

Menkhoff, Manuel

Working Paper

Belief Updating and AI Adoption: Experimental Evidence from Firms

CESifo Working Paper, No. 12291

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Menkhoff, Manuel (2025) : Belief Updating and AI Adoption: Experimental Evidence from Firms, CESifo Working Paper, No. 12291, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/334647>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CES ifo

**12291
2025**

November 2025

Working Papers

Belief Updating and AI Adoption: Experimental Evidence from Firms

Manuel Menkhoff

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Belief Updating and AI Adoption: Experimental Evidence from Firms

Manuel Menkhoff*

November, 2025

Abstract

Using a large German firm survey, I randomize information on documented AI productivity gains and industry adoption rates and track firms over time. Beliefs about AI's productivity potential rise significantly after the treatments across the prior distribution without reducing uncertainty. These treatment-induced belief shifts map into behavior: in firms where the respondent has high decision authority, AI adoption is more likely one year later. Information about competitor adoption has direct effects on actions: incumbent adopters cut prices, while not-yet adopters revise business expectations upward. Together, the results highlight the role of expectations, strategic considerations, and informational frictions in shaping technology diffusion and its macroeconomic impact.

JEL Classification: D22, D84, E22, E31, O33

Keywords: artificial intelligence, technological change, technology adoption, firm expectations, RCT, belief updating, price-setting

*University of Copenhagen, ifo Institute, CEBI, and CESifo, mmk@econ.ku.dk. I thank Rüdiger Bachmann, Lea Best, Ronnie Chatterji, Mathias Dolls, Oliver Falck, Ingrid Haegele, Matthias Heinz, Anders Humlum, Gernot Müller, Jens Oehlen, Morten Olsen, Andreas Peichl, Davud Rostam-Afschar, Chris Roth, Emilie Vestergaard, and Johannes Wohlfart for helpful comments. I am grateful to the survey department of the ifo Institute, in particular Katrin Demmelhuber, Felix Leiss, Agnesa Nimanaj, and Klaus Wohlrabe for their help and the opportunity to add new questions to the existing firm surveys. I thank the LMU-ifo Economics & Business Data Center (EBDC) for assistance with the data. The randomized control trial is registered at the AEA RCT Registry (AEARCTR-0013696). The activities of CEBI are financed by the Danish National Research Foundation, grant DNRFF134. Manuel Menkhoff is affiliated with the Danish Finance Institute (DFI) which financially supports researchers in Finance.

1 Introduction

The rise of artificial intelligence is a promising driver of economic growth potential. At the same time, the adoption of new technologies faces barriers such as information frictions and strategic considerations that operate through firms' beliefs (e.g., [Hall, 2006](#)). Understanding how decision makers in firms form expectations about potential gains from AI and how this shapes the probability of adoption and behavior is therefore of first-order importance. Beyond the specific case of AI, this perspective offers broader insights into how firms process and respond to technological news shocks (e.g., [Beaudry and Portier, 2006](#)).

While previous studies have documented the effects of AI on individual productivity, typically in single-firm settings (e.g., [Brynjolfsson, Li and Raymond, 2025](#)), we rarely observe technology-specific beliefs across many firms or exogenous shocks to those beliefs. Leading firm surveys elicit macro expectations rather than AI-specific productivity beliefs, and most AI studies randomize tools for workers or analyze a single firm's rollout. I address these gaps with a large, representative survey of firms that randomizes information about documented AI productivity gains and industry adoption rates, and then tracks belief updates, subsequent adoption, and other outcomes at the firm level over the following year.

This paper makes three key contributions. First, I show that firms' beliefs about the productivity potential of AI shift strongly when they receive new information about realized AI productivity gains or industry-specific adoption rates. The effects are particularly strong in the upper tail of firms' belief distribution and do not reduce uncertainty. Second, these shifts in beliefs translate into actual adoption: firms where the respondent has high decision authority or that have high internal financing capacity are significantly more likely to adopt AI one year later. Third, information on competitor adoption shapes strategic behavior: adopters cut prices consistent with anticipating stronger competition, while not-yet adopters revise business expectations upward, in line with technological-news-driven optimism.

In more detail, I implement a randomized control trial (RCT) in which one randomly selected group of firms receives information on AI-driven productivity gains documented in other studies, while another group is provided with information on the share of firms in their industry that already use AI or plan to adopt it. These two treatments test information frictions and strategic factors as adoption barriers. Before treatment, I elicit firms' beliefs about AI-driven macro productivity gains and about industry adoption rates. After the treatment, I elicit beliefs about the expected productivity gains from AI for their own business and track AI adoption one year later. For each belief, firms report a lower and an upper bound; the midpoint serves as the expectation, and the range measures subjective uncertainty.

This survey experiment is embedded in the Ifo Business Survey, a well-established, large-scale survey of German firms (Sauer, Schasching and Wohlrabe, 2023). The panel structure of the survey, along with its regular questions, allows me to track firms over time and analyze the effects of the treatments on a range of outcomes, including their business expectations, pricing and investment decisions, as well as the adoption of AI.

Information about productivity gains and competitors' adoption of AI has a significant and sizable effect on firms' beliefs about the productivity gains from AI for their own business. The effect is not present among firms that have already adopted, but arises strongly among firms that have not yet adopted AI. For example, while firms in the control group that currently discuss the implementation of AI expect an average productivity increase of 7.5% from AI over the next five years, this expectation rises by 1.4 percentage points (pp) for firms at the margin of adoption that received the productivity treatment and by 3.5 pp for those that received the competition treatment. This effect is driven by updating lower and, in particular, upper bound beliefs and thereby does not decrease firm-specific uncertainty about AI productivity potential.

The information treatment significantly increases AI adoption in firms where the respondent has high decision authority—i.e., is a top-level manager or the owner and the firm is a family firm (few decision makers). Treated firms that had only discussed AI adoption or had not considered it at all are 8 pp more likely to adopt AI one year later. The effect operates through changes in firms' beliefs about AI-related productivity gains and is particularly pronounced in firms that typically rely on internal funds for investment, consistent with greater flexibility to act quickly when an opportunity arises. Reassuringly, the effect is concentrated in industries where occupations with AI-exposed tasks—such as programming, writing, and customer service—are overrepresented.

The effect of the competition treatment on firms' business expectations and price decisions strongly depends on the AI adoption status of the firms. Firms that have already implemented AI tend to revise their business expectations downward and set lower prices after the competition treatment. For instance, treated firms that have already adopted AI are 5.4 pp less likely to increase prices over the following months. This effect is particularly driven by firms that initially expected a low AI adoption rate among competitors and had high uncertainty about it. This pattern is consistent with Bertrand-type competition under imperfect information about competitors' AI adoption, where an increase in expected competitive pressure leads to more pessimistic business expectations and more cautious price-setting.

Both the competition and productivity treatments increase business expectations over the following months for firms that have not yet considered adopting AI. These positive effects provide new evidence that technological news shocks influence firms' sentiment. Even though productivity gains may take years to materialize, these firms become more optimistic in the short run. Consistent with

this interpretation, the increase in business expectations is not driven solely by changes in firms' own AI-related productivity beliefs, but treated firms also report higher GDP growth expectations.

This paper is structured as follows. First, I relate the paper to the literature. Then, Section 2 introduces the survey and the RCT set-up, and Section 3 discusses descriptive patterns. In Section 4, I study the effects of the RCT on AI productivity beliefs, Section 5 turns to the effects on adoption, and Section 6 analyzes effects on sentiment and price-setting. Section 7 concludes.

Related literature This paper contributes to several strands of the literature. First, it relates to work on the adoption and diffusion of new technologies (e.g., [Griliches, 1957](#), [Hall, 2006](#), [Kalyani et al., 2025](#), [Rogers, 2003](#)). I add to this literature by directly eliciting firms' productivity beliefs about AI and study how firms update their beliefs when presented with information about expected productivity improvements and adoption patterns in their industry. The observed belief updating provides empirical evidence on the important role of information frictions as adoption barriers that operate through firms' beliefs. In addition, in line with the literature on social learning and herding ([Banerjee, 1992](#), [Bikhchandani, Hirshleifer and Welch, 1992](#), [Conley and Udry, 2010](#), [Young, 2009](#)), information about competitors' adoption triggers adoption.

Second, more concretely, the paper complements the growing and so far mostly descriptive literature on AI adoption (e.g., [Bick, Blandin and Deming, 2025](#), [Chatterji et al., 2025](#), [Humlum and Vestergaard, 2025b](#), [Licht and Wohlrabe, 2024](#), [McElheran et al., 2024](#)).¹ I show that firms in which the respondent has high decision authority are ultimately more likely to adopt AI following treatments on productivity improvements and on industry adoption. This complements [Humlum and Vestergaard \(2025b\)](#), who find that productivity treatments shift workers' beliefs but not their AI adoption, due to employer constraints and training needs, frictions that are less relevant for CEOs or owners. This underscores the central role of firms in the diffusion of new technologies.

Third, the paper connects to the growing literature on AI and firm behavior. For example, [Brynjolfsson, Li and Raymond \(2025\)](#) analyze the productivity effect of the introduction of an AI assistance for customer-support agents. [Otis et al. \(2023\)](#) study the impact of AI on entrepreneurial performance. [Babina et al. \(2024\)](#) show that firms that invest in AI have higher growth due to increased product innovation. This paper contributes to the literature by documenting productivity beliefs of AI and tracking firms' expectations and price-setting over time in response to randomized treatments that influence their beliefs about the productivity benefits for their business. Complementarily, [Brüll, Mäurer and Rostam-Afschar \(2025\)](#) provide experimental evidence from tax consulting firms and

¹There is also a rapidly growing literature on the labor market effects of AI, see e.g., [Brynjolfsson, Chandar and Chen \(2025\)](#), [Freund and Mann \(2025\)](#), [Hampole et al. \(2025\)](#), [Humlum and Vestergaard \(2025a\)](#).

show that correcting underestimation of automatability shifts beliefs upward and raises expected profits, without immediate adjustments in wages or number of employees.

Fourth, more broadly, the paper connects to two strands of the macroeconomic literature. First, it is related to the literature on technological news shocks (e.g., Barsky, Basu and Lee, 2015, Beaudry and Portier, 2006, Cascaldi-Garcia and Vukotic, 2022, Enders, Kleemann and Müller, 2021, Schmitt-Grohé and Uribe, 2012). The setup in this paper offers a clean firm-level test of technology-news shocks, complementing existing time-series evidence. Business expectations rise after treatment among firms that had not previously considered adopting AI. Second, methodologically, this paper is inspired by the empirical macro literature that applies RCTs—typically about inflation or growth outlooks—in firm surveys to understand the updating behavior of managers (e.g., Coibion, Gorodnichenko and Kumar, 2018, Coibion, Gorodnichenko and Ropele, 2020). A novel contribution of this paper is that I use AI adoption rates in firms’ industries as a treatment to shift managers’ beliefs not only about the potential benefits of this technology, but also about the competitive pressure associated with AI adoption. Relatedly, Bursztyjn et al. (2025) show that demand for AI-based learning tools increases with the share of other teenagers using the technology. In their setting, however, the social information mainly operates through social pressure rather than social learning, as higher peer adoption has little effect on beliefs about AI quality. This also relates to the household literature that studies belief updating about relative standing after providing targeted information on peers’ income (Hvidberg, Kreiner and Stantcheva, 2023, van Rooij et al., 2024).

2 Data and survey design

In this section, I introduce the firm survey and the survey experiment.

ifo Business Survey I use the ifo Business Survey, a long-running monthly panel of roughly 6,000 German firms, which is widely used in economic research (e.g., Bachmann et al., 2025, Enders, Hünnekes and Müller, 2022, Link et al., 2023). The main product of the survey is a business sentiment indicator for Germany—the widely recognized ifo Business Climate Index—which is cited by Bloomberg as one of the “12 Global Economic Indicators to Watch” and among the top 20 macroeconomic news events influencing global asset prices (Kerssenfischer and Schmeling, 2024). Administered by the ifo Institute, the survey benefits from established relationships with participating firms and above-average response rates (Born et al., 2023); many firms have participated in over 100 waves. This makes the survey a natural vehicle for conveying credible information.

The survey covers a broad set of industries and a wide range of firm sizes and is designed to be representative of German firms (Hiersemenzel, Sauer and Wohlrabe, 2022). The median firm in

the sample has 36 employees (p10: 3, p90: 350) and annual revenues of 7 million euros (p10: 240 thousand, p90: 97 million). 40% of firms are in services, 36% in manufacturing, and 24% in wholesale and retail trade. 12% of firms are located in former East Germany.

A key advantage of the firm survey is that the questions are typically answered by the actual decision-maker (Hennrich, Sauer and Wohlrabe, 2023). In parts of the analysis, I selectively focus on a subsample of respondents with high decision authority to tease out the effects of the information treatments on adoption and behavior. For this, I use items indicating who answered the survey and whether the firm is family-owned.

In June 2023, June 2024, and May 2025, there is a special question on the adoption of AI in the respondent's own firm. The question asked for the status of AI adoption, whether it is (i) used, (ii) planned, (iii) discussed, or (iv) no topic.

Randomized information treatments In June 2024, I embedded a randomized control trial (RCT) within the survey. Specifically, I randomly allocated firms into three groups, stratified within major industries: manufacturing, services, and retail/wholesale. The first group serves as a control group and received no information. The other two groups each received a different piece of information. Appendix Table A1 shows that group assignment cannot be predicted by observable firm characteristics.

The second group receives information on productivity gains. In particular, I refer to well-known papers in the academic literature—also used in the macroeconomic productivity analysis of Acemoglu (2025)—that estimate productivity gains in the context of programming (Peng et al., 2023), writing (Noy and Zhang, 2023), and customer service (Brynjolfsson, Li and Raymond, 2025). Additionally, I provide hyperlinks to these papers. The treatment reads as follows:

Productivity Treatment (Group 2): *The use of AI can lead to significant time and cost savings, as shown by recognized scientific studies, some of which document large productivity gains through AI (e.g. 56% time savings in programming, 40% time savings in writing texts or 14% time savings in customer service).*

This treatment is designed to reduce missing information about the gains from AI and, in this setting, provides a causal test of information frictions as a barrier to adoption (Hall, 2006). By doing so, it may increase expected profitability and lower uncertainty about this technology.

The third group receives tailored information on the adoption of AI in their broadly defined industry. Specifically, I use data from the AI adoption question in the ifo survey conducted in June 2023. The treatment is presented as follows:

Competition Treatment (Group 3): *A year ago, by the way, around [30%/19%/19%] of firms in the German [manufacturing/services/retail&wholesale] industry stated in the ifo survey that they were using AI or were already planning to do so.*

This treatment is intended to convey to firms that AI is relevant for their own business. It provides industry-level adoption information to make the competitive relevance of AI salient for the firm. In addition, it may trigger responses consistent with social learning and competitive considerations (herd behavior), which are known to matter for technology adoption (Banerjee, 1992, Bikhchandani, Hirshleifer and Welch, 1992, Conley and Udry, 2010, Young, 2009).

To gauge how surprising the competition treatment is and to capture firms' uncertainty about competitors' AI adoption, I elicit prior beliefs as a range before the treatment. Specifically, firms report a lower and an upper bound for the current adoption share in their (broad) industry:

Competition prior (only Group 3): *What do you think is the current share of companies in the German [manufacturing/services/retail&wholesale] industry that are already using or planning to use artificial intelligence (AI)? Between ___ % and ___ %*

Survey questions on productivity beliefs Before the treatments, I elicit firms' beliefs about the macroeconomic productivity potential of AI. I use a five-year horizon to capture medium-term beliefs, since implementation can take time. To measure uncertainty, all firms report a range for expected productivity gains. The question reads as follows:

Prior: *What do you think is the overall economic productivity potential of artificial intelligence (AI) in the next 5 years?*

- *No effect on productivity*
- *Increase in productivity between ___ % and ___ %*

After the treatments, I elicit firms' beliefs about the productivity potential of AI for their own business. The question format is exactly the same:

Posterior: *What do you think is the productivity potential of AI in your company in the next 5 years?*

- *No effect on productivity*
- *Increase in productivity between ___ % and ___ %*

To calculate expected productivity gains, I take the average between the lower and upper bound value. The productivity beliefs are winsorized at the 95th percentile. When I study the uncertainty

that is associated with those beliefs, I consider the span from the lower to the upper bound value. This follows the elicitation approach of [Bachmann et al. \(2025\)](#), who also ask firms to state lower and upper bounds—in their case for expected sales growth—to capture subjective uncertainty.

3 Descriptive analysis

In this section, I present descriptive statistics on firms' AI adoption and their productivity beliefs about AI.

AI adoption dynamics Figure 1 presents the distribution of firms' AI adoption status as reported in June 2023, June 2024, and May 2025. In June 2023, 13% of firms had adopted AI, while for 46% of firms, AI adoption was not yet a consideration. By June 2024, the adoption rate increased to 24%, and the share of firms that did not consider AI at all declined to 27%. In May 2025, more than one third of firms reported using AI. Hence, AI adoption among firms was highly dynamic over the period studied. For more detailed summary statistics, see [Licht and Wohlrabe \(2024\)](#).

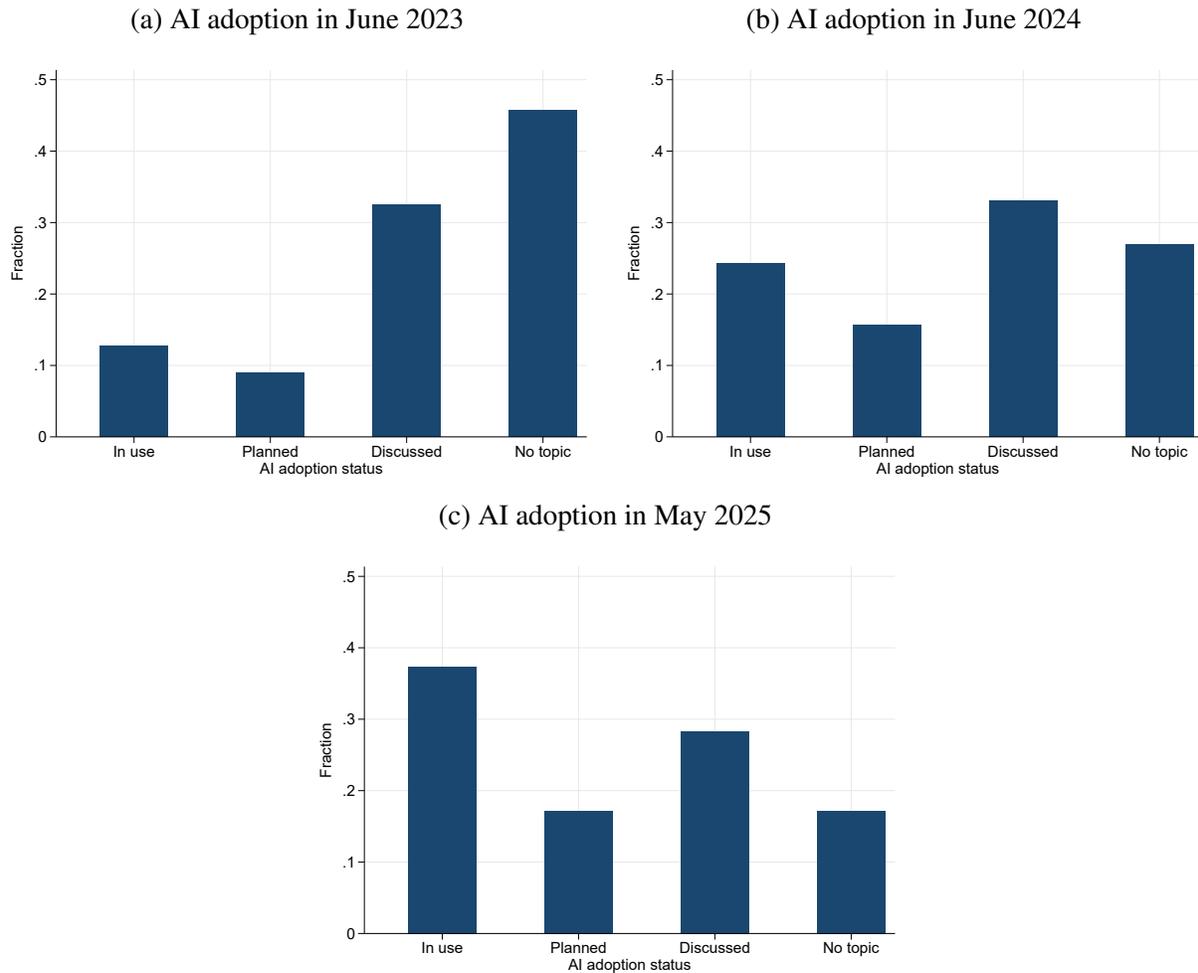
For the analysis of treatment effects by adoption status below, I use firms' AI adoption status from June 2023 to mitigate endogeneity concerns and ensure that adoption decisions were finalized before measuring their impact. When this information is missing, I supplement it with responses from June 2024, collected prior to the treatments, to maximize the sample size.

Firms' beliefs about AI productivity gains The panels of Figure 2 depict the distribution of firms' beliefs in the control group about AI's productivity potential for their own business (left panel) and the macroeconomy (right panel) over the next five years. These beliefs vary widely across firms. On average, firms' productivity expectations for their own business are lower (mean = 7%) than for the aggregate economy (mean = 11%). Close to 40% of firms anticipate no productivity benefits from AI over this period.

Firms' productivity expectations for their own business are strongly associated with their beliefs about AI's impact on the broader economy ($\rho = 0.66$). This observed correlation suggests that firms may extrapolate from their own experience to the broader economy, as modeled in macroeconomic island frameworks ([Lorenzoni, 2009](#), [Lucas, 1972](#)). This finding is consistent with a high degree of uncertainty about AI-driven productivity gains.

Appendix Figure A1 illustrates the distribution of the span between firms' lower and upper bound estimates, a direct measure of subjective uncertainty. There is substantial variation in subjective

Figure 1: Firms' AI adoption status

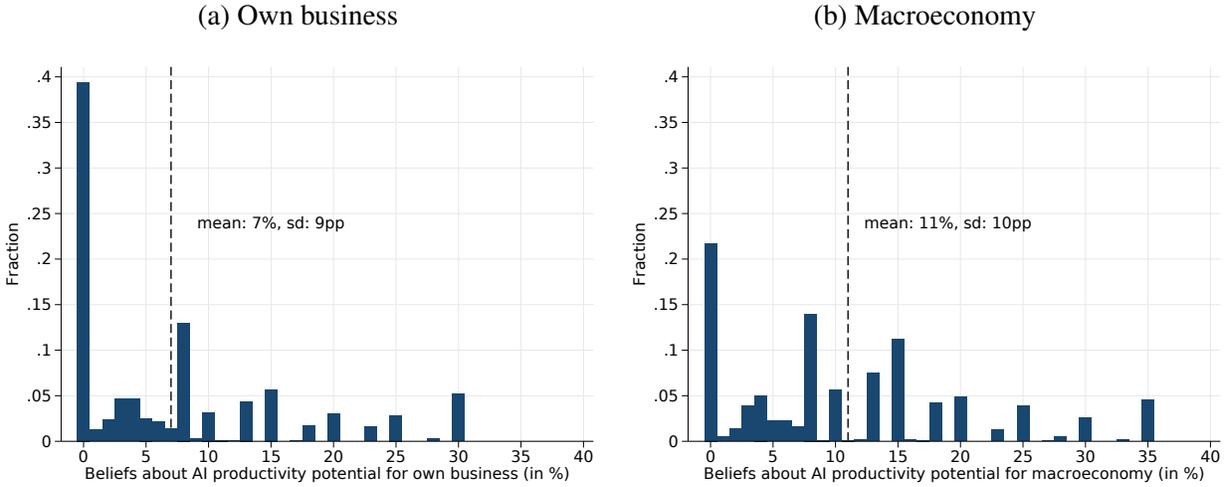


Notes: The figure shows the distribution of firms' adoption status elicited in June 2023, June 2024, and May 2025.

uncertainty regarding AI's productivity effects, with uncertainty being higher for the macroeconomy (mean span = 10.4 pp) than for firms' own business (mean span = 8.4 pp).

Validation of productivity beliefs I validate firms' productivity beliefs by constructing an AI exposure index based on the three studies used in the productivity treatment. For each study, I identify the most strongly affected ISCO-08 three-digit occupation and map it to industries. Specifically, the strong productivity gains in programming primarily affect software and applications developers and analysts (ISCO 251), the productivity gains in writing are most relevant for authors, journalists and linguists (ISCO 264), and the productivity gains in customer service particularly concern client information workers / customer service clerks (ISCO 422).

Figure 2: Firms' AI productivity beliefs



Notes: The figure shows histograms of firms' AI productivity beliefs for their own business (left panel) and for the aggregate economy (right panel). These beliefs are elicited in terms of ranges in June 2024. The average between the lower and upper bound value is used. The sample is restricted to the control group. The survey questions are in the text.

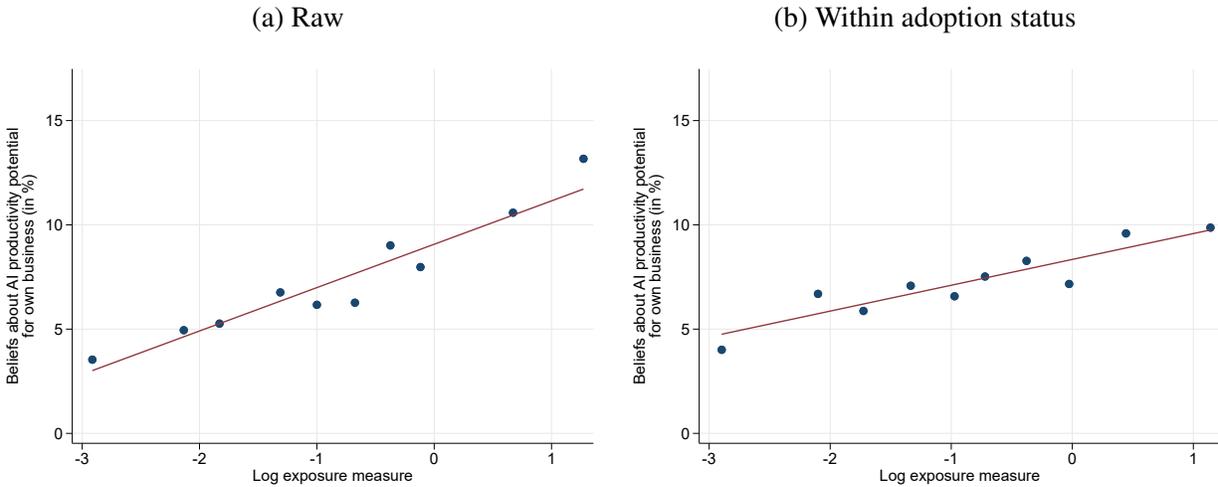
I then use a matrix provided by the German Federal Employment Agency (Bundesagentur für Arbeit) that includes the number of workers for each industry (WZ08 4-digit)–occupation (ISCO-08 3-digit) combination in 2024 to match the occupation data to the survey, where I only observe industry information. I then assess whether the respective three occupations are under- or overrepresented in a given industry by computing the probability that workers in a given occupation are employed in an industry relative to the unconditional probability of being employed in that industry:

$$\text{LogExposure} = \ln \left(\frac{P(I | O)}{P(I)} \right). \quad (1)$$

When taking logs, a value smaller (larger) than zero indicates that the occupation is under-(over)represented in an industry. I then take the average of this exposure measure over the three occupations to derive a single index. In the adoption analysis, I also consider an indicator that equals one if at least one of the three occupations is overrepresented in a 4-digit industry.

The elicited productivity beliefs of managers for their own firm are strongly predicted by variation in the exposure index across four-digit industries. Figure 3 displays a binned scatterplot of firms' expected productivity gains for their own business against the average log exposure measure over the three occupations described above. I restrict the sample to firms in the control group. Firms in industries where AI-exposed occupations are particularly overrepresented expect, on average, productivity gains of 10% or more over the next five years, whereas firms in low-exposure industries expect around 5%. This strong relationship is partly explained by adoption status, but even conditional on adoption the relationship remains significant, as shown in Panel (b). This

Figure 3: Relationship between firms' AI productivity beliefs and AI industry exposure



Notes: The figure shows binned scatter plots between firms' AI productivity beliefs and an AI exposure measure at the 4-digit NACE industry level. The sample is restricted to the control group. The AI exposure measure represents an over/under-representativeness (larger/smaller than 0) of occupations that have been shown to have strong AI productivity benefits: software and applications developers and analysts (ISCO 251), journalists and linguists (ISCO 264), and client information workers / Customer service clerks (ISCO 422). Details are in the text. Panel (a) shows the raw relations. In Panel (b), it is controlled for AI adoption status fixed effects.

indicates that the elicited subjective beliefs of managers about AI productivity across firms contain meaningful variation.

As an alternative validation check, I consider the AI industry exposure (AIIE) measure from [Felten, Raj and Seamans \(2021\)](#) at the 2-digit NACE industry level. This measure links common AI applications to occupational abilities based on a crowd-sourced data and then aggregates them at the industry level. Notably, the measure itself is agnostic regarding whether AI complements or substitutes labor. Appendix Figure A2 presents binned scatter plots of firms' productivity beliefs for firms in the control group against the AIIE. Again, I find that firms' productivity beliefs for their own business are systematically related to the AI industry exposure measure, indicating that these beliefs contain meaningful variation, driven by expectations of higher labor productivity.

Beliefs on AI adoption of competitors To what extent do firms think that other firms in their industry have already adopted AI? Appendix Figure A3 plots the distribution of expected adoption rates (defined as the share of firms that already use AI or plan to use it) and associated uncertainty about it (span between lower and upper bound estimate). Firms strongly disagree on expected adoption rates with an interquartile range spanning from 10% to 35%. They also show high uncertainty about it, as the most often reported ranges amount to 10 and 20 pp. Consistent with imperfect knowledge about competitors' adoption, the adoption rate beliefs are systematically related to firms' own adoption status: if firms already use AI, they expect the adoption rate to be 15

pp higher compared to firms, where AI adoption is no topic yet, as shown in Appendix Figure A4. Later, I exploit heterogeneity in those prior beliefs on adoption rates—conditional on adoption status—to study the effects of the competition treatment on firms’ price setting.

4 Effects of information treatments on AI productivity beliefs

The information treatments significantly raise firms’ beliefs about AI’s productivity potential for their own business. Figure 4 plots prior beliefs about AI-driven productivity gains for the macroeconomy over the next five years against posterior beliefs about AI’s productivity gains for their own business, split by treatment groups. The left panel shows firms that are planning, discussing, or not currently considering adoption; the right panel shows firms that already use AI.

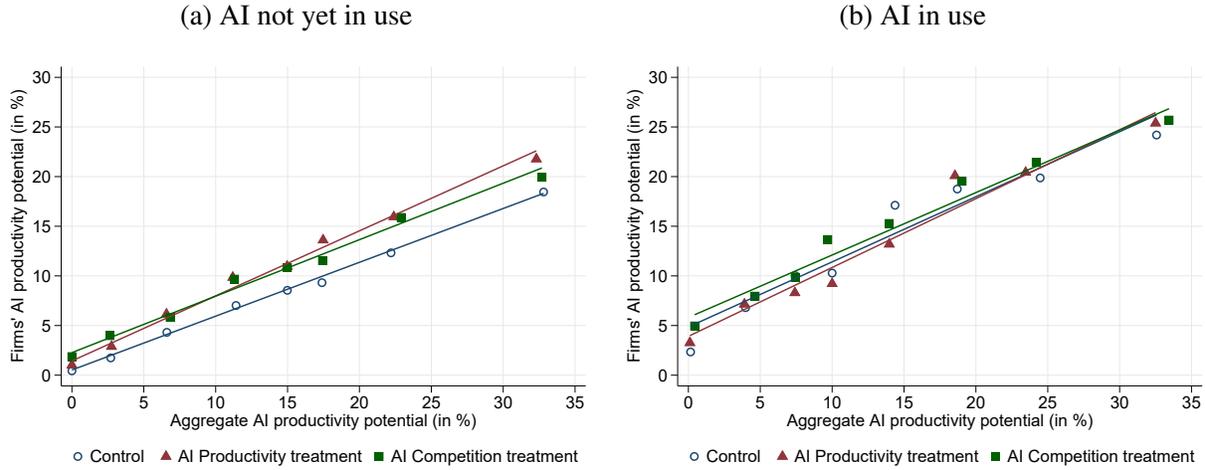
For firms that do not yet use AI, the productivity and competition treatments systematically shift AI productivity beliefs upward. Interestingly, the updating behavior deviates from Bayesian updating: it is not the case that firms with a highly optimistic prior (regarding aggregate AI productivity potential) update less or even downward while only those with a more moderate prior adjust upward. Instead, the upward shift is of similar magnitude across the entire prior distribution. This pattern is consistent with the randomized treatments conveying credible firm-level relevance, i.e., that AI matters not only in general but also for their own business, inducing across-the-board (non-Bayesian) upward revisions.

In contrast, for firms already using AI, the treatments leave productivity beliefs unchanged. This is consistent with these firms being well informed about AI’s productivity effects, so additional information does not shift beliefs. This corroborates the treatment design and suggests that experimenter-demand effects are unlikely.

Table 1 presents the average differences between the treatment groups in a regression framework and separately by prior adoption status. Across all firms, the productivity treatment increases beliefs about AI productivity gains for their own business by 1.2 pp, while the competition treatment raises them by almost 3 pp (Column 1). For both treatments, the effects are insignificant for firms that already use AI (Column 2). Among not-yet adopters, the productivity treatment raises productivity beliefs by 1.4–2.2 pp, and the competition treatment by 3.2–4.1 pp. Relative to the control group, effects are largest for firms that have not yet considered AI: the productivity treatment nearly doubles perceived productivity potential, and the competition treatment more than doubles it (control mean: 2.6%), see Column 5.

Next, I test for heterogeneity among firms that have not yet adopted AI. Overall, treatment effects are quite homogeneous across firm characteristics, see Appendix Table A2, where I interact the

Figure 4: Effects of information treatments on beliefs about firms' own productivity beliefs



Notes: The figure shows binned scatter plots of firms' beliefs about the productivity gains from AI for the macroeconomy (before the treatments) and for their own business (after the treatments). The relationship is shown separately by AI adoption status (left panel for firms that plan / discuss the use of AI or AI is no topic; right panel for firms that already use AI) and treatment groups: control group (blue dots), productivity treatment group (red triangle), and competition treatment group (green square). The survey questions are shown in the text.

treatment dummies with firm characteristics and additionally control for adoption group (planned / discussed / no topic) and prior productivity beliefs about AI's macroeconomic potential. Updating is broadly similar across firm size and, if anything, the competition treatment has slightly larger effects for large firms (more than 50 employees), consistent with a higher relevance of competitors' adoption decisions for these firms. There is no significant heterogeneity with respect to the manager's education, the use of external funds for investment, or the respondent's role (i.e., whether the respondent has high decision authority).

Firms in high-exposure industries—where at least one of the three occupations targeted by the treatment (programming, writing, customer service) is overrepresented at the 4-digit industry level—do not update their AI productivity beliefs more. A plausible explanation is that these industries both (i) tend to benefit more from AI (as shown above and by the positive, significant coefficient on the non-interacted exposure indicator) and (ii) were better informed ex ante. Hence, a weak response to information during the information provision experiment is a sign of a high importance of being informed about it outside the survey, see [Mackowiak and Wiederholt \(2025\)](#) for a formalization of rational inattention during RCTs. These forces appear to offset each other. I revisit exposure when analyzing the probability of adoption in the next section. Finally, firms that are located in East Germany show significantly weaker updating in response to the productivity treatment, which is consistent with lower trust in public signals.

Table 1: Treatment effects on beliefs about productivity gains of AI for own business

	Beliefs about AI productivity potential for own business				
	(1)	(2)	(3)	(4)	(5)
	All firms	By AI adoption status:			
		In use	Planned	Discussed	No topic
Productivity treatment	1.24*** (0.34)	-0.27 (0.97)	2.16** (1.05)	1.40** (0.57)	1.95*** (0.38)
Competition treatment	2.95*** (0.36)	1.03 (1.02)	4.10*** (1.00)	3.45*** (0.57)	3.20*** (0.45)
Constant	6.94*** (0.37)	14.1*** (0.87)	11.5*** (0.75)	7.53*** (0.43)	2.61*** (0.26)
Observations	4288	579	473	1451	1785

Notes: The table displays the treatment effects on firms’ beliefs about productivity gains of AI for their own business. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. The treatment effect is shown separately by AI adoption status in Columns 2 to 5. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A key takeaway from the analysis is that providing firms with information on (i) the benefits of new technologies and (ii) adoption rates within their industry increases their perceived benefits of adopting the technology for their own business. This finding provides empirical evidence that both imperfect information (e.g., [Hall, 2006](#)) as well as social learning and competitive considerations (herd behavior) (e.g., [Banerjee, 1992](#)) are important barriers to technology adoption that can be mitigated by information treatments about the technology and about others’ adoption, respectively.

Uncertainty about productivity gains Rather than reducing uncertainty, the treatments shift the entire distribution of perceived productivity gains to the right, with the largest movement in the upper tail. Table 2 repeats the analysis of Table 1, but now, instead of using average productivity beliefs as the outcome, it looks at the lower- and upper-bound estimates separately. The sample is smaller because firms that report “no effect on productivity” do not provide lower and upper bounds. On average, the lower bound increases by 0.8 pp under the productivity treatment and by 2.2 pp under the competition treatment. The effects on the upper bound are even stronger, amounting to 1.5 pp and 3.0 pp, respectively. The stronger effect on the upper compared to the lower bound is robust across AI-adoption status among not-yet adopters. Hence, subjective uncertainty—measured by the span between the lower and upper bounds—does not decline after the information provision. Nonetheless, greater upside potential can spur adoption, via both a shift in the first moment and, in particular, an increase in upside risk, consistent with models of R&D investment under ambiguity ([Beauchêne, 2019](#)).

Table 2: Treatment effects on lower/upper bound beliefs about productivity gains of AI for own business

	(1)	(2)	(3)	(4)	(5)
	All firms	In use	By AI adoption status: Planned	Discussed	No topic
<i>Panel A: Lower bound (minimum)</i>					
Productivity treatment	0.77** (0.31)	-0.27 (0.85)	1.20 (0.90)	1.19** (0.51)	1.93*** (0.48)
Competition treatment	2.21*** (0.36)	0.46 (0.94)	3.31*** (0.96)	2.45*** (0.50)	3.63*** (0.53)
<i>Panel B: Upper bound (maximum)</i>					
Productivity treatment	1.46** (0.64)	-0.054 (1.52)	3.30** (1.65)	2.26*** (0.82)	2.53*** (0.95)
Competition treatment	2.98*** (0.61)	-0.54 (1.71)	5.12*** (1.50)	3.81*** (0.84)	4.89*** (0.87)
Observations	2822	522	428	1143	729

Notes: The table displays the treatment effects on firms' beliefs about lower bound (Panel A) and upper bound (Panel B) productivity gains of AI for their own business. "Productivity treatment" and "Competition treatment" are dummies indicating the treatment groups as introduced in the text. The treatment effect is shown separately by AI adoption status in Columns 2 to 5. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Effects of information treatments on AI adoption

The fact that firms significantly increase their productivity beliefs following the treatments, as documented in the previous section, suggests that they are more likely to adopt AI. In the following analysis, I examine to what extent the information treatments ultimately affect the probability of AI adoption and whether this operates through firms' beliefs. Specifically, I focus on firms that had either only discussed adopting AI or for which AI was not yet a topic. I measure AI adoption using an indicator variable equal to one if the firm reports using AI in the May 2025 survey wave. I then regress this indicator on a treatment dummy, include 4-digit industry fixed effects, and cluster standard errors at the industry level.

Identifying treatment effects on actual AI use is demanding: it requires detecting a statistically significant change in firms' adoption behavior resulting from brief information treatments delivered one year earlier. In addition, the sample size declines due to attrition, and many other factors may have influenced firms' adoption decisions in the interim. Therefore, I focus on firms where the

Table 3: Treatment effects on AI adoption

	(1)	(2)	(3)	(4)	(5)
	OLS			IV	
Both treatments	0.017 (0.022)	0.082*** (0.030)	0.080*** (0.030)		
Productivity treatment				0.098*** (0.037)	
Competition treatment				0.062* (0.035)	
AI productivity beliefs (own business)					0.034** (0.013)
Observations	2376	1230	1230	1230	1230
F-Stat					19
High decision authority sample	no	yes	yes	yes	yes
AI productivity beliefs ctrl (macro)	no	no	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes

Notes: The table displays the treatment effects on firms’ probability to adopt AI in May 2025. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. In Column 5, firms’ AI productivity beliefs about their own business are instrumented with the two treatment dummies. Columns 2-5 focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

survey respondent arguably has high decision-making authority in the firm. Specifically, I classify a respondent as having “high decision authority” if both of the following hold: (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned, which typically implies more centralized decision-making (Bertrand and Schoar, 2006).

Treated firms where the survey respondent has high decision-making authority, are significantly more likely to adopt AI one year later. While this relationship is not statistically significant in the full sample (see Column 1 of Table 3), it becomes highly significant in this subsample (Column 2). The coefficient remains stable when controlling for firms’ prior beliefs about macroeconomic AI productivity gains (Column 3) and is driven by both treatments, particularly the productivity treatment, as shown in Column 4.

The effect is sizable in magnitude. Firms with high decision-making authority in the treatment group are 8 pp more likely to adopt AI compared to those in the control group. Given that the baseline probability of AI adoption in the control group is 21%, this corresponds to an increase of about 40% for treated firms.

The last column of Table 3 shows that the higher probability of AI adoption operates through changes in firms' productivity beliefs. I implement an instrumental variables (IV) approach, using the two treatment dummies to instrument for firms' productivity beliefs, thereby generating exogenous variation (first-stage F-statistic = 19). This exogenous variation in beliefs, in turn, leads to a higher probability of adopting AI. Specifically, a 1 pp increase in AI productivity beliefs increases the probability of adoption by 3.4 pp among firms with high decision-making authority.

The significant effects on adoption after just one year suggest that firms typically rely on off-the-shelf AI solutions that require little customization and can be implemented quickly. Supporting this interpretation, [Hampole et al. \(2025\)](#) show that real-time business intelligence tools—automated, data-driven decision aids—are the most common AI application across industries. Likewise, [Chatterji et al. \(2025\)](#) report that, within ChatGPT, workers primarily use it for writing, consistent with reliance on readily available, easy-to-integrate chat-based applications. That the effect on adoption is significant only when respondents have high decision-making authority further suggests that adoption depends on organizational environments that facilitate rapid implementation.

Heterogeneity in AI adoption by exposure The effects on adoption of the productivity treatment are predominantly present for firms where at least one of three occupations, which are covered in the treatment (programming, writing, customer service), is overrepresented in their 4-digit industry. Table 4 presents the regression results, where the treatments are interacted with an indicator for whether at least one of these three occupations is overrepresented. In the full sample, firms that are particularly exposed to the productivity treatment show a weakly significant effect (Column 1). Among respondents with high decision-making authority, firms in these exposed industries are 16 pp more likely to adopt after the treatments (Column 2), whereas firms in other industries show no significant difference relative to control firms in the same industry. The difference is even more pronounced for the productivity treatment, where I estimate a 25 pp higher probability of adoption (Column 3). Also for the competition treatment, the effect is statistically significant only for firms in exposed industries, yet it is not statistically different from the effect for firms in low-exposure industries (Column 4). Instrumenting productivity beliefs with the treatments, I find that higher productivity beliefs translate into a higher probability of adoption one year later only among firms in exposed industries (Column 5), with no pass-through among firms in low-exposure industries (Column 6).

The heterogeneity by AI exposure corroborates the baseline finding. In particular, high-exposure firms—those most likely to benefit given the productivity information—are the ones that adopt after the intervention. Interestingly, productivity beliefs become more optimistic for all firms, regardless of exposure. The gap between broad belief updates and concentrated adoption likely reflects implementation lags among low-exposure firms. Consistent with this interpretation, coding

and writing use cases face lower adoption frictions because off-the-shelf tools enable quick uptake among high-exposure firms. A related explanation for the divergence is short-lived post-treatment optimism that recedes once practical hurdles become salient.

Heterogeneity in AI adoption by financing source The treatment effects are particularly pronounced for firms that typically invest using internal funds. Appendix Table A3 presents results splitting firms by whether they invested using only internal funds or also partly external funds in the last year.² For firms that relied solely on internal funds, the productivity and competition treatments increase the probability of AI adoption by 15 and 11 pp, respectively (Column 1). The results remain significant when considering the full sample (Column 3), not only firms with high decision-making authority.

Heterogeneity analysis by firm size and investment volume (Appendix Table A4) suggests that neither small firms nor a low investment volume explains the heterogeneity by investment financing source. This finding is in line with the interpretation that firms relying on internal funds can act more quickly when an opportunity arises, as they can commit funds immediately without external approval. The flexibility to implement value-enhancing investments is a central motive for holding cash in corporate finance (Froot, Scharfstein and Stein, 1993, Opler, Pinkowitz and Williamson, 1999).

Investment plans in software An alternative concrete measure of firms' AI implementation is their investment in technologies that support and enhance its performance. A regular question in the November wave (five months after the treatments) of the survey asks about firms' investment plans in software and databases for the following year. This question is qualitative and framed in terms of planned changes relative to the current year (increase, no change, decrease). To assess the magnitude, the variables are standardized (z-score), enabling interpretation in terms of changes measured in standard deviations. I control for firms' business situation to reduce noise in the estimates and control for firms' total investment plans (same qualitative format) to account for overall investment plans.

Firms with high decision-making authority that receive the treatments are significantly more likely to plan disproportionate investment in software and databases over the next year. Appendix Table A5 presents the regression results. As with the adoption outcome, I only observe a significant effect for firms with high decision-making authority. The treatments increase the likelihood of investing in software and databases by 0.12 standard deviations (Column 2); mostly driven by the productivity treatment (Column 4).

²Firms that only used internal funds have, on average, 54% higher cash-to-asset ratios (26% vs. 17%), based on a one-time survey question fielded in 2020.

Table 4: Treatment effects on AI adoption by exposure

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS				IV	
	both		productivity	competition	both	
					AI exposure:	
					high	low
Treatment ×						
low AI exposure	-0.037 (0.031)	0.015 (0.042)	0.013 (0.047)	0.029 (0.051)		
high AI exposure	0.058* (0.029)	0.16*** (0.041)	0.25*** (0.055)	0.087** (0.044)		
AI productivity beliefs (own business)					0.077*** (0.024)	0.0062 (0.018)
Observations	2147	1117	718	740	415	702
P-value	0.03	0.01	< 0.01	0.39		
F-Stat					8.5	7.9
High decision authority sample	no	yes	yes	yes	yes	yes
AI prod. beliefs ctrl (macro)	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes

Notes: The table displays the treatment effects on firms’ probability to adopt AI in May 2025. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. The treatments are interacted with an indicator for whether at least one of the three occupations—programming, writing, or customer service—for which the productivity treatment provided information on AI productivity gains is overrepresented in the firm’s industry. In Columns 5 and 6, firms’ AI productivity beliefs about their own business are instrumented with the two treatment dummies. Columns 2-6 focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6 Effects of information treatments on business expectations and price-setting

The panel dimension of the survey allows me to track treatment effects on two key firm-level outcomes after the intervention: (i) business expectations and (ii) realized price-setting decisions.

Business expectations provide a measure of sentiment and allow me to study whether anticipated productivity gains—either for the firm itself or for the wider economy—make firms more optimistic in the short run. This setting offers a clean case study of technological news shocks and contributes to the literature that uses structural models and time-series methods to analyze the macroeconomic role of such shocks.

Beyond beliefs, the survey also captures realized price changes. Ex ante, the treatments' impact on price setting is ambiguous. AI can lower marginal costs and increase competitive pressure for incumbents, both pushing prices downward. At the same time, technological news shocks can be inflationary in models where demand rises before productivity materializes, putting upward pressure on prices. As a result, the net effect on price-setting depends on the relative strength of these opposing forces.

To estimate the effects of the information treatments on firms' business expectations and price-setting, I exploit the panel dimension of the ifo Business Survey. The two variables are elicited monthly as part of the core questionnaire. I consider the remaining six months of 2024 following the treatment in June. To improve the precision of the estimates, I include 4-digit NACE industry fixed effects and month fixed effects. Standard errors are clustered at the firm level to allow for serial correlation and heteroskedasticity in the error terms within firms over time.

The outcome variables for business expectations and price-setting are ordinal measures coded as -1 (decrease), 0 (no change), and 1 (increase). While OLS treats these as continuous variables—implicitly assuming equal spacing between categories—it provides a straightforward estimate of the treatment effect. To assess the magnitude, the variables are standardized (z-score), enabling interpretation in terms of changes measured in standard deviations.

6.1 Business expectations

Do the AI productivity and competition treatments shape firms' business sentiment? In Table 5, I analyze the treatment effects on firms' business expectations. This is a core monthly question that asks whether firms expect business conditions to improve, remain the same, or deteriorate over the

Table 5: Treatment effects on business expectations

	Business expectations (z-score, Jul–Dec 2024)				
	(1)	(2)	(3)	(4)	(5)
	All firms	By AI adoption status:			
		In use	Planned	Discussed	No topic
Productivity treatment	-0.0032 (0.030)	-0.16* (0.093)	-0.052 (0.11)	0.016 (0.054)	0.078* (0.045)
Competition treatment	0.057** (0.029)	-0.15* (0.091)	0.11 (0.11)	0.093* (0.051)	0.13*** (0.044)
Constant	-0.064*** (0.020)	0.24*** (0.056)	0.035 (0.068)	-0.063* (0.035)	-0.21*** (0.030)
Observations	22071	2833	2377	7471	9390

Notes: The table presents the treatment effects on firms’ business expectations for the next six months in the pooled sample from July to December 2024 controlling for 4-digit NACE industry fixed effects and month fixed effects. Business expectations are measured monthly using a qualitative question with three possible responses: rather better, no change, rather worse. The variable is standardized as a z-score. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. The treatment effect is shown separately by AI adoption status in Columns 2 to 5. Clustered standard errors at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

next six months. It constitutes the forward-looking half of the business cycle indicator, the “ifo Index,” which is typically used as a sentiment indicator.

Firms for which AI implementation is not a topic have significantly higher business expectations in the months following the treatments. The effect is largest for the competition treatment, where business expectations increase by 0.13 standard deviations (Column 5). Among firms currently discussing AI implementation, the effects also tend to be positive and are significant at the 10% level for the competition treatment. The point estimates tend to be negative for firms that have already implemented AI. Appendix Table A6 shows that the increase in business expectations among firms not yet considering AI is particularly strong when respondents have high decision-making authority and is driven by fewer negative business expectations.

In Table 6, I study the effects on business expectations for firms that have not yet considered AI in more detail, both in the full sample (Panel A) and among firms where the respondent has high decision authority (Panel B). Column 1 again shows a positive effect of the treatments on firms’ sentiment. Column 2 reports IV estimates where productivity beliefs are instrumented with the treatments; the results indicate that changes in business expectations operate through changes in AI productivity beliefs. Controlling for posterior AI productivity beliefs (Column 3) attenuates—but

does not eliminate—the treatment effect on business expectations. This pattern is consistent with beliefs mediating part of the effect, with additional channels (e.g., broader demand) leaving a residual treatment impact.

In line with an improved aggregate outlook that may raise expected product demand, firms also report higher real GDP growth expectations after treatment (Appendix Table A7). In particular, treated firms that have not yet considered AI report 0.13 pp higher expected real GDP growth for 2025, as elicited in November 2024. Inflation expectations for 2025 are unaffected.

Lastly, I exploit that for a large share of firms I observe AI adoption status one year after the treatment. Restricting to firms that had not adopted by May 2025 ensures the effect on business expectations is anticipatory rather than driven by rapid implementation (see Column 4, Table 6). I still find positive, significant effects, particularly in the high decision authority sample.

The findings establish a significant role of anticipatory productivity effects on firms' sentiment. Technological news shocks—revisions to expected future productivity without contemporaneous productivity gains—predict a short-run rise in sentiment. This is similar in spirit to [Schmitt-Grohé and Uribe \(2012\)](#), who use time-series methods to estimate the role of anticipated innovations. The increase in optimism following AI news also raises the possibility of bubbles and ultimately unfulfilled expectations.

Table 6: Treatment effects on business expectations in firms not yet considering AI

	Business expectations (z-score, Jul–Dec 2024)			
	Full sample			Not adopted in May 2025
	(1)	(2)	(3)	(4)
<i>Panel A: Full sample</i>				
Both treatments	0.11*** (0.040)		0.088** (0.040)	0.10** (0.050)
AI productivity beliefs (own business)		0.042** (0.018)	0.0033 (0.0032)	
AI productivity beliefs (macro)		-0.012 (0.0079)	0.0038 (0.0026)	
Observations	8929	8929	8929	6194
F-Stat		66		
<i>Panel B: High decision authority sample</i>				
Both treatments	0.14** (0.054)		0.11** (0.056)	0.17** (0.066)
AI productivity beliefs (own business)		0.053** (0.023)	0.0072 (0.0045)	
AI productivity beliefs (macro)		-0.018* (0.010)	0.000071 (0.0036)	
Observations	4914	4914	4914	3458
F-Stat		39		

Notes: The table presents the treatment effects on firms’ business expectations for the next six months in the pooled sample from July to December 2024 controlling for 4-digit NACE industry fixed effects and month fixed effects. The sample focuses on firms, where AI is not yet a topic and is restricted to observations with non-missing data for variables used in all columns. Business expectations are measured monthly using a qualitative question with three possible responses: rather better, no change, rather worse. The variable is standardized as a z-score. “Both treatments” indicate being treated with the “Productivity treatment” or the “Competition treatment” as introduced in the text. In Column 2, AI productivity beliefs about firms’ own business are instrumented with the two treatments. Clustered standard errors at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6.2 Price-setting decisions

Table 7: Treatment effects on price setting

	Price changes (z-score, Jul–Dec 2024)				
	(1)	(2)	(3)	(4)	(5)
	All firms	In use	By AI adoption status:		
			Planned	Discussed	No topic
Productivity treatment	-0.021 (0.024)	-0.12 (0.079)	-0.16* (0.092)	0.069 (0.042)	-0.0065 (0.037)
Competition treatment	-0.017 (0.024)	-0.19*** (0.068)	-0.026 (0.090)	-0.010 (0.041)	0.045 (0.038)
Constant	-0.052*** (0.017)	0.077 (0.049)	-0.016 (0.053)	-0.12*** (0.029)	-0.059** (0.025)
Observations	22025	2820	2367	7462	9376

Notes: The table presents the treatment effects on firms’ realized price-setting in the pooled sample from July to December 2024 controlling for 4-digit NACE industry fixed effects and month fixed effects. Price decisions are measured monthly using a qualitative question with three possible responses: increase, no change, decrease. The variable is standardized as a z-score. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. The treatment effect is shown separately by AI adoption status in Columns 2 to 5. Clustered standard errors at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Next, I turn to the effects of the treatments on firms’ price-setting. In the ifo Business Survey, each month firms are asked whether they increased, decreased, or left their prices unchanged in the previous month. Again, I use the z-score as the outcome variable. Table 7 presents the regression results.

On average, there is no significant effect on firms’ price-setting behavior in the months following the treatments (see Column 1). However, prices decrease significantly by 0.19 standard deviations for firms that already use AI and receive the competition treatment (Column 2).³ The effect size increases to 0.43 standard deviations when focusing on respondents with high decision authority (see Column 2 of Appendix Table A8). A linear probability model using an indicator for a monthly price increase yields consistent results: for AI users, the probability of a price increase falls by 5.4 pp relative to a 11% probability in the control group (see Column 3 of Appendix Table A8).

³The weakly significant negative effect of the productivity treatment among firms that already plan to adopt AI (Column 3) might be driven by lower expected marginal costs.

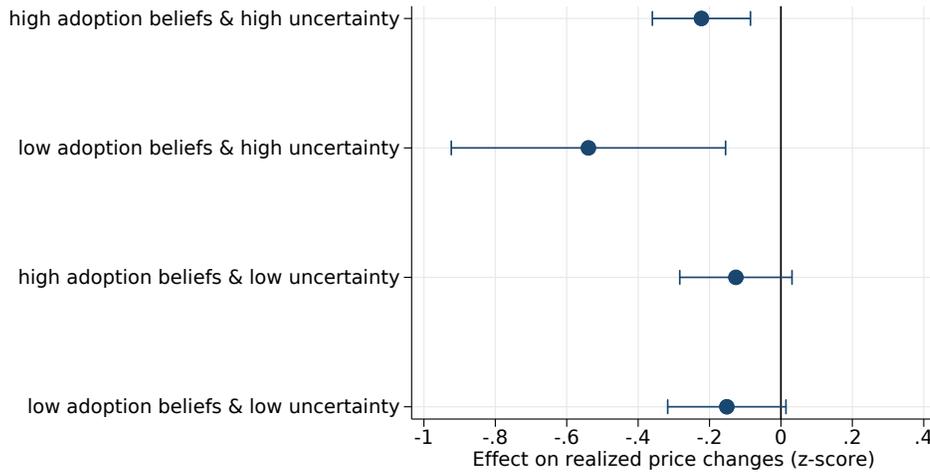
Interpretation through a stylized Bertrand framework with beliefs about AI adoption The negative price effects for firms that already use AI are consistent with the idea that these firms anticipate stronger competitive pressure when informed about competitors' AI adoption. To formalize this intuition, I refer to a stylized Bertrand framework described in Appendix C. Intuitively, in this setting, firms that already use AI must form beliefs about competitors' future AI adoption when setting prices. Because the use of AI raises productivity and thereby lowers marginal costs, a firm's optimal price p_1 decreases with the expected productivity of its competitors $\mathbb{E}[A_2]$. Put differently, higher expected AI adoption among competitors leads firms to expect lower rival costs and, consequently, to set lower prices themselves: $\frac{\partial p_1}{\partial \mathbb{E}[A_2]} < 0$.

Bayesian updating predicts that prior beliefs about competitors' adoption shape firms' price responses. In particular, firms with a flat prior and low average adoption beliefs are expected to react strongly. To investigate this hypothesis, I examine the price effects based on firms' prior beliefs about the share of firms already using AI or planning to adopt it. Specifically, I perform median splits on firms' average prior beliefs about industry adoption rates (median = 30%) and the associated uncertainty (median span = 10 pp), creating four groups.

In line with Bayesian updating of competitive pressure, Figure 5 shows that the downward pressure on price-setting is particularly pronounced for firms that initially expected a low adoption rate and exhibited high uncertainty about it. For these firms, the share of competitors adopting or planning to adopt AI in the competition treatment comes as a negative surprise: more competitors than expected will use the new technology in the near future. The high initial uncertainty indicates that these firms lack superior information that could otherwise diminish the treatment's impact on shifting prior beliefs.

Interestingly, point estimates are negative for all subgroups. A natural interpretation is that the signal sharpens beliefs about a state with relatively strong future competitive pressure. In Bayesian terms, firms move to a more precise belief that competition will be intense, which leads them to put more weight on this risk and cut prices even when their prior already implied substantial adoption. In this view, the increase in precision pushes price-setting down for all firms, while differences in mean updating generate the stronger response for firms with low and uncertain prior beliefs.

Figure 5: Heterogeneity of competition treatments on price-setting by prior AI adoption beliefs



Notes: The figure shows the heterogeneity of the competition treatment effects on price-setting by prior beliefs on AI adoption rates. Firms are split by above/below median average prior beliefs about industry adoption rates (median = 30) and the associated uncertainty (median span = 10). The sample focuses on firms that already use AI. Clustered standard errors at the firm level. Bars represent 90% confidence intervals.

7 Conclusion

This paper establishes three results. First, providing firms with information on realized AI productivity gains and on industry adoption rates causally shifts managers' beliefs about own-firm AI productivity upward without reducing uncertainty. Second, these belief shifts translate into action: treated firms with high decision-making authority are substantially more likely to adopt AI within a year. Third, information about competitor adoption shapes behavior: incumbent adopters lower prices, while non-adopters revise business expectations upward, consistent with increased competitive pressure for adopters and technological-news-driven optimism among not-yet adopters.

The findings highlight the important role of informational frictions and strategic factors in the diffusion of new technologies and suggest that targeted information provision can influence firms' beliefs and accelerate adoption. The centrality of firm-level adoption is consistent with the AI adoption literature, which highlights the role of employers as gatekeepers for widespread worker-level AI use (e.g. [Humlum and Vestergaard, 2025b](#)).

References

- Acemoglu, Daron.** 2025. “The Simple Macroeconomics of AI.” *Economic Policy*, 40(121): 13–58.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson.** 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics*, 151: 103745.
- Bachmann, Rüdiger, Kai Carstensen, Stefan Lautenbacher, Manuel Menkhoff, and Martin Schneider.** 2025. “Uncertainty and Change: Survey Evidence of Firms’ Subjective Beliefs.” NBER Working Paper 29430.
- Banerjee, Abhijit V.** 1992. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, 107(3): 797–817.
- Barsky, Robert B, Susanto Basu, and Keyoung Lee.** 2015. “Whither News Shocks?” *NBER Macroeconomics Annual*, 29(1): 225–264.
- Beauchêne, D.** 2019. “Is Ambiguity Aversion Bad for Innovation?” *Journal of Economic Theory*, 183: 1154–1176.
- Beaudry, Paul, and Franck Portier.** 2006. “Stock Prices, News, and Economic Fluctuations.” *American Economic Review*, 96(4).
- Bertrand, Marianne, and Antoinette Schoar.** 2006. “The Role of Family in Family Firms.” *Journal of Economic Perspectives*, 20(2): 73–96.
- Bick, Alexander, Adam Blandin, and David J. Deming.** 2025. “The Rapid Adoption of Generative AI.” National Bureau of Economic Research NBER Working Paper 32966.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch.** 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy*, 100(5): 992–1026.
- Born, Benjamin, Zeno Enders, Gernot J. Müller, and Knut Niemann.** 2023. “Firm expectations about production and prices: Facts, determinants, and effects.” In *Handbook of Economic Expectations*, ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 355–383. Elsevier. Section: 12.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen.** 2025. “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.” Working Paper.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond.** 2025. “Generative AI at Work.” *Quarterly Journal of Economics*, 140(2): 889–942.
- Brüll, Eduard, Samuel Mäurer, and Davud Rostam-Afschar.** 2025. “When the Tax Machine Cometh: Who Benefits from Automation?” Mimeo, Mannheim.
- Burszty, Leonardo, Alex Imas, Rafael Jiménez-Durán, Aaron Leonard, and Christopher Roth.** 2025. “Social Dynamics of AI Adoption.” NBER Working Paper 34488.

- Cascaldi-Garcia, Danilo, and Marija Vukotic.** 2022. “Patent-Based News Shocks.” *Review of Economics and Statistics*, 104(1): 51–66.
- Chatterji, Aaron, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman.** 2025. “How People Use ChatGPT.” NBER Working Paper 34255.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar.** 2018. “How Do Firms Form Their Expectations? New Survey Evidence.” *American Economic Review*, 108(9): 2671–2713.
- Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele.** 2020. “Inflation Expectations and Firm Decisions: New Causal Evidence.” *Quarterly Journal of Economics*, 135(1): 165–219.
- Conley, Timothy G, and Christopher R Udry.** 2010. “Learning about a New Technology: Pineapple in Ghana.” *American Economic Review*, 100(1): 35–69.
- Enders, Zeno, Franziska Hünnekes, and Gernot J. Müller.** 2022. “Firm Expectations and Economic Activity.” *Journal of the European Economic Association*, 20(6): 2396–2439.
- Enders, Zeno, Michael Kleemann, and Gernot J. Müller.** 2021. “Growth Expectations, Undue Optimism, and Short-Run Fluctuations.” *Review of Economics and Statistics*, 103(5): 905–921.
- Felten, Edward, Manav Raj, and Robert Seamans.** 2021. “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses.” *Strategic Management Journal*, 42(12): 2195–2217.
- Freund, Lukas B., and Lukas F. Mann.** 2025. “Job Transformation, Specialization, and the Labor Market Effects of AI.” CESifo Working Paper 12072.
- Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein.** 1993. “Risk Management: Coordinating Corporate Investment and Financing Policies.” *Journal of Finance*, 48(5): 1629–1658.
- Griliches, Zvi.** 1957. “Hybrid Corn: An Exploration in the Economics of Technological Change.” *Econometrica*, 25(4).
- Hall, Bronwyn H.** 2006. “Innovation and Diffusion.” In *The Oxford Handbook of Innovation.*, ed. Jan Fagerberg and David C. Mowery.
- Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller.** 2025. “Artificial Intelligence and the Labor Market.” National Bureau of Economic Research NBER Working Paper 33509.
- Hennrich, Jonas, Stefan Sauer, and Klaus Wohlrabe.** 2023. “Who Reports the Mood in German Boardrooms? Evidence from the Ifo Business Survey.” CESifo Working Paper 10571.
- Hiersemenzel, Magdolna, Stefan Sauer, and Klaus Wohlrabe.** 2022. “On the Representativeness of the ifo Business Survey.” CESifo Working Paper 9863.
- Humlum, Anders, and Emilie Vestergaard.** 2025a. “Large Language Models, Small Labor Market Effects.” NBER Working Paper 33777.

- Humlum, Anders, and Emilie Vestergaard.** 2025b. “The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities Among Workers.” *Proceedings of the National Academy of Sciences*, 122(1): e2414972121.
- Hvidberg, Kristoffer B, Claus T Kreiner, and Stefanie Stantcheva.** 2023. “Social Positions and Fairness Views on Inequality.” *Review of Economic Studies*, 90(6): 3083–3118.
- Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun.** 2025. “The Diffusion of New Technologies.” *The Quarterly Journal of Economics*, 140(2): 1299–1365.
- Kerssenfischer, Mark, and Maik Schmeling.** 2024. “What Moves Markets?” *Journal of Monetary Economics*, 145: 103560.
- Licht, Thomas, and Klaus Wohlrabe.** 2024. “AI Adoption Among German Firms.” CESifo Working Paper 11459.
- Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart.** 2023. “Information Frictions Among Firms and Households.” *Journal of Monetary Economics*, 135: 99–115.
- Lorenzoni, Guido.** 2009. “A Theory of Demand Shocks.” *American Economic Review*, 99(5): 2050–2084.
- Lucas, Robert E.** 1972. “Expectations and the Neutrality of Money.” *Journal of Economic Theory*, 4(2): 103–124.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2025. “Rational Inattention during an RCT.” *American Economic Review: Insights*, forthcoming.
- McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas.** 2024. “AI Adoption in America: Who, What, and Where.” *Journal of Economics & Management Strategy*, 33(2): 375–415.
- Noy, Shakked, and Whitney Zhang.** 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science*, 381(6654): 187–192.
- Opler, Tim, Lee Pinkowitz, and Rohan Williamson.** 1999. “The Determinants and Implications of Corporate Cash Holdings.” *Journal of Financial Economics*, 52(1): 3–46.
- Otis, Nicholas G, Berkeley Haas, Rowan Clarke, and Rembrand Koning.** 2023. “The Uneven Impact of Generative AI on Entrepreneurial Performance.” Working Paper.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer.** 2023. “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.” arXiv Working Paper.
- Rogers, Everett.** 2003. *Diffusion of Innovations*. . 5 ed., Free Press.
- Sauer, Stefan, Moritz Schasching, and Klaus Wohlrabe.** 2023. *Handbook of ifo Surveys. ifo Beiträge zur Wirtschaftsforschung*.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2012. “What’s News in Business Cycles.” *Econometrica*, 80(6): 2733–2764.

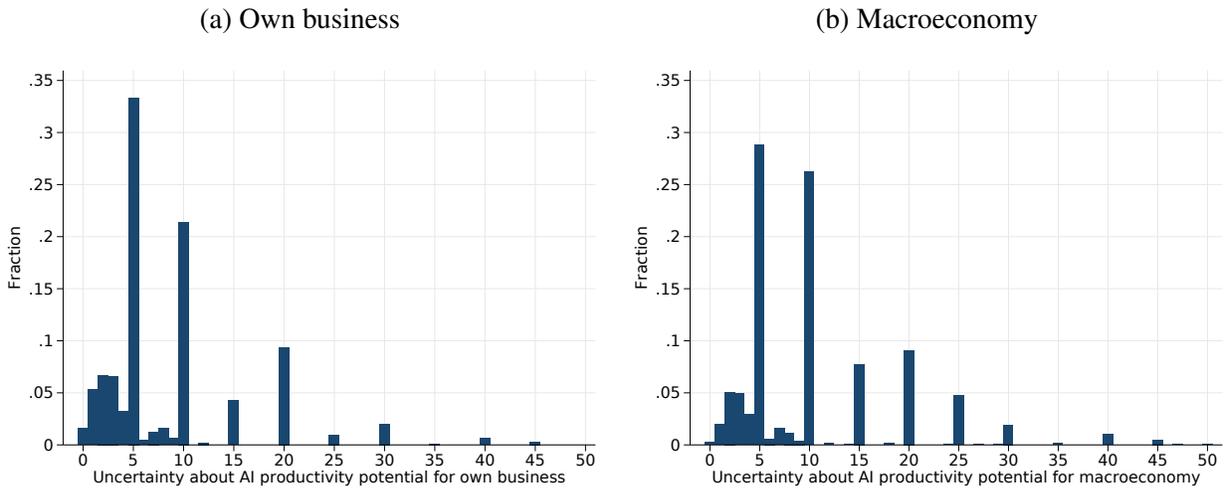
van Rooij, Maarten, Olivier Coibion, Dimitris Georgarakos, Bernardo Candia, and Yuriy Gorodnichenko. 2024. “Keeping up with the Jansens: Causal Peer Effects on Household Spending, Beliefs and Happiness.” NBER Working Paper 32107.

Young, H. Peyton. 2009. “Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning.” *American Economic Review*, 99(5): 1899–1924.

Appendix

A Additional figures

Figure A1: Firms' subjective uncertainty about AI productivity gains



Notes: The figure shows histograms of firms' subjective uncertainty about AI productivity gains for their own business (left panel) and for the aggregate economy (right panel). These beliefs are elicited in terms of ranges in June 2024. Subjective uncertainty is measured by the span between firms' lower and upper bound estimates of productivity gains. The sample is restricted to the control group. The survey questions are in the text.

Figure A2: Relationship between firms' AI productivity beliefs and AI industry exposure (Felten, Raj and Seamans (2021) measure)

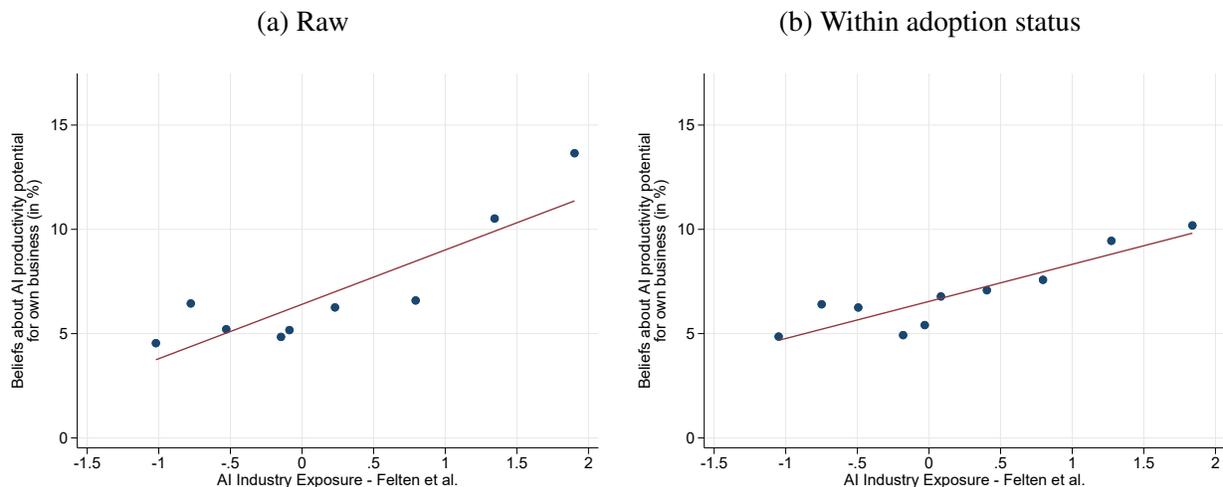


Figure A3: Firms' beliefs on AI adoption of competitors

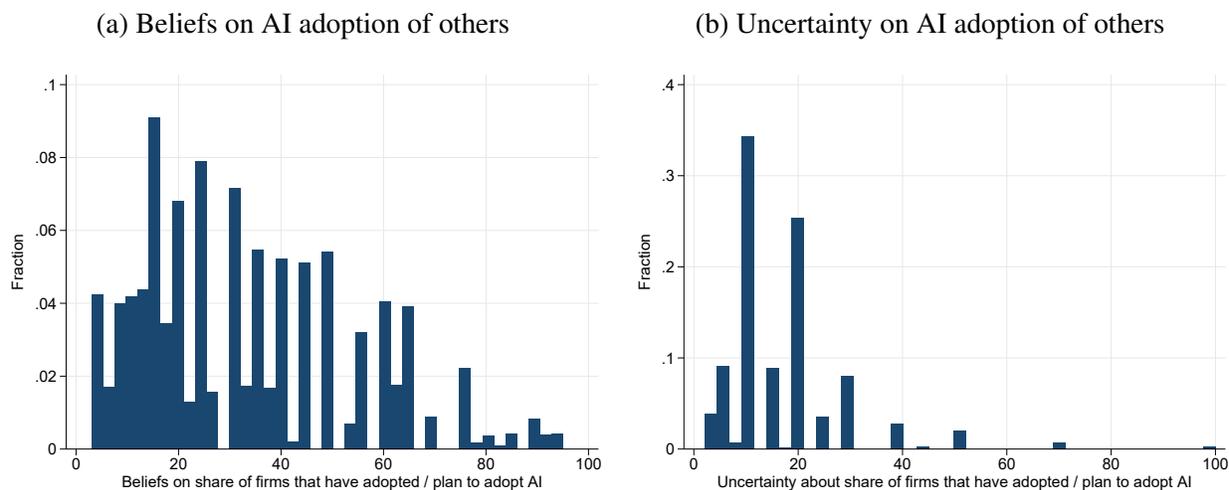
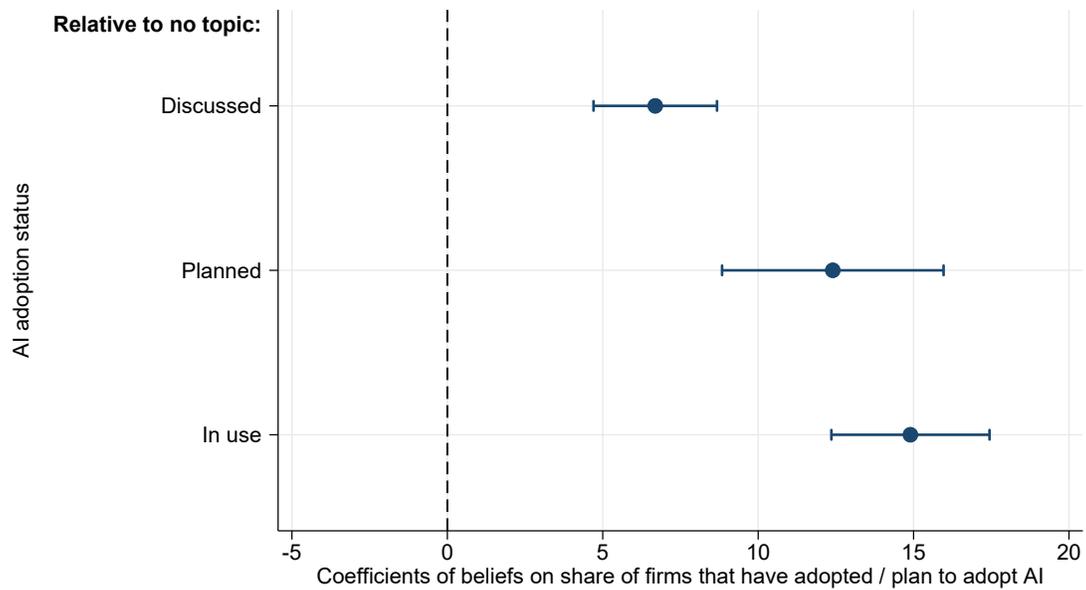


Figure A4: Firms' AI adoption and beliefs on AI adoption of competitors



Notes: The figure shows the coefficients from a regression of firms' beliefs on competitors' AI adoption and own AI adoption status, relative to firms where AI adoption is no topic. All specifications include 4-digit NACE industry fixed effects. Clustered standard errors at the 4-digit NACE industry level. Bars represent 90% confidence intervals.

B Additional tables

Table A1: Predictors of treatment status

	Full sample			High decision authority sample		
	(1)	(2)	(3)	(4)	(5)	(6)
	Control	Productivity treatment	Competition treatment	Control	Productivity treatment	Competition treatment
Firm size (omitted category: large)						
medium	0.012 (0.030)	0.014 (0.028)	-0.026 (0.029)	-0.024 (0.048)	0.046 (0.046)	-0.022 (0.046)
small	-0.0081 (0.037)	0.0013 (0.033)	0.0068 (0.034)	-0.034 (0.053)	0.021 (0.053)	0.013 (0.051)
Industry (omitted category: manufacturing)						
Retail/wholesale	-0.014 (0.023)	-0.018 (0.024)	0.032 (0.023)	-0.022 (0.030)	-0.021 (0.029)	0.043 (0.027)
Services	0.020 (0.022)	-0.016 (0.023)	-0.0046 (0.020)	-0.024 (0.028)	-0.00083 (0.034)	0.024 (0.030)
Region (omitted category: west)						
East	-0.018 (0.025)	0.029 (0.026)	-0.010 (0.023)	-0.034 (0.038)	0.028 (0.035)	0.0056 (0.036)
Prior AI Adoption status (omitted category: no topic)						
In use	0.029 (0.031)	-0.032 (0.028)	0.0039 (0.027)	-0.023 (0.037)	0.023 (0.042)	0.00064 (0.035)
Planned	0.031 (0.032)	-0.022 (0.030)	-0.0093 (0.030)	0.037 (0.044)	-0.026 (0.041)	-0.011 (0.042)
Discussed	-0.0065 (0.020)	-0.0063 (0.022)	0.013 (0.020)	-0.015 (0.028)	0.032 (0.030)	-0.017 (0.028)
Log Revenues 2023	-0.0044 (0.0061)	-0.0015 (0.0056)	0.0059 (0.0052)	-0.0071 (0.0086)	-0.0048 (0.0081)	0.012 (0.0075)
Observations	3251	3251	3251	1727	1727	1727
R^2	0.003	0.002	0.003	0.002	0.004	0.004
F-Stat	.89	.67	.87	.32	.65	.92
P-value	.53	.74	.55	.97	.75	.51

Notes: The table displays results from linear probability models, where the dependent variable is an indicator for assignment to a given treatment or control group. Clustered standard errors at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2: Testing for heterogeneity in treatment effects on AI productivity beliefs

	(1)	(2)	(3)	(4)	(5)	(6)
Productivity treatment	1.93*** (0.34)	2.14*** (0.28)	1.88*** (0.28)	2.11*** (0.48)	1.98*** (0.38)	1.82*** (0.49)
Competition treatment	1.68*** (0.38)	2.24*** (0.31)	2.08*** (0.42)	2.87*** (0.63)	2.09*** (0.46)	2.14*** (0.60)
Large firms	0.11 (0.32)					
× Productivity treatment	0.061 (0.44)					
× Competition treatment	1.01* (0.53)					
East		0.48 (0.42)				
× Productivity treatment		-1.40** (0.60)				
× Competition treatment		-0.99 (0.70)				
High exposure			0.86** (0.40)			
× Productivity treatment			0.093 (0.58)			
× Competition treatment			-0.033 (0.64)			
High education				0.63 (0.44)		
× Productivity treatment				-0.22 (0.61)		
× Competition treatment				-1.11 (0.77)		
External funds					0.024 (0.36)	
× Productivity treatment					-0.17 (0.51)	
× Competition treatment					0.051 (0.59)	
High decision authority						-0.18 (0.40)
× Productivity treatment						0.088 (0.62)
× Competition treatment						-0.064 (0.70)
Observations	3605	3605	3261	1922	3107	2716

Notes: The table displays the treatment effects on firms’ beliefs about productivity gains of AI for their own business interacted with various firm characteristics. Only firms that do not yet use AI are considered. I use the following indicator definitions. “Large firm”: firms with more than 50 employees. “East”: firms located in former East Germany. “High exposure”: at least one of the three occupations—programming, writing, or customer service—for which the productivity treatment provided information on AI productivity gains is overrepresented in the firm’s industry. “High education”: higher education than a bachelor degree. “External funds”: using external funds for investment. “High decision authority”: respondent has high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. The treatment effect is shown separately by AI adoption status in Columns 2 to 5. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3: Treatment effects on AI adoption by financing source of investment

	(1)	(2)	(3)	(4)
	Invest with		Invest with	
	Internal funds	External funds	Internal funds	External funds
Productivity treatment	0.15*** (0.055)	0.029 (0.056)	0.063* (0.036)	-0.041 (0.041)
Competition treatment	0.11** (0.050)	-0.015 (0.063)	0.069** (0.031)	-0.045 (0.045)
Constant	0.19*** (0.030)	0.26*** (0.034)	0.24*** (0.019)	0.30*** (0.026)
Observations	567	508	1117	921
High decision authority sample	yes	yes	no	no
Industry FE	yes	yes	yes	yes

Notes: The table displays the treatment effects on firms' probability to adopt AI in May 2025. "Productivity treatment" and "Competition treatment" are dummies indicating the treatment groups as introduced in the text. In Columns 1 and 3, firms that financed last year's investment solely with internal funds are considered; in Columns 2 and 4 firms that also used external funds. Columns 1-2 focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A4: Treatment effects on AI adoption by size and investment volume

	(1)	(2)	(3)	(4)
	Firm size		Investment/revenue ratio	
	Large	Small	High	Low
Productivity treatment	0.18** (0.071)	0.044 (0.047)	0.15*** (0.053)	0.050 (0.055)
Competition treatment	0.11* (0.064)	0.024 (0.040)	0.092* (0.049)	0.0030 (0.052)
Constant	0.20*** (0.037)	0.23*** (0.024)	0.19*** (0.028)	0.25*** (0.029)
Observations	409	747	523	546
High decision authority sample	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes

Notes: The table displays the treatment effects on firms' probability to adopt AI in May 2025. "Productivity treatment" and "Competition treatment" are dummies indicating the treatment groups as introduced in the text. In Columns 1, large firms (at least 50 employees) are considered; in Columns 2, small firms (less than 50 employees); in Columns 3 firms with a high investment/revenue ratio in the last two years (higher than 3%); in Columns 4 firms with a low investment/revenue ratio in the last two years (max. 3%). All columns focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A5: Treatment effects on investment plans in software and databases

	(1)	(2)	(3)	(4)	(5)
Both treatments	0.034 (0.040)	0.12** (0.056)	0.11** (0.057)		
Productivity treatment				0.15** (0.063)	
Competition treatment				0.079 (0.066)	
AI productivity beliefs (own business)					0.045* (0.025)
Observations	1770	895	895	895	895
F-Stat					12
High decision authority sample	no	yes	yes	yes	yes
AI prod. beliefs ctrl (macro)	no	no	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes

Notes: The table displays the treatment effects on firms' investment plans in software and databases. Investment plans are elicited in the November wave. The question is qualitative and framed in terms of planned changes relative to the current year (increase, no change, decrease). The variable is standardized (z-score). "Productivity treatment" and "Competition treatment" are dummies indicating the treatment groups as introduced in the text. "Both treatments" is a dummy indicating both treatment groups. In all regressions, it is controlled for firms' business situation and total investment plans. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6: Treatment effects on business expectations: robustness

	Business expectations (Jul–Dec 2024)					
	z-score		Increase indicator		Decrease indicator	
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity treatment	0.078* (0.045)	0.11* (0.064)	0.014 (0.013)	-0.0026 (0.021)	-0.036 (0.023)	-0.068** (0.033)
Competition treatment	0.13*** (0.044)	0.17*** (0.059)	0.017 (0.012)	0.0025 (0.020)	-0.070*** (0.022)	-0.10*** (0.031)
Constant	-0.21*** (0.030)	-0.30*** (0.040)	0.079*** (0.0083)	0.094*** (0.014)	0.43*** (0.015)	0.49*** (0.021)
Observations	9390	5156	9390	2599	9390	5156
High decision authority sample	no	yes	no	yes	no	yes

Notes: The table presents the treatment effects on firms’ business expectations in the pooled sample from July to December 2024 controlling for 4-digit NACE industry fixed effects and month fixed effects. The sample focuses on firms, where AI is not yet a topic. Business expectations are measured monthly using a qualitative question with three possible responses: rather better, no change, rather worse. In Columns 1-2, the variable is standardized as a z-score. In Columns 3-4 (5-6), the outcome variable is an indicator for higher (lower) business expectations. Columns 2, 4, and 6 focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. Clustered standard errors at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A7: Treatment effects on GDP growth expectations in firms not yet considering AI

	GDP Growth 2025		Inflation 2025	
	(1)	(2)	(3)	(4)
<i>Panel A: Full sample</i>				
Both treatments	0.13** (0.056)	0.12** (0.057)	0.024 (0.045)	0.038 (0.047)
AI productivity beliefs (own business)		-0.00067 (0.0048)		-0.0034 (0.0036)
Observations	1257	1257	1257	1257
<i>Panel B: High decision authority sample</i>				
Both treatments	0.29*** (0.091)	0.29*** (0.094)	-0.034 (0.073)	-0.024 (0.076)
AI productivity beliefs (own business)		-0.00029 (0.0066)		0.00056 (0.0049)
Observations	658	658	658	658
Industry FE	yes	yes	yes	yes
AI productivity beliefs ctrl (macro)	no	yes	no	yes

Notes: The table presents the treatment effects on firms' GDP growth and Inflation expectations for 2025, elicited in November 2024, controlling for 4-digit NACE industry fixed effects. The sample focuses on firms, where AI is not yet a topic. P1/P99 outliers of firms' GDP growth and Inflation expectations are dropped. "Both treatments" indicate being treated with the "Productivity treatment" or the "Competition treatment" as introduced in the text. The sample focuses on firms, where AI is not yet a topic. Panel B restricts the sample to respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. Standard errors clustered at the 4-digit NACE industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A8: Treatment effects on price-setting: robustness

	Price changes (Jul–Dec 2024)					
	z-score		Increase indicator		Decrease indicator	
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity treatment	-0.12 (0.079)	-0.069 (0.17)	-0.032 (0.028)	-0.0011 (0.069)	0.020 (0.018)	0.029 (0.025)
Competition treatment	-0.19*** (0.068)	-0.43*** (0.14)	-0.054** (0.022)	-0.095** (0.047)	0.026 (0.018)	0.091** (0.037)
Constant	0.077 (0.049)	0.22** (0.092)	0.11*** (0.017)	0.15*** (0.038)	0.043*** (0.011)	0.023* (0.012)
Observations	2820	1166	2820	1166	2820	1166
High decision authority sample	no	yes	no	yes	no	yes

Notes: The table presents the treatment effects on firms’ realized price-setting in the pooled sample from July to December 2024 controlling for 4-digit NACE industry fixed effects and month fixed effects. The sample focuses on firms, where AI is used. Price decisions are measured monthly using a qualitative question with three possible responses: increase, no change, decrease. In Columns 1-2, the variable is standardized as a z-score. In Columns 3-4 (5-6), the outcome variable is an indicator for price increases (decreases). Columns 2, 4, and 6 focus on respondents with high decision-making authority, defined as (i) the respondent is a top manager (C-suite or head of department) or the owner; and (ii) the firm is family-owned. “Productivity treatment” and “Competition treatment” are dummies indicating the treatment groups as introduced in the text. Clustered standard errors at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Stylized Bertrand framework

C.1 Environment

Two firms $j \in \{1, 2\}$ compete in prices for differentiated products. Market demand for firm j is

$$q_1 = a - bp_1 + dp_2, \quad q_2 = a - bp_2 + dp_1,$$

with $a > 0$ and $b > d > 0$. Parameters b and d capture own- and cross-price sensitivities, respectively. The condition $b > d$ ensures that products are imperfect substitutes and demand is well behaved.

Each firm has constant marginal cost c_j , $j \in \{1, 2\}$.

C.2 Bertrand pricing and equilibrium

Firms simultaneously choose prices p_1, p_2 to maximize static profits

$$\pi_1 = (p_1 - c_1)q_1, \quad \pi_2 = (p_2 - c_2)q_2.$$

The first-order conditions are

$$\begin{aligned} \frac{\partial \pi_1}{\partial p_1} &= a + bc_1 - 2bp_1 + dp_2 = 0, \\ \frac{\partial \pi_2}{\partial p_2} &= a + bc_2 - 2bp_2 + dp_1 = 0. \end{aligned}$$

Solving this linear system yields the unique Nash equilibrium prices

$$\begin{aligned} p_1^*(c_1, c_2) &= \frac{2ab + ad + 2b^2c_1 + bdc_2}{4b^2 - d^2}, \\ p_2^*(c_1, c_2) &= \frac{2ab + ad + 2b^2c_2 + bdc_1}{4b^2 - d^2}, \end{aligned}$$

where $4b^2 - d^2 > 0$ under $b > d$.

Differentiating Firm 1's equilibrium price with respect to Firm 2's marginal cost gives

$$\frac{\partial p_1^*}{\partial c_2} = \frac{bd}{4b^2 - d^2} > 0.$$

C.3 Price response to beliefs

Let

$$c_j = \frac{w}{A_j},$$

where $w > 0$ is the wage and A_j is firm j 's productivity. Firm 1 is an adopter with high productivity $A_1 = A^{AI}$ and marginal cost $c_1 = w/A^{AI}$. Firm 2's future productivity A_2 is uncertain from Firm 1's perspective.

Let F denote Firm 1's belief distribution over A_2 . Given these beliefs, Firm 1's expected marginal cost of Firm 2 is

$$\bar{c}_2(F) \equiv \mathbb{E}_F \left[\frac{w}{A_2} \right].$$

An information shock (competition treatment) that makes Firm 1 believe competitors are more likely to adopt AI corresponds to a shift in beliefs from F to F' such that probability mass moves from low to high values of A_2 . In other words, under F' high productivity states of A_2 become more likely and low productivity states become less likely. Since the function $a \mapsto w/a$ is strictly decreasing, any such shift implies

$$\mathbb{E}_{F'}[A_2] \geq \mathbb{E}_F[A_2] \quad \text{and} \quad \mathbb{E}_{F'} \left[\frac{w}{A_2} \right] < \mathbb{E}_F \left[\frac{w}{A_2} \right],$$

so that

$$\bar{c}_2(F') < \bar{c}_2(F).$$

From the Bertrand solution above we know that Firm 1's equilibrium price is strictly increasing in Firm 2's marginal cost,

$$\frac{\partial p_1^*(c_1, c_2)}{\partial c_2} > 0.$$

Evaluating at $c_2 = \bar{c}_2(F)$ and $c_2 = \bar{c}_2(F')$ then yields

$$p_1^*(c_1, \bar{c}_2(F')) < p_1^*(c_1, \bar{c}_2(F)).$$

In words, when the information shock makes Firm 1 expect higher competitor productivity—or, more generally, increases its confidence that competitors are in a high productivity (AI-adopting) state—this lowers the expected marginal cost of competitors and induces Firm 1 to set a lower optimal price.