. Alternative networks based on sibling ties yield qualitatively similar results.

Beyond the diffusion of AI itself, we exploit the social spillovers channel as an instrument for AI adoption to estimate its causal effects on firm-level outcomes. Our IV estimates indicate that AI adoption leads to a significant increase in firm employment, suggesting that the productivity gains from AI dominate the displacement of workers at the firm level. These results contribute to a growing body of evidence on the firm-level effects of AI, and point toward a net positive impact on labor demand and value added, at least for the firms induced to adopt by network exposure.

Our findings carry an important potential policy implication. While a large body of work has documented spillovers in the *invention* of new technologies, much less attention has been paid to spillovers in the *adoption* and *diffusion* of technologies. Our results suggest that adoption spillovers are empirically significant, raising the possibility that the pace of AI diffusion may be inefficiently slow. If the social returns to adoption indeed exceed the private returns, as the existence of positive spillovers suggests, then the traditional focus of innovation policy on subsidizing invention may need to be complemented by measures that also support the diffusion of new technologies. Such measures could include subsidies, training programs or coordination mechanisms to share best practices for the adoption of certain technologies. A full characterization of these policy implications is an important avenue for future research.

[7] In ongoing work, we extend this analysis to the case of asymmetric personal networks across firms and heterogeneous cost structures across firms. The optimal Pigouvian subsidy then varies across firms and depends on the full social network structure.

References

Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo, "Competing with Robots: Firm-Level Evidence from France," *AEA Papers and Proceedings*, 2020, *110*, 383–388.

_ , Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas, "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey," Working Paper 30659, National Bureau of Economic Research 2022.

_ , _ , _ , _ , _ , _ , _ , _ , _ , _ , _ , and _ , "Advanced Technology Adoption: Selection or Causal Effects?," *AEA Papers and Proceedings*, 2023, *113*, 210–214.

Aghion, Philippe and Gilles Saint-Paul, "Virtues of Bad Times: Interaction Between Productivity Growth and Economic Fluctuations," *Macroeconomic Dynamics*, 1998, *2* (3), 322–344.

_ and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

_ , Céline Antonin, Simon Bunel, and Xavier Jaravel, "Modern Manufacturing Capital, Labor Demand and Product Market Dynamics: Evidence from France," CEP Discussion Paper 1910 2023.

_ , Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard, "How Different Uses of AI Shape Labor Demand: Evidence from France," *AEA Papers and Proceedings*, 2025, *115*, 62–67.

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

Bandiera, Oriana and Imran Rasul, "Social Networks and Technology Adoption in Northern Mozambique," *The Economic Journal*, 2006, *116* (514), 869–902.

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Besley, Timothy and Anne Case, "Modeling Technology Adoption in Developing Countries," *American Economic Review*, 1993, *83* (2), 396–402.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, "Firm-Level Automation: Evidence from the Netherlands," *AEA Papers and Proceedings*, 2020, *110*, 389–393.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bilgin, Nuriye Melisa, Ester Faia, and Gianmarco I. P. Ottaviano, "Technology Spillovers, Diffusion and Rivalry in Firm Networks," CEPR Discussion Paper 19804 2024.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia, and Ioannis Papadakis, "Artificial Intelligence and Jobs: Evidence from US Commuting Zones," *Economic Policy*, 2025, *40* (121), 145–194.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Working Paper, Stanford Digital Economy Lab 2025.

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

Choi, Jonathan H., Amy B. Monahan, and Daniel Schwarcz, "Lawyering in the Age of Artificial Intelligence," *Minnesota Law Review*, 2024, *109* (1), 147–218.

Conley, Timothy G. and Christopher R. Udry, "Learning about a New Technology: Pineapple in Ghana," *The American Economic Review*, 2010, *100* (1), 35–69.

Curtis, E. Mark, Daniel G. Garrett, Eric C. Ohrn, Kevin A. Roberts, and Juan Carlos Suárez Serrato, "Capital Investment and Labor Demand," Working Paper 29485, National Bureau of Economic Research 2021.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," Working Paper 24-013, Harvard Business School 2023.

Dixon, Jay, Bryan Hong, and Lynn Wu, "The Robot Revolution: Managerial and Employment Consequences for Firms," *Management Science*, 2021, *67* (9), 5586–5605.

Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich, "Threats and Opportunities in the Digital Era: Automation Spikes and Employment Dynamics," *Research Policy*, 2021, *50* (7), 104137.

Foster, Andrew D. and Mark R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 1995, *103* (6), 1176–1209.

Golub, Benjamin and Evan Sadler, "Learning in Social Networks," in Yann Bramouillé, Andrea Galeotti, and Brian W. Rogers, eds., *The Oxford Handbook of the Economics of Networks*, Oxford University Press, 2016, chapter 19, pp. 504–542.

Granovetter, Mark S., "The Strength of Weak Ties," *American Journal of Sociology*, 1973, *78* (6), 1360–1380.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," Working Paper 33509, National Bureau of Economic Research 2025.

Hosseini, Seyed M. and Guy Lichtinger, "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," SSRN Scholarly Paper 5425555 2025.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," Working Paper, University of Chicago 2021.

_ **and Emilie Vestergaard**, "Large Language Models, Small Labor Market Effects," Working Paper 33777, National Bureau of Economic Research 2025.

Koch, Michael, Ilya Manuylov, and Marcel Smolka, "Robots and Firms," *Economic Journal*, 2021, *131* (638), 2553–2584.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, 1993, *60* (3), 531–542.

Munshi, Kaivan, "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution," *Journal of Development Economics*, 2004, *73* (1), 185–213.

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

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Wing, Coady, Seth M. Freedman, and Alex Hollingsworth, "Stacked Difference-in-Differences," Working Paper 32054, National Bureau of Economic Research 2024.

A Appendix

A.1 The ICT Enterprise Survey in Denmark

The data on AI adoption used in this paper are drawn from the European survey on ICT usage and e-commerce in enterprises, administered in Denmark by Statistics Denmark (*Danmarks Statistik*). This appendix provides additional details on the survey's institutional framework, methodology, and the specific questions used to measure AI adoption.

Institutional framework. The survey is part of a harmonized European data collection effort coordinated by Eurostat, the Statistical Office of the European Union. Its legal basis is Regulation (EU) No. 2019/2152 of the European Parliament and of the Council on European business statistics, implemented through annual Commission implementing regulations that specify the data requirements for each reference year. The survey is conducted annually by the National Statistical Institutes of participating countries on the basis of a model questionnaire provided by Eurostat, ensuring cross-country comparability. In Denmark, the survey is mandatory, and non-response is subject to legal penalties, which contributes to an exceptionally high response rate of approximately 96%.

Target population and sampling. The target population consists of enterprises with 10 or more employees operating in private, non-financial industries. The sectoral coverage spans manufacturing, construction, wholesale and retail trade, transportation, accommodation and food services, information and communication, real estate, professional and scientific activities, and administrative and support services (NACE Rev. 2 Sections C through N, as well as Group 95.1). Financial and insurance activities are excluded. In Denmark, the target population comprises approximately 20,000 enterprises per year over our period of interest. The sample is drawn using a stratified random sampling design, with strata defined along two dimensions: firm size (small: 10–49 employees; medium: 50–249; large: 250 or more) and industry. The survey is exhaustive for large enterprises (100 or more employees), meaning that all such firms are surveyed every year. For small and medium enterprises, a random sample is drawn within each stratum. The resulting sample size is approximately 4,500 enterprises per year, covering roughly 25% of the target population.

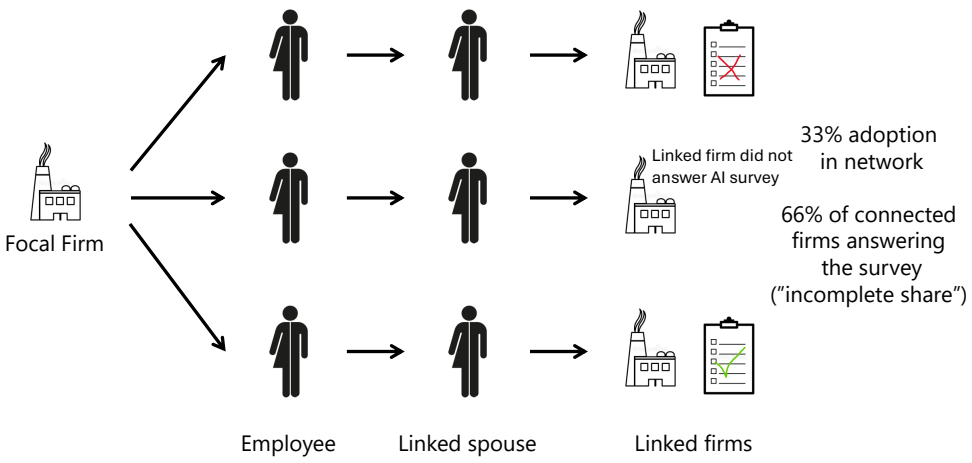
Measurement of AI adoption. Questions on AI technologies were first included in the survey in 2017. The survey defines AI adoption as the use by the enterprise of at least one of the following technologies: text mining and natural language processing; speech recognition; computer vision and image recognition; natural language generation; machine learning (including

deep learning); and technologies for physical movement (autonomous robots, self-driving vehicles). This technology-specific definition, based on the Eurostat model questionnaire, ensures that respondents anchor their answers to concrete and well-identified technologies, rather than relying on a broad or subjective interpretation of what constitutes artificial intelligence. Questions on AI adoption typically refer to the current situation at the time of the survey, i.e., the first months of the reference year.

A.2 Accounting for network structure

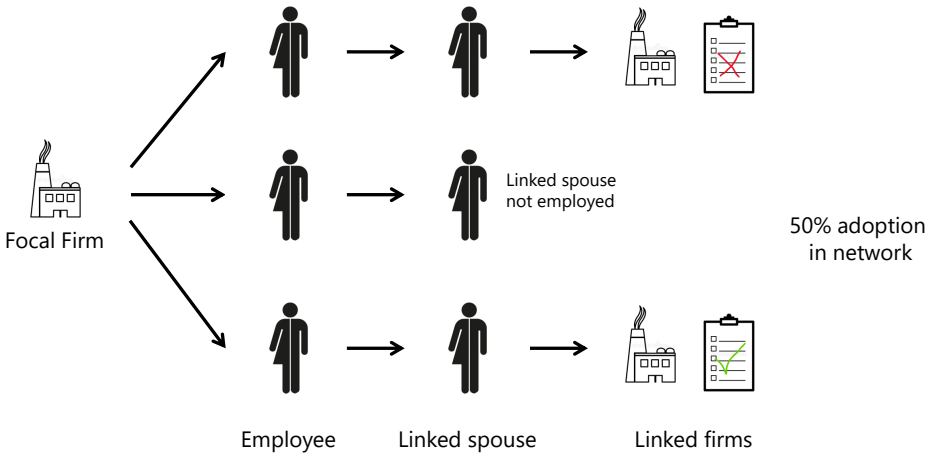
The following figures illustrate the different cases of network structure that may affect the computation of the exposure measure z_i, as discussed in Section 3.2.

Appendix Figure A1. Exposure shares with incomplete network information: a stylized example



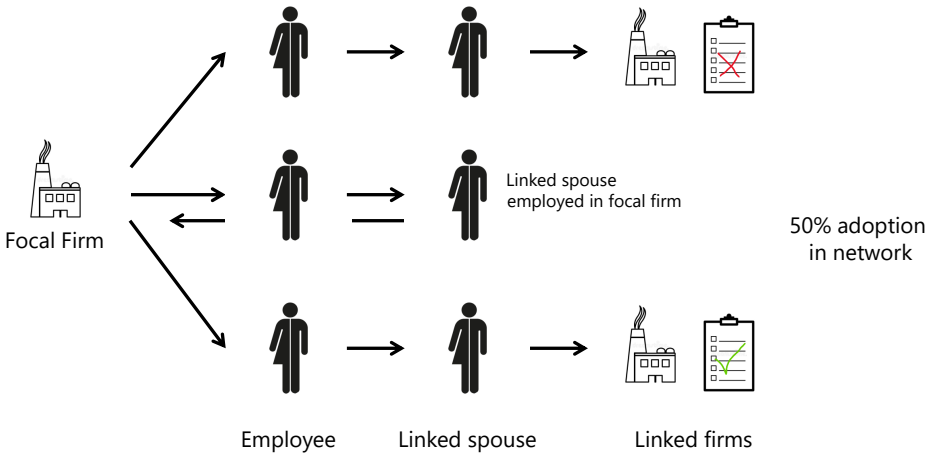
Notes: This figure illustrates the case of incomplete shares. The focal firm has three employees with spouses working at other firms, but only two of these linked firms responded to the ICT Enterprise Survey. Among the two observed linked firms, one is an early AI adopter, yielding a measured exposure of 33% (one out of three) and an incomplete share of 66% (two out of three linked firms answering the survey).

Appendix Figure A2. Exposure shares with an unemployed spouse: a stylized example



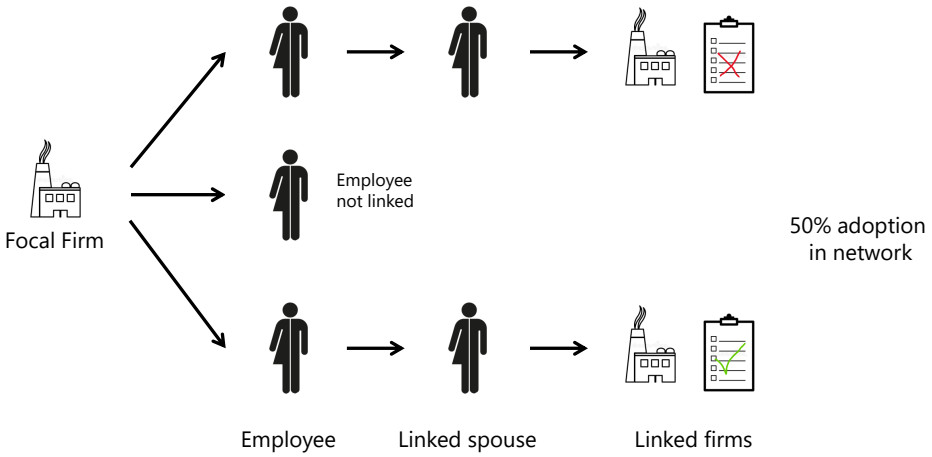
Notes: This figure illustrates the case in which one employee's spouse is unemployed or out of the labor force, and therefore does not generate a link to any outside firm. The measured exposure is computed over the remaining connected employees only.

Appendix Figure A3. Exposure shares with a spouse employed at the focal firm: a stylized example



Notes: This figure illustrates the case in which one employee's spouse is also employed at the focal firm itself. This self-link does not generate any exposure to external AI adoption and is excluded from the computation of the exposure measure.

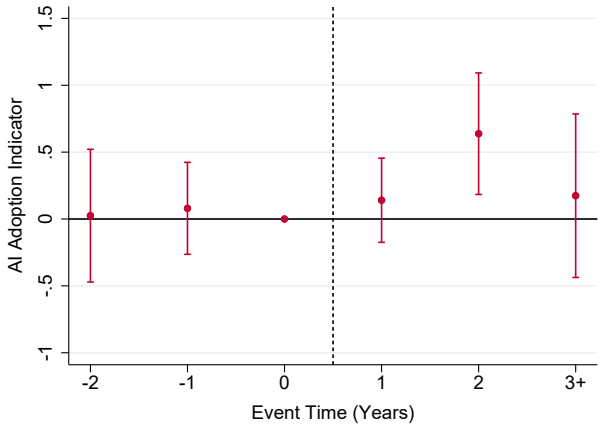
Appendix Figure A4. Exposure shares with an unlinked employee: a stylized example



Notes: This figure illustrates the case in which one employee does not have a spouse, and is therefore entirely absent from the spousal network. The measured exposure is computed over the subset of employees who have a partner working at another firm.

A.3 Additional Results

Appendix Figure A5. Effect of education network exposure on AI adoption



Notes: This figure plots the estimated coefficients $\hat{\delta}_t$ from equation (4), where the outcome is an indicator for AI adoption and the exposure measure is constructed using education network ties (co-enrollment in the same undergraduate cohort). The dashed vertical line marks the base year. Confidence intervals are at the 95% level.