

Automation and Demand for Labor

Experimental Evidence from White Collar Jobs *

Eduard Brüll[†], Samuel Mäurer[‡], Davud Rostam-Afschar[§]

February 26, 2025

Abstract

How do employers respond to automation shocks? We investigate this question using a randomized information intervention that exogenously shifts employers' beliefs about automation rates of their workforce. Focusing on the tax consulting and auditing sectors where well-defined job titles and high exposure to generative language models create credible automation potential we assign firms to one of three treatment groups or a control group. Treated firms revise revenue and profit expectations upward but do not alter hiring or firing decisions, suggesting automation enhances efficiency without immediate labor displacement. Notably, wages remain unchanged, indicating firms intend to retain productivity gains rather than share them with employees.

Keywords: Automation, Labor Demand, White Collar Jobs, Human Capital Investment, Technological Displacement

JEL classification: J23, J24, O33, C93

*We thank xxx. The study was pre-registered in the AEA RCT registry (RCT ID: AEARCTR-0014788). This experiment was reviewed by the Ethics Committee of the University of Mannheim (EC-8/2024). We declare that we have no interests, financial or otherwise, that relate to the research described in this paper. We are grateful to the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for financial support through CRC *TRR 266 Accounting for Transparency* (Project-ID 403041268).

[†]ZEW Leibniz Centre for European Economic Research, 68161 Mannheim, Germany (e-mail: eduard.bruell@zew.de),

[‡]University of Mannheim, 68131 Mannheim, Germany, Germany (e-mail: samuel.maeurer@uni-mannheim.de),

[§]University of Mannheim, 68131 Mannheim, Germany; GLO; IZA; NeST (e-mail: rostam-afschar@uni-mannheim.de),

1 Introduction

Digital technologies are a key driver of economic growth, yet their diffusion fundamentally reshapes labor markets. Emerging technologies such as generative artificial intelligence (GenAI) are increasingly capable of automating complex cognitive tasks, raising critical questions about their impact on employment strategies, wage structures, and skill requirements.

Historically, automation has primarily affected blue-collar jobs in manufacturing and manual labor-intensive sectors, where robotics and mechanization replaced routine tasks (e.g., Acemoglu and Restrepo, 2020; Dauth et al., 2021). However, the latest wave of AI-driven automation differs markedly: highly educated, high-income white-collar occupations are now among the most exposed (e.g. Eloundou et al., 2023). Recent studies suggest that AI assistance can significantly enhance worker productivity in professional settings. For example, Noy and Zhang (2023) show that AI-assisted writers complete tasks faster and produce higher-quality output, while Brynjolfsson et al. (2023) document a 14% productivity increase among customer service agents using generative AI, with the largest gains among less-experienced workers. At the same time, Felten et al. (2023) highlight that occupations relying on communication, analysis, and creative abilities—once considered resistant to automation—are now highly susceptible to AI-driven disruption.

Despite the growing literature on automation’s effects, most existing studies focus on workers rather than employers. How firms adapt to automation risks remains less understood. This paper addresses this gap by investigating how employers in white-collar industries respond when their beliefs about automation rates are exogenously shifted. Using a randomized information intervention, we study firms in the tax consulting and auditing sectors—industries characterized by well-defined occupational roles and high exposure to generative AI. Our survey, which targeted the entire population of tax advisors and auditors listed in Germany’s official register, allows us to analyze how belief updates about automation potential influence employment strategies, wage expectations, and business outlooks. Additionally, we examine firm-level outcomes such as revenue, profits, and costs to assess whether employers view automation as an opportunity for efficiency gains or a disruptive force necessitating workforce reductions.

Our results reveal several striking patterns. First, employers systematically underestimate automation risks for their profession. Initial beliefs about the share of tasks that could be automated within the next decade are significantly lower than expert assessments. After receiving objective information on automation potential, respondents revise their beliefs upward, particularly for lower-skilled roles such as tax clerks and certified tax assistants. Belief updating is weaker for higher-skilled occupations like auditors and tax

advisors, suggesting that firms perceive greater barriers to automation at the top of the professional hierarchy.

Second, despite updating their beliefs about automation rates, firms do not adjust hiring or firing plans for the next three years. However, treated firms revise their revenue and profit expectations upward, consistent with the notion that automation enhances efficiency without triggering immediate labor displacement. Interestingly, cost expectations also rise, indicating anticipated investments in new technologies or upskilling initiatives. However, wage expectations remain unchanged, suggesting that firms intend to retain productivity gains rather than pass them on to employees.

Third, firms exposed to new information on automation not only reassess existing job roles but also anticipate new tasks emerging as a consequence of AI adoption. In particular, employers expect increased demand for legal tech expertise, compliance monitoring, and AI interaction skills such as prompt engineering. This aligns with a growing recognition that generative AI does not merely replace existing jobs but also reshapes job content and skill requirements.

Our study contributes to three strands of economic research. First, we extend the literature on automation and labor markets. Much of this literature focuses on the effects of automation in manufacturing and manual labor-intensive sectors. For example, Acemoglu and Restrepo (2020) find that increased robot adoption in U.S. manufacturing is associated with significant reductions in employment and wages, with localized displacement effects that are not fully offset by gains in other sectors. In contrast, Dauth et al. (2021) examine the impact of robots in Germany, finding job losses in manufacturing but compensating employment gains in service industries. Further studies uncover interesting heterogeneities in automation’s effects (see Aghion et al., 2022, for a survey of the recent literature). For instance, Bessen et al. (2020) demonstrate that firms adopting automation technologies often save labor while maintaining wage growth. Similarly, Koch et al. (2021) find that robot adoption in Spanish manufacturing firms leads to significant output gains and net job creation, suggesting that automation can enhance productivity without necessarily reducing employment. Aghion et al. (2020) further show that automation in French manufacturing increases employment at the firm and industry levels, with productivity effects outweighing displacement effects. At the same time, recent evidence suggests that AI adoption may increase wage inequality, with high-wage workers benefiting disproportionately while low-wage and production workers face negative employment effects (Bonfiglioli et al., 2024). Unlike robotics, which primarily automated routine manual tasks, AI influences both routine and non-routine cognitive tasks. Gathmann et al. (2024) show that AI reduces abstract tasks like information gathering while increasing the need for high-level routine tasks that require monitoring and process oversight. Our study

extends this literature to white-collar industries, suggesting that GenAI-driven automation might enhance efficiency without triggering immediate job losses. Moreover, unlike manufacturing automation, which primarily substitutes for low-skill tasks, generative AI influences a broader spectrum of occupational roles, including highly skilled professions.

Second, our study takes a novel employer-centered perspective, which allows us to detect potential adjustments in hiring plans, wage strategies, and skill investment decisions before they materialize as measurable labor market outcomes. This is particularly valuable because the existing literature, while rich in documenting the impacts of automation at the worker and firm levels, often examines outcomes only after automation technologies have been implemented. For example, Bessen et al. (2025) examine worker-level outcomes following firm-level automation expenditures, finding significant impacts on worker displacement and cumulative wage losses. Similarly, Acemoglu et al. (2022) analyze the adoption of AI using vacancy-level data, demonstrating shifts in hiring patterns and skill requirements at AI-exposed establishments between 2010 and 2018. However, their analysis does not extend to the most recent wave of generative AI adoption, leaving open questions about how firms anticipate and adapt to these transformative technologies. While these approaches are invaluable for understanding post-adoption consequences, it does not shed light on how firms plan for or adapt to automation before investments are made. Our study instead captures firms anticipatory responses, showing that automation beliefs influence business expectations and investment strategies before observable adjustments occur.

Our study methodologically builds upon the literature employing information interventions to examine how accurate data can correct misperceptions and influence economic preferences and behaviors (e.g. Alesina et al., 2022; Kuziemko et al., 2015; Coibion et al., 2018; Wiswall and Zafar, 2015). Within this framework, recent research has focused on how information about automation affects individual expectations and behaviors. For instance, Arntz et al. (2022) find that while many workers fear job losses due to automation, providing information about neutral net employment effects can reduce these concerns. Similarly, Jeffrey (2021) shows that the framing of automation influences policy preferences, with narratives of inevitable displacement increasing support for redistribution. Furthermore, Lergetporer et al. (2023), which is closest to our study, since it relies on the same automatability measure, demonstrate that workers often underestimate the automatability of their occupations. They also show that providing personalized information about automation risks increases their willingness to engage in further training. Our study extends this literature by shifting the focus to employers, examining how updated beliefs about automation potential influence business expectations and labor market strategies in a white-collar industry.

The remainder of this paper is organized as follows. Section 2 describes the survey, experimental setup, and estimation strategy. Section 3 presents the results, while Section 4 concludes.

2 Survey and Experimental Setup

2.1 The GBP Tax Advisor and Auditor Survey

Our analysis draws on a specialized survey of tax advisors and auditors conducted between November 2024 and January 2025. As part of the German Business Panel (GBP), this survey targeted all professionals listed in Germany's official register of licensed tax advisors (*Steuerberater*) and auditors (*Wirtschaftsprüfer*), leveraging over 100,000 email addresses from the register.

Maintained by professional chambers, the official register is both mandatory and exhaustive, covering all individuals and firms authorized to practice as tax advisors and auditors in Germany. Since registration is a legal requirement, this dataset is uniquely complete, unlike alternative sources, which are often fragmented or incomplete.

In Germany, tax advisors and auditors are classified as *Freiberufler* (liberal professionals), a designation that differentiates them from traditional firms. While they operate independently under distinct legal and tax frameworks, they frequently employ significant numbers of workers, playing a vital role in the labor market. However, standard firm-level datasets, such as those used in business or employer-employee panel studies, typically exclude *Freiberufler*, creating a significant data gap. The mandatory register allows us to bridge this gap, offering direct insights into this unique professional group and their responses to automation and other economic trends.

Survey Modality and Data Collection Process The survey was conducted online and designed to be accessible across multiple devices, including desktop computers, tablets, and smartphones. The interface was optimized for both large and small screens to ensure a seamless user experience across different devices.

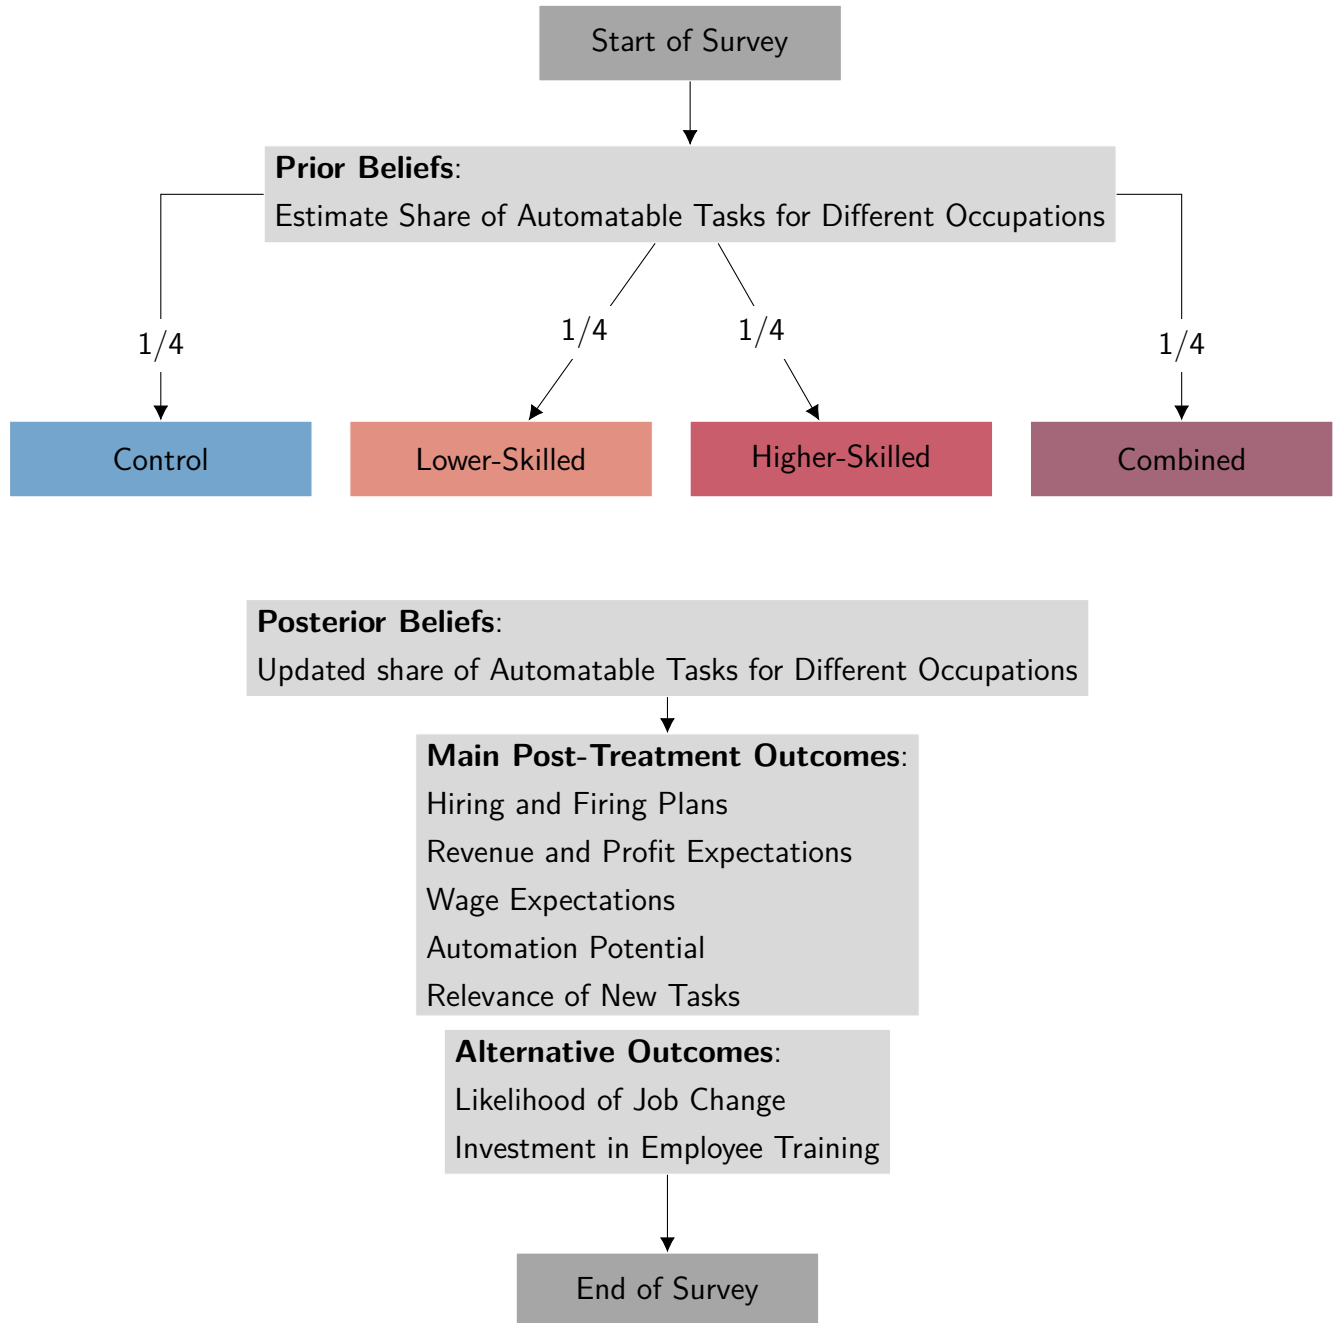
To manage outreach effectively, we distributed the survey in weekly batches starting in November 2024. Each respondent received up to two reminders if they had not completed the survey: the first two weeks after the initial invitation and the second one month later. This structured follow-up approach helped maximize response rates while preventing excessive survey fatigue.

The median survey completion time was 684 seconds (approximately 11.4 minutes), which aligns well with the expected 15-minute duration. While most respondents completed the survey promptly after receiving their invitation, a subset only finalized their responses after receiving a reminder. These delayed completions resulted in some extreme duration values, as certain participants resumed the survey days or even weeks after initially opening it.

Ongoing Data Collection This paper presents a preliminary version of our analysis, with survey fieldwork still ongoing. As new responses are collected, the dataset is dynamically updated to incorporate additional observations. At the current stage, the rough sample includes 1,657 responses before applying the filtering steps described in Section 2.3. After excluding economically inactive respondents and ensuring data consistency, the final sample currently consists of 1,229 observations. Future iterations of this paper will integrate the complete dataset as fieldwork progresses.

2.2 Experimental Setup

Figure 1: Experimental Design



Notes: The figure illustrates the design of our information treatment

The key part of our survey is a randomized information intervention designed to examine how tax advisory and auditing firms respond to updated information on the automation potential of their workforce. The experiment follows a structured sequence, visually represented in Figure 1.

Before the intervention, we collect respondents employment levels and their prior beliefs about the automatability of four tax-related occupations: tax clerks, certified tax assistants, tax advisors, and auditors. Participants estimate the percentage of core activities within each occupation that they believe can be automated within the next ten years.

Following this, respondents are randomly assigned to one of three treatment groups or a control group, each with equal probability. The treatment groups receive personalized information about the automatability of each occupation, based on occupation-level estimates from the IAB Job-Futuromat.

The IAB Job Futuromat The IAB Job-Futuromat is a tool developed by the Institute for Employment Research (IAB), a research division of Germanys Federal Employment Agency. It provides a systematic assessment of how digital technologies impact various occupations by evaluating the degree to which specific tasks within those roles can be automated. It covers approximately 4,000 occupations and is based on expert-driven task analyses, making it one of the most detailed and policy-relevant resources on labor automation.

The automatability scores in the Job-Futuromat are built on BERUFENET data, an expert database maintained by the German Federal Employment Agency, which documents occupational tasks, required skills, and competencies for career guidance and job placement. The methodology behind BERUFENET, as described by Dengler et al. (2014), follows a task-based approach similar to O*NET in the U.S., systematically mapping occupations to their core tasks and assessing their substitutability by automation. This expert-driven approach offers a robust alternative to survey-based task measurements, ensuring that occupational analyses remain consistent and reliable over time.

Research using Job-Futuromat data has revealed that occupations with higher substitutability potential tend to experience lower employment growth on average (e.g. Dengler and Matthes, 2021; Grienberger et al., 2020).¹ It has also been used in an information experiment by Lergetporer et al. (2023), who study how workers adjust their training and upskilling demand when they learn the automatability of their occupation.

The automation potential estimates for the four tax-related occupations considered in this study are strikingly high (see table 1). According to the Job-Futuromat, tax clerks face a complete automation risk (100%), while certified tax specialists also exhibit a high substitutability potential (90%). Even among higher-skilled roles, auditors (50%) and tax

¹However, some highly automatable professions have still seen employment growth, indicating that factors beyond technological feasibility, such as economic demand, regulatory environments, and skill shortages, play a crucial role in the adoption of automation.

Table 1: Automation potential of tax occupations according to the Job-Futuromat.

Occupation	Automation Potential
Tax Clerk (Steuerfachangestellter)	100%
Certified Tax Specialist (Steuerfachwirt)	90%
Auditor (Wirtschaftsprüfer)	50%
Tax Advisor (Steuerberater)	30%


advisors (30%) show considerable exposure to automation.

The Information Treatment To examine how individuals respond to such information, we implement an information treatment that randomly assigns respondents into one of four groups, each receiving a different type of expert-provided automation assessment in a custom visualization:

1. **Control Group:** Only respondents' own estimates are displayed.
2. **Lower-Skilled Treatment:** Respondents estimates are compared with expert assessments for tax clerks and certified tax assistants.
3. **Higher-Skilled Treatment:** Respondents estimates are compared with expert assessments for tax advisors and auditors.
4. **Combined Treatment:** A comprehensive visualization comparing prior beliefs and expert assessments for all four listed occupations.

Figure 2: Example Screenshot of the Information Treatment

Hier sehen Sie Ihre Antworten zur Automatisierungsquote sowie die Einschätzungen des Instituts für Arbeits- und Berufsforschung (IAB).



Quelle: Die Automatisierungsquoten des IAB stammen aus dem Job Futuromat 2024.

Note: This figure presents a screenshot of the combined treatment animation assuming the respondent choose only 0 as priors for all occupations.

Source: German Business Panel Qualtrics Screenshot

Respondents assigned to one of the three treatment arms receive an animated visualization comparing their own estimates of the automation potential in their occupation to expert assessments from the Institute for Employment Research (IAB). The treatment heading states: *"Here you can see your answers on the automation rate, along with the assessments of the Institute for Employment Research (IAB)."*

The visualization consists of a dumbbell plot (see 2 for an example screenshot for the “combined treatment” arm) where each occupation is represented by a horizontal line connecting two color-coded points: the respondents own assessment (initially displayed) and the IAB estimate (revealed through animation).² The animation unfolds smoothly, starting with respondents own estimates and then progressively revealing the objective

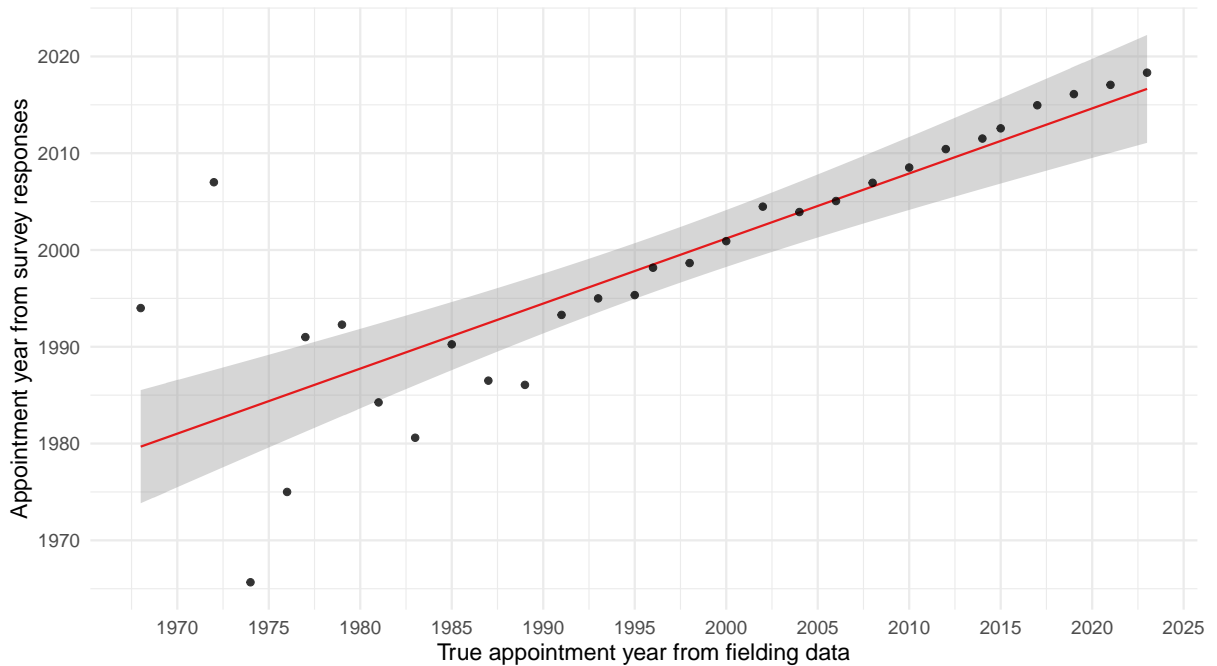
²The visualization is implemented using `d3.js`, a JavaScript library for producing dynamic, interactive data visualizations in web browsers (Bostock et al., 2011).

IAB values, visually emphasizing the gap between the two. The animation design follows best practices in visual perception research, using motion to guide attention while avoiding excessive cognitive load. For the control group only a static plot is displayed, showing the own estimates of the respondents graphically.

After the experiment, we ask whether respondents want to update their beliefs to elicit a posterior for all 4 occupations. We then proceed with several questions on hiring and firing, revenue, profit and cost and wage expectations as well as perceived automation potentials and new tasks due to automation.

2.3 Data Quality and Plausibility Checks

Figure 3: Survey and register data



Note: This figure presents a binned scatter plot comparing self-reported appointment years from the survey with official register data. Each point represents the average self-reported appointment year within 30 equally sized bins of the true appointment year from the register data. A linear fit in red, demonstrates a strong positive correlation, indicating high consistency between self-reported and official records. Deviations are most pronounced among respondents with early appointment years, likely reflecting inactive professionals who retain their designation.

Source: German Business Panel Tax Advisor and Auditor Survey 2025 and German Registers of Tax Advisors and Auditors


Ensuring the reliability and representativeness of our survey data is crucial for deriving meaningful insights about tax advisory and auditing firms. While the dataset shows

strong initial quality and alignment with expectations, we undertook a series of targeted checks to validate its plausibility, assess its representativeness, and refine it to accurately target the population of interest.

Consistency with Register Data A critical test of our datasets accuracy is the alignment between self-reported and official register data. To this end, we compare survey answers on the appointment year as tax advisor to the official register entry for each respondent. We this check in a binned scatter-plot in Figure 3, which reveals a strong positive correlation, reflecting the reliability of responses. The red trend line and confidence interval suggest that, for the majority of respondents, self-reported data closely matches the fielding data from the register. Extreme deviations are rare and especially present for the oldest respondents in our sample (i.e. the early appointment years), who are unlikely to still be economically active as tax advisors. To prevent inconsistencies, we restricted the data to cases where the reported appointment years deviated by no more than five years from the register.

Sample Representativeness Across Key Demographics To ensure our survey sample reflects the broader population of tax advisory and auditing firms, we compared sample proportions with population benchmarks from the official statistics of the chamber of tax advisors. Figure 2 illustrates deviations across self-employment status, gender, age categories, and federal states. The results demonstrate that the survey largely captures the target population, with most deviations falling within acceptable ranges. These differences primarily stem from the nature of the official register and our focus on economically active tax and auditing firms. For instance, since the titles of tax advisor and auditor are lifelong, many older professionals retain their designation despite no longer being active. Consequently, the underrepresentation of respondents over the age of 70 in our survey is expected.

Figure 4: Representativity of the survey



Note: This figure compares key demographic characteristics of survey respondents with population benchmarks from the official statistics of the Chamber of Tax Advisors, with whiskers representing 95% confidence intervals. The plotted coefficients represent differences in sample proportions relative to the population across categories such as self-employment status, gender, age groups, and federal states.

Source: German Business Panel Tax Advisor and Auditor Survey 2025 and Official Statistics of the Chamber of Tax Advisors

Filtering and Cleaning Steps While the overall data quality is high and the survey is representative of the target population we intend to capture, some filtering and cleaning was still needed. Across all filtering steps, our goal was to ensure that only active tax advisory and auditing firms remained in the dataset while excluding respondents who are in the professional register but do not operate in the relevant business segment.

First, we screened for information in open occupation and legal form fields in our survey to excluded respondents whose occupations (e.g., retirees, university professors) or organizational roles (e.g., heads of large corporate tax departments) did not align with the

target population.

Second, since not all respondents completed these fields, we applied revenue plausibility checks. Firms reporting revenue below 25,000 EUR were excluded, as such values indicate economic inactivity.³ Likewise, firms with revenues exceeding 15 million EUR or unusually high revenue per employee were flagged for manual review. Many of these cases involved corporate tax departments of firms in other economic sectors rather than independent tax firms, introducing potential bias. Where open responses or contact details confirmed this, we excluded them from the sample.

In addition, respondents who reported employment figures above 150 or revenues above 10 million Euros at their firms were reviewed, as they likely represented outliers compared to typical tax firms. Interestingly, most of these respondents are working at large international auditing firms. While we include these observations in the survey, we run most of the analyses only for smaller firms, since we only have 148 observations for these larger firms.

Finally, we excluded a small number observations with obviously erroneous answers. These included only a small number respondents who stated that they fire more than 100% of their workforce, expect revenue and cost decreases of more than 100% or cross-referenced appointment year data with the appointment years stated in the register data, removing cases where reported appointment years deviated by more than five years.

2.4 Descriptive Statistics and Covariate Balance

While our filtering and cleaning steps ensured that the final dataset accurately captured the population of active tax advisory and auditing firms, it is equally important to assess the overall characteristics of the surveyed firms. The following section presents a detailed descriptive analysis, offering insights into key demographic, employment, and revenue distributions.

Summary Statistics The summary statistics presented in Table A.1 provide a comprehensive overview of the firms surveyed in the German Business Panel Tax Advisor and Auditor Survey 2025. The average respondent is around 51 years old, with a strong representation of self-employed professionals (75%). Most respondents are tax advisors (98%), with a minority working in related auditing roles. Firm sizes vary significantly, with an average of 70 employees when all firms are included, though the distribution

³These respondents were typically beyond the typical retirement age for tax advisors, reflecting the ability of registered advisors to maintain small advisory roles past retirement.

is highly skewed. The median employment is only 9 employees. Revenue statistics exhibit a similar pattern, with a median revenue of approximately 1 million euros but a mean exceeding 3 million euros due to high-revenue outliers among large firms. Female representation stands at 32%, reflecting broader industry demographics.

Revenue and Employment Distribution Figure A.1 visualizes the distribution of firm revenue (left) and total employees (right) in more detail for our main target group of smaller tax firms with less than 150 employees and revenues below 10 million Euros. The majority of firms have revenue below 2.5 million euros and employ fewer than 50 individuals, though some large firms, primarily large multinational auditing firms, contribute to a long right tail in both distributions.

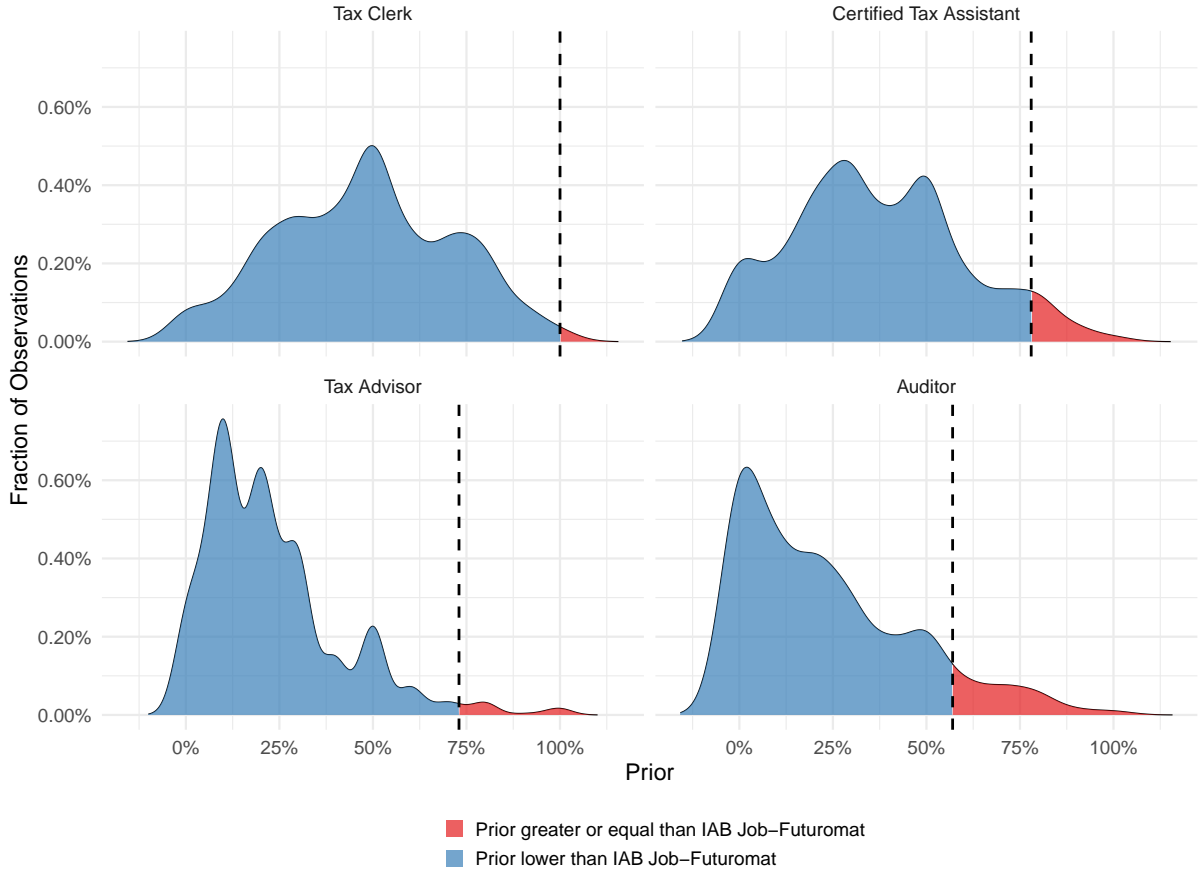
Self-reported AI usage We also elicited self-reported ai usage. Figure A.3 in the Appendix shows the current use of generative AI across firms of different sizes. Among the smallest firms (04 employees), almost half report never using generative AI, and only 14.9% use it often or always. In contrast, larger firms (11144 employees) show higher adoption rates, with less than 15% never using AI and 21.8% using it often or always.

Covariate Balance The covariate balance plot in Figure A.2 verifies the success of the randomization process. Mean differences between treatment arms and the control group remain small across all key firm characteristics, with confidence intervals largely overlapping zero. This ensures that any treatment effects observed in later analyses are not driven by pre-existing differences in firm size, revenue, or regional distribution. The balance in employment and revenue distributions further underscores the robustness of the experimental design.

3 Results

Before we analyze how information affects tax advisors' expectations and decision-making, we first examine their prior beliefs about automation risks across different occupations.

Figure 5: Distribution of Prior Beliefs



Note: This figure displays the distribution of prior beliefs about job automatability for four tax-related occupations: Tax Clerk, Certified Tax Assistant, Tax Advisor, and Auditor. The x-axis represents the subjective probability (prior) that a given occupation will be automated, while the y-axis indicates the fraction of respondents reporting each probability level. The dashed vertical lines denote the automatability estimates from the IAB Job-Futuromat benchmark. The shading differentiates between respondents whose prior beliefs are below (blue) or at least as high (red) as the benchmark estimate.


Source: German Business Panel Tax Advisor and Auditor Survey 2025

To this end, figure 5 shows the distribution of participants prior estimates of the share of automatable tasks within the next 10 years for tax clerks, certified tax assistants, tax advisors, and auditors. The vertical dashed lines represent expert assessments from the the Institute for Employment Research (IAB), allowing us to compare perceived risks and expert assessments. The results indicate that most respondents underestimate the likelihood of automation for all occupations, with a only a fraction overestimating these risks. This initial misalignment in beliefs serves as the foundation for the subsequent information updating and its consequences.

3.1 Information Updating

To understand how tax firms adjust their beliefs about automatability of their workforce when exposed to new information, we compare prior and posterior beliefs across treatment and control groups. Figure 6 plots this relationship with the x-axis representing prior beliefs and the y-axis showing posterior beliefs on the share of automateable tasks for Certified Tax Assistants. The dashed 45° line indicates no belief updating, where respondents' posteriors perfectly align with their priors.

Figure 6: Prior and Posterior Beliefs for Certified Tax Assistants



Note: This figure illustrates belief updating about automation rates for Certified Tax Assistants. The x-axis represents respondents' prior beliefs, and the y-axis shows their posterior beliefs about automation rates. Light blue dots represent individual values in the control group, while red triangles denote those in the treatment group. Larger, darker markers indicate averages for 10 quantile bins within each group. The dashed 45° line represents no belief updating.

Source: German Business Panel Tax Advisor and Auditor Survey 2025

The control group (blue dots) largely adheres to this diagonal, suggesting that in the absence of new information, respondents' beliefs remain stable. In contrast, those being in any of the four treatment arm and receiving an information treatment (red triangles) tend to shift upward, suggesting that structured information interventions lead employers to revise upward their estimates of how automatable the task of certified tax assistants are. Similar adjustments can also be observed across other tax occupations, as shown in Figure A.4 in the appendix. This adjustment is most pronounced for lower-skilled roles,

such as tax clerks and certified tax assistants, while belief updating is weaker for higher-skilled roles, such as tax advisors and auditors. However, the extent of these adjustments varies systematically with the strength of prior beliefs.

Bayesian Learning Framework for Belief Updating We formalize this belief updating process using a Bayesian learning framework, where agents combine prior beliefs with new information to form posteriors (e.g. Coibion et al., 2025):

$$\mathbf{posterior}_i = \alpha + \beta \mathbf{prior}_i + \delta \cdot \mathbf{treated}_i + \gamma \cdot (\mathbf{treated}_i \times \mathbf{prior}_i) + \varepsilon_i, \quad (1)$$

where $\mathbf{posterior}_i$ denotes the updated belief about the share of automatable tasks for respondent i , while \mathbf{prior}_i represents their initial estimate before receiving the information treatment. The variable $\mathbf{treated}_i$ is a binary indicator that takes a value of one if the respondent was assigned to any of the treatment arms and zero otherwise. The parameter δ captures the direct effect of the information treatment on belief updating, while γ measures how prior strength moderates this updating process, indicating whether individuals with stronger priors discount new information more heavily. This finding is consistent with Bayesian updating and models of rational inattention (Sims, 2003), where agents update their beliefs only when new information is perceived as sufficiently precise or novel.

Empirical Evidence of Belief Adjustment We estimate this model separately for each occupation and multiple definitions of the treatment, with results presented in Table A.2 in the appendix.⁴ The "any treatment" specification pools all respondents who received any version of the information treatment, while the "lower-skilled and combined treatment" specification isolates respondents who were exposed to information specifically about lower-skilled occupations (Tax Clerks, Certified Tax Assistants) or a combination of lower- and higher-skilled roles. Similarly, the "higher-skilled and combined treatment" specification includes respondents who received information about higher-skilled occupations (Tax Advisors, Auditors) or a mix of high- and low-skilled occupations.

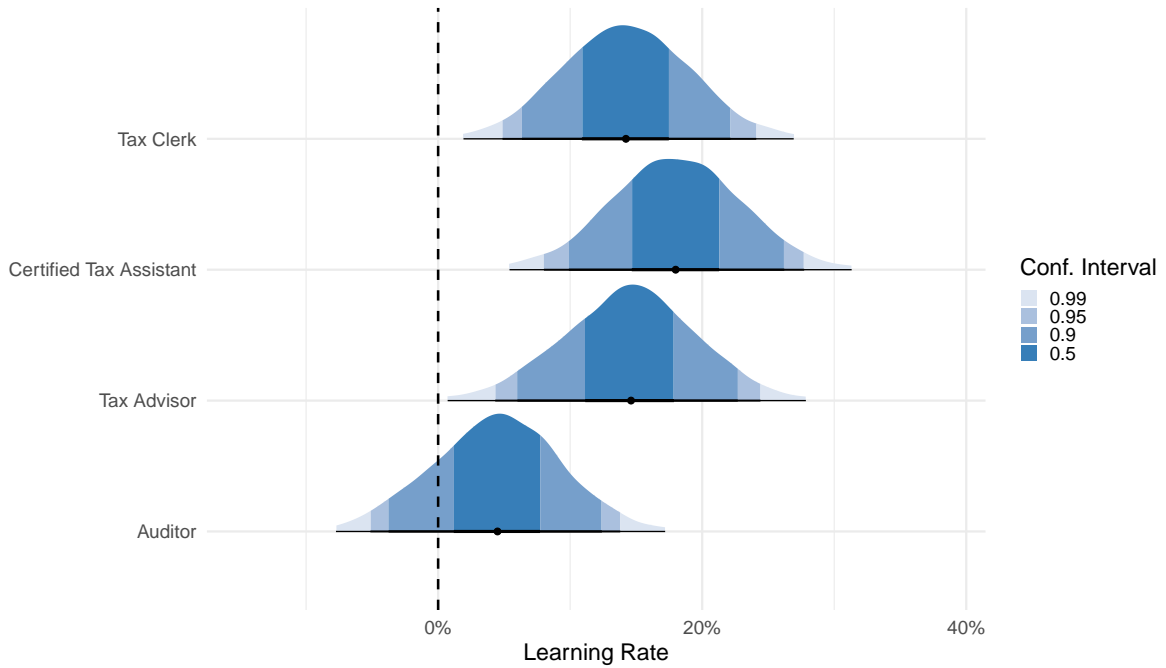
Across all specifications, the coefficient β on \mathbf{prior}_i is close to one, confirming that in the control group, firms rely almost entirely on their prior expectations when forming posterior beliefs. However, for those receiving an information treatment, the coefficient on the interaction term γ is negative and statistically significant for Tax Clerks and Tax

⁴All estimations in this table and throughout the paper were implemented in R using the `tidyverse` (Wickham et al., 2019), the `fixest` package for regressions (Bergé, 2018), and result summaries created with `modelsummary` (Arel-Bundock, 2022).

Assistants across all treatment specifications, whereas for Tax Advisors, it is significant only in the *"any treatment"* and *"higher-skilled and combined treatment"* arms. For auditors, γ is negative but not statistically significant in the *"any treatment"* arm, whereas it becomes significant in the *"higher-skilled and combined treatment"* specification.

Learning Rates and Responsiveness to Information These results suggest that belief updating is strongest and most consistent for lower-skilled occupations, where firms likely have greater uncertainty about automation risks. For higher-skilled occupations, belief updating appears to be more sensitive to how the information is framed, as evidenced by the fact that auditors update their beliefs only when exposed to a treatment specifically emphasizing automation risks for high-skilled jobs.

Figure 7: Learning rates



Note: This figure presents the estimated learning rates for different occupations in the tax advisory and auditing sector. Learning rates are derived from a regression-based Bayesian updating framework, where belief shifts in automation potential are modeled as a function of prior expectations and information treatment exposure. The density plots visualize the distribution of estimated learning rates across occupations, with shading indicating different confidence intervals (50%, 90%, 95%, and 99%). A higher learning rate suggests greater responsiveness to new information about automation rates.

Source: German Business Panel Tax Advisor and Auditor Survey 2025

A key implication of these results is that γ serves as a measure of the learning rate when $\beta \approx 1$. Since β is close to unity, prior beliefs explain most of the variation in posterior

beliefs. In this case, γ captures the extent to which respondents discount new information relative to their priors, effectively measuring how much weight is placed on the signal. A more negative γ suggests that stronger priors reduce responsiveness to new information, implying a lower learning rate.

Figure 7 visualizes this relationship by plotting the learning rate implied by the estimated interaction coefficient across occupations with their confidence intervals for the "*any treatment*" specification. The figure illustrates that learning rates are highest for tax clerks and certified tax assistants, where firms appear most responsive to new information, and lowest for auditors, where firms exhibit minimal adjustment. The confidence intervals reinforce that while all lower-skilled occupations show significant updating, the learning rate for auditors is not distinguishable from zero, confirming the weak effect of new information on beliefs about this role.

Variation in Belief Updating Across Occupations The muted updating for auditors in the "*any treatment*" specification, reflected in the negative but insignificant coefficient on γ , suggests that firms already hold relatively fixed expectations about the non-automatable nature of this role. Unlike tax clerks or assistants, where employers may have greater uncertainty about task automation, auditors are widely perceived to require high levels of expertise, regulatory compliance, and client interaction, factors that inherently limit the relevance of automation. Consequently, firms may not expect substantial automation in the first place, making new information less influential in shifting beliefs.

However, in the higher-skilled and combined treatment specification, γ is negative and statistically significant for auditors, indicating that belief updating occurs when automation risks for high-skilled occupations are explicitly emphasized. One possible explanation is that the salience of the information differs across treatments. When automation risks are framed in general terms, respondents may dismiss the information as already known or not relevant to auditors. However, when automation risks for high-skilled professionals are made explicit, respondents may perceive this information as more surprising or relevant, leading to greater belief revision.

3.2 Employment Plans and Revenue Expectations

While the results above demonstrate that employers systematically revise their beliefs about the automatability of their workforce in response to new information, shifting expectations about automation does not necessarily translate into changes in firm behavior.

Whether belief updates influence concrete decisions such as hiring, wage setting, or investment strategies remains an open question.

Instrumenting Automation Beliefs To identify the causal effect of automation beliefs on these firm-level outcomes, we exploit the experimental variation from the information treatment as an instrument. Specifically, we estimate a two-stage least squares (2SLS) model, where posterior beliefs about automation rates endogenously determined by priors and treatment assignments serve as an instrumented predictor of firm behavior. This approach allows us to isolate the exogenous variation in belief shifts and rule out potential confounders that could simultaneously affect both expectations and firm decisions.

Table 2: Employment Plans and Revenue Expectations

	(1)	(2)	(3)	(4)	(5)	(6)
	Firing	Hiring	Revenue	Profit	Costs	Wages
Predicted Posterior	-0.0003 (0.0005)	0.0004 (0.0004)	0.0013*** (0.0003)	0.0011*** (0.0003)	0.0007* (0.0003)	0.0001 (0.0001)
Constant	0.2387*** (0.0262)	0.1096*** (0.0163)	0.0684*** (0.0122)	0.0479*** (0.0116)	0.0577*** (0.0143)	0.0541*** (0.0064)
First Stage F-Stat	161.09	172.49	142.22	142.22	142.22	142.22
Prediction at Mean	-1.36 p.p.	2.04 p.p.	5.88 p.p.	5.13 p.p.	2.9 p.p.	0.5 p.p.
<i>N</i>	638	642	503	503	503	503
Adjusted R^2	-0.003	0.001	0.048	0.056	0.009	0.002

Note: This table reports the second-stage results from an instrumental variables (IV) regression estimating the effect of updated automation beliefs on firm-level outcomes. Each column corresponds to a different dependent variable: firing and hiring plans as a share of current employment (Columns 1 and 2), as well as expected future relative revenue, profit, cost, and wage growth (Columns 3-6). The instrument for posterior automation beliefs is the randomized information treatment, using a treatment definition where assignment to any treatment arm is considered treated. Additionally, priors and posteriors for certified tax assistants are used. Standard errors are reported in parentheses. $p < 0.1$, $p < 0.05$, and $p < 0.01$ denote significance at the 10%, 5%, and 1% levels, respectively. The "Prediction at Mean" row reports the estimated treatment effect for respondents with the mean fitted posterior belief, rather than the marginal effect of an incremental belief update.

Table 2 presents the results of this IV estimation for key firm-level outcomes. The first-stage used for these estimation is roughly similar to the relation displayed in Figure 6.

Likewise, we use the prior and posterior beliefs about the automatability of Certified Tax Assistants, where the learning rate is the highest, as well as a treatment indicator that considers respondents that are assigned to any treatment arm as treated. However,

we explore choices of occupational priors/posteriors in the first stage in table A.3 in the appendix. These robustness checks confirm that the main results hold when varying the occupational group used to define prior and posterior beliefs and the results are qualitatively the same for different specifications.

The Effects on Employment Plans and Revenue Expectations Each column in the table corresponds to a different outcome variable. Columns (1) and (2) report estimated effects on firing and hiring plans between 2025 and 2027, expressed as a share of current employment. Columns (3) through (6) display expectations regarding future relative revenue, profit, cost, and wage growth over the same period.

Notably, the first-stage F-statistics across all specifications are well above conventional thresholds, confirming the strength of our instrument and ensuring that the variation in posterior beliefs induced by the treatment provides a valid source of exogenous variation for identifying the causal effects on firm outcomes.

The results reveal that updated automation beliefs have no significant effect on firing or hiring decisions, suggesting that firms do not anticipate immediate employment adjustments in response to learning about the automatability of their workforce.⁵

However, belief shifts are associated with a significant increase in expected revenue and profit growth (Columns 3 and 4), consistent with the idea that firms anticipate efficiency gains from automation but do not expect immediate labor displacement.

Interestingly, cost expectations also increase significantly (Column 5), though at a smaller magnitude than revenue and profit, suggesting that firms foresee some additional expenditures, potentially on AI adoption or skill investments. At the same time, wage expectations remain unaffected (Column 6), indicating that firms do not plan to share productivity gains with employees in the form of higher wages, at least in the near term.

We also report the "Prediction at Mean" for each estimation, which provides an estimate of the treatments average effect for firms with the mean fitted posterior belief, rather than reflecting the impact of an incremental increase in the posterior belief itself. For instance, the estimated 5.88 percentage point increase in revenue expectations suggests that, at the average learning rate, treated firms anticipate a substantial productivity-driven revenue gain. Similarly, the estimated 5.13 percentage point increase in profit expectations underscores that firms view automation as a net positive for profitability.

By contrast, the estimated effect for costs, while positive and significant, is more modest,

⁵This is line with research by Grienberger et al. (2020), who found that employment in several tax occupations with very high automation potential was still growing in 2023.

indicating that firms expect efficiency gains but foresee only limited additional expenditures relative to revenue growth. Meanwhile, the near-zero estimate for wage expectations suggests that despite expecting higher revenue and profit, firms do not intend to pass these gains to employees through higher salaries. These findings align with the research on rent-sharing, that shows that firms often do not share productivity gains with employees through higher wages, especially when such gains are derived from automation or technological advancements. For instance, Kline et al. (2019) found that workers capture only a fraction of the surplus generated by patents, with significant disparities based on tenure and position within the firm. Similarly, Cho and Krueger (2022) observed that rent-sharing within firms is uneven, favoring higher-earning employees.⁶

3.3 Revenue Expectation by Occupation

To further examine how belief updating affects firms economic expectations, Table 3 presents the estimated effect of updated posterior beliefs on expected revenue growth across different tax-related occupations due to automation. We again use the same IV strategy, compared any treatment arm with the control group, and use the appropriate prior/posterior pair for each occupation

The estimates reveal that across all occupations, higher posterior beliefs about automatability are associated with significantly higher expected revenue growth. This suggests that firms anticipating greater automation in these roles expect efficiency gains to translate into revenue increases. However, the magnitude of this effect varies across occupations. Using the Mean Posterior increase for each role, we find that for Tax Clerks, this increase in revenue per hour is at 27.85 percentage points. For Certified Tax Assistants, the expected revenue per hour increase is 20.41 percentage points. Among higher-skilled roles, the expected revenue effect is smaller.

⁶This might be in part because the wage schedule for most tax occupations is fixed. This is in line with research by Franceschelli et al. (2010), who show that productivity gains under performance pay schemes translate more directly into higher wages compared to fixed-wage schemes, where productivity improvements have a more limited effect on employee compensation, even when both types of workers achieve similar productivity increases.

Table 3: Expected Revenue per Hour Changes by Occupation

	(1)	(2)	(3)	(4)
	Certified			
	Tax Clerk	Tax Assistant	Tax Advisor	Auditor
Predicted Posterior	0.0048*** (0.0009)	0.0044*** (0.0011)	0.0032*** (0.0009)	0.0039*** (0.0012)
Constant	-0.0062 (0.0460)	0.0196 (0.0430)	0.0658*** (0.0218)	0.0140 (0.0303)
First Stage F-Stat	162.81	157.73	159.11	137.72
Prediction at Mean	27.85 p.p.	20.41 p.p.	9.49 p.p.	12.07 p.p.
<i>N</i>	597	565	611	505
Adjusted R^2	0.083	0.090	0.092	0.063

Note: This table reports the IV second stage estimates for the relationship between predicted posterior beliefs about automation rates and expected revenue per hour across four tax-related occupations: Tax Clerks, Certified Tax Assistants, Tax Advisors, and Auditors. The dependent variable is expected revenue growth per hour. The respective posteriors and priors beliefs are used for each occupation. Standard errors are reported in parentheses. Asterisks denote statistical significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

These results again indicate that firms revise their revenue expectations more strongly for lower-skilled occupations, such as Tax Clerks and Certified Tax Assistants, where belief updating is more pronounced. In contrast, the revenue per hour adjustment is more muted for higher-skilled occupations, such as Tax Advisors and Auditors, where prior beliefs were already relatively fixed, limiting the scope for large belief revisions. These findings align with our previous results, showing that firms update their beliefs most strongly for roles where automation uncertainty is highest and, in turn, anticipate greater revenue growth from automation in these occupations.

3.4 Automation Potential and New Tasks

The results thus far have demonstrated that firms systematically update their beliefs about automation risks when exposed to new information. The evidence presented in Table 4 further reveals that this belief updating extends beyond broad occupational categories to more granular task-level assessments.

For this table, we again rely on our standard instrumental variable specification to estimate the effect of updated automation beliefs on the perceived automation potential of specific tax-related tasks. Here, the dependent variables are binary indicators equal to one if a respondent believes that a given task such as tax filing, payroll accounting, tax consulting, succession advisory, or international tax advisory has automation potential.

Table 4: Tasks with Automation Potential

	(1)	(2)	(3)	(4)	(5)
	Tax Filing	Payroll Accounting	Tax Consulting	Succession Advisory	International Tax
Predicted Posterior	0.0047*** (0.0007)	0.0038*** (0.0007)	0.0022*** (0.0007)	0.0018*** (0.0005)	0.0016*** (0.0006)
Constant	0.5197*** (0.0475)	0.5988*** (0.0472)	0.0929** (0.0431)	-0.0279 (0.0252)	-0.0073 (0.0303)
First Stage F-Stat	211.11	211.11	211.11	211.11	211.11
Prediction at Mean	26.79 p.p.	21.81 p.p.	12.42 p.p.	10.54 p.p.	9.41 p.p.
<i>N</i>	852	852	852	852	852
Adjusted R^2	0.066	0.042	0.009	0.009	0.011

Note: This table reports the second-stage results from an instrumental variables (IV) regression estimating the effect of updated automation beliefs on a set of indicator variables for questions referring to tasks in tax occupation. Each question elicits whether respondent see automation potential in a specific part of their occupations. The instrument for posterior automation beliefs is the randomized information treatment, using a treatment definition where assignment to any treatment arm is considered treated. Additionally, priors and posteriors for certified tax assistants are used. Standard errors are reported in parentheses. $p < 0.1$, $p < 0.05$, and $p < 0.01$ denote significance at the 10%, 5%, and 1% levels, respectively. The "Prediction at Mean" row reports the estimated treatment effect for respondents with the mean fitted posterior belief, rather than the marginal effect of an incremental belief update.

Higher Automation Potential for Lower-Skilled Tasks Consistent with our earlier findings, respondents exposed to the information treatment report significantly higher automation potential for lower-skilled tasks, particularly those traditionally performed by Tax Clerks and Certified Tax Assistants. The estimated treatment effects at the mean learning rate are largest for Tax Filing (Mean Effect: 26.79 p.p.) and Payroll Accounting (Mean Effect: 21.81 p.p.), two tasks that are already subject to considerable automation through existing software solutions. This aligns with our previous results, where belief updating and learning rates were strongest for these occupations (see Figure 7).

Interestingly, even in the control group, baseline automation expectations for these tasks are already high. The constant terms for Tax Filing (0.5197) and Payroll Accounting

(0.5988) suggest that more than half of respondents already associate these tasks with automation potential, even without exposure to new information. This reinforces the idea that automation in compliance-heavy, structured processes is well understood by tax professionals.

Some Perceived Automation Potential for Higher-Skilled Tasks Interestingly, treated respondents also report increased automation potential for tasks typically performed by Tax Advisors and Auditors, such as Succession Advisory (Mean Effect: 10.54 p.p.) and International Tax Consulting (Mean Effect: 9.41 p.p.). While the magnitudes of these effects are lower, they remain statistically significant, suggesting that even for traditionally high-skilled, judgment-intensive roles, respondents see some scope for automation.

Unlike lower-skilled tasks, baseline automation expectations for these high-skilled roles are close to zero in the control group, as reflected in the near-zero constant estimates (-0.0279 for Succession Advisory and -0.0073 for International Tax). This suggests that respondents do not naturally perceive automation as relevant for these tasks unless they receive explicit information highlighting high automation rates for their profession. The fact that treated respondents revise their beliefs even for high-skilled tasks suggests that AI-based tools and automation solutions are beginning to shape expectations beyond purely routine work. While automation expectations remain strongest for procedural and compliance-related work, firms do not entirely discount the potential for AI-driven automation even in high-skilled advisory roles.

Beyond revising their expectations about the automation potential of existing tasks, respondents also anticipate the emergence of new tasks as a consequence of automation. This raises the question of how firms expect job roles to evolve in response to automation—whether they foresee a net displacement of tasks or an expansion into new responsibilities that complement AI-driven workflows.

To explore this, we extend our analysis to investigate whether respondents see new tasks emerging as a result of automation. Table 5 presents results from an instrumental variables (IV) regression, estimating the effect of updated automation beliefs on the perceived relevance of specific new tasks that could gain importance in an increasingly automated work environment. The dependent variables are binary indicators for whether respondents expect Legal Tech, Compliance, Prompt Engineering, or Quality Assurance to become relevant as part of their evolving job responsibilities.

Table 5: New Tasks due to Automation

	(1)	(2)	(3)	(4)
	Legal		Prompt	Quality
	Tech	Compliance	Engineering	Assurance
Predicted Posterior	0.0026*** (0.0008)	0.0022** (0.0010)	0.0022** (0.0009)	0.0016** (0.0007)
Constant	0.0826* (0.0474)	0.4152*** (0.0603)	0.1550*** (0.0524)	0.7844*** (0.0435)
First Stage F-Stat	178.45	178.45	178.45	178.45
Prediction at Mean	14.96 p.p.	12.94 p.p.	12.61 p.p.	9.14 p.p.
<i>N</i>	797	797	797	797
Adjusted R^2	0.011	-0.001	0.010	0.007

Note: This table reports the second-stage results from an instrumental variables (IV) regression estimating the effect of updated automation beliefs on a set of indicator variables for questions referring to new tasks due to automation. Each question elicits whether respondent see find new tasks relevant in case of automation. The instrument for posterior automation beliefs is the randomized information treatment, using a treatment definition where assignment to any treatment arm is considered treated. Additionally, priors and posteriors for certified tax assistants are used. Standard errors are reported in parentheses. $p < 0.1$, $p < 0.05$, and $p < 0.01$ denote significance at the 10%, 5%, and 1% levels, respectively. The "Prediction at Mean" row reports the estimated treatment effect for respondents with the mean fitted posterior belief, rather than the marginal effect of an incremental belief update.

The results reveal that higher automation beliefs significantly increase the likelihood of respondents considering new tasks relevant, though the effect sizes vary across task types. At the mean learning rate, the estimated treatment effects are largest for Legal Tech (14.96 p.p.) and Compliance (12.94 p.p.), suggesting that tax professionals increasingly view legal automation tools and regulatory compliance work as key areas where their responsibilities might expand.

Similarly, Prompt Engineering (12.61 p.p.) the skill of designing and refining AI-generated output emerges as another notable area where respondents foresee potential task shifts. While still a relatively new concept in professional tax work, this suggests that some tax professionals are beginning to anticipate the growing role of AI interaction and optimization as part of their job, reflecting broader labor market trends where demand for AI-related skills has increased across diverse occupations, often accompanied by wage

premiums (Alekseeva et al., 2021).

In contrast, Quality Assurance (9.14 p.p.) shows a more muted, yet still significant effect. The high baseline constant for Quality Assurance (0.7844) indicates that even in the control group, many respondents already see this as an important part of tax work. Unlike Legal Tech or Prompt Engineering, where automation may create entirely new areas of expertise, Quality Assurance may be perceived as a natural extension of existing responsibilities, focused on ensuring accuracy in AI-assisted tax processes rather than fundamentally transforming job roles.

The variation across tasks highlights different levels of perceived complementarity between AI and human expertise. Legal Tech and Compliance are seen as clear areas for task expansion, potentially requiring upskilling in legal automation and regulatory monitoring, while Prompt Engineering suggests early recognition of AI interaction skills as a new job component. Quality Assurance, in contrast, remains a core responsibility, likely focusing on mitigating AI errors rather than creating entirely new workstreams. Rather than replacing professionals, automation appears to be driving a transition toward augmented work, where human oversight and AI-driven processes coexist. This aligns with the notion that automation often complements rather than substitutes for human labor, particularly by creating new tasks where human expertise remains essential (e.g. Acemoglu and Restrepo, 2019; Brynjolfsson and McAfee, 2014). As Acemoglu and Restrepo (2019) argue, the net impact of automation on labor demand depends on the balance between task displacement and the emergence of novel roles that leverage human comparative advantages. Similarly, Brynjolfsson and McAfee (2014) highlight that while routine tasks are increasingly automated, new roles emerge that require advanced cognitive and interactive skills, reinforcing the idea that technology reshapes job content rather than simply eliminating work.

4 Preliminary Conclusion

Our findings reveal several important insights into how tax advisory and auditing firms perceive and respond to automation trends. First, we observe a significant gap between initial employer beliefs and expert assessments regarding automation potential. While firms tend to underestimate automation risks, our information intervention successfully prompts belief updating, particularly for lower-skilled occupations like tax clerks and certified tax assistants. For higher-skilled roles such as tax advisors and auditors, belief adjustments are more limited, suggesting that firms perceive greater barriers to automation at the top of the professional hierarchy.

Despite updating their beliefs about automation, firms do not immediately revise their hiring or firing plans, indicating that automation-induced workforce reductions are not a primary concern in the near term. However, firms that update their automation expectations anticipate higher revenue and profit growth, consistent with the notion that efficiency gains from automation may enhance firm performance rather than lead to immediate labor displacement. Interestingly, while firms foresee increased costs potentially due to investment in new technologies or upskilling initiatives, wage expectations remain largely unchanged, implying that anticipated productivity gains are not expected to translate into higher employee compensation in the short run.

Moreover, our results highlight that automation is not merely perceived as a labor-replacing force but as a driver of job transformation. Firms exposed to updated automation information expect new tasks to emerge, particularly in areas like legal tech, compliance, and AI interaction roles such as prompt engineering. This suggests that while automation reshapes job content, it also creates opportunities for skill development and specialization rather than rendering professional roles obsolete.

These findings contribute to the broader economic literature on automation and labor markets in three key ways. First, they extend existing research beyond manufacturing and manual labor-intensive industries, showing how generative AI might influence a broad spectrum of white-collar occupations. Second, by taking an employer-centered perspective, our study captures anticipatory responses to automation before they materialize as observable labor market outcomes, offering a forward-looking view of technological adaptation. Finally, our results underscore the nuanced impact of automation on task composition, revealing that rather than reducing overall employment, automation may shift the skill demands of the workforce in ways that require continued investment in complementary human capital.

References
















- ACEMOGLU, D., D. AUTOR, J. HAZELL, AND P. RESTREPO (2022): “Artificial Intelligence and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 40(S1), S293–S340. Cited on page 4.
- ACEMOGLU, D., AND P. RESTREPO (2019): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33(2), 3–30. Cited on page 28.
- (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128(6), 2188–2244. Cited on pages 2 and 3.
- AGHION, P., C. ANTONIN, S. BUNEL, AND X. JARAVEL (2020): “What are the labor and product market effects of automation? New evidence from France,” Discussion Paper DP14443, Centre for Economic Policy Research (CEPR). Cited on page 3.
- AGHION, P., C. ANTONIN, S. BUNEL, AND X. JARAVEL (2022): “The Effects of Automation on Labor Demand: A Survey of the Recent Literature,” Discussion Paper DP16868, Centre for Economic Policy Research (CEPR). Cited on page 3.
- ALEKSEEVA, L., J. AZAR, M. GINÉ, S. SAMILA, AND B. TASKA (2021): “The demand for AI skills in the labor market,” *Labour Economics*, 71, 102002. Cited on page 28.
- ALESINA, A., A. MIANO, AND S. STANTCHEVA (2022): “Immigration and Redistribution,” *The Review of Economic Studies*, 90(1), 1–39. Cited on page 4.
- AREL-BUNDOCK, V. (2022): “modelssummary: Data and Model Summaries in R,” *Journal of Statistical Software*, 103(1), 1–23. Cited on page 18.
- ARNTZ, M., S. BLESSE, AND P. DOERRENBURG (2022): “The End of Work is Near, Isn’t It? Survey Evidence on Automation Angst,” *ZEW - Centre for European Economic Research Discussion Paper*, 22(036). Cited on page 4.
- BERGÉ, L. (2018): “Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm,” Discussion paper, CREA Discussion Papers 13. Cited on page 18.
- BESSEN, J., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2020): “Firm-level automation: Evidence from the Netherlands,” in *AEA Papers and Proceedings*, vol. 110, 389–393. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203. Cited on page 3.
- (2025): “What Happens to Workers at Firms that Automate?,” *The Review of Economics and Statistics*, 107(1), 125–141. Cited on page 4.
- BONFIGLIOLI, A., R. CRINÒ, G. GANCIA, AND I. PAPADAKIS (2024): “Artificial intelligence and jobs: evidence from US commuting zones*,” *Economic Policy*, 40(121), 145–194. Cited on page 3.
- BOSTOCK, M., V. OGIEVETSKY, AND J. HEER (2011): “D³: Data-Driven Documents,”

- IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2301–2309. Cited on page 10.
- BRYNJOLFSSON, E., D. LI, AND L. R. RAYMOND (2023): “Generative AI at Work,” Working Paper 31161, National Bureau of Economic Research. Cited on page 2.
- BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York. Cited on page 28.
- CHO, D., AND A. B. KRUEGER (2022): “Rent Sharing within Firms,” *Journal of Labor Economics*, 40(S1), S17–S38. Cited on page 23.
- COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018): “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review*, 108(9), 2671–2713. Cited on page 4.
- COIBION, O., Y. GORODNICHENKO, AND M. WEBER (2025): “Tell Me Something I Dont Already Know: Learning in Low- and High-Inflation Settings,” *Econometrica*, 93(1), 229–263. Cited on page 18.
- DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2021): “The Adjustment of Labor Markets to Robots,” *Journal of the European Economic Association*, 19(6), 3104–3153. Cited on pages 2 and 3.
- DENGLER, K., AND B. MATTHES (2021): “Folgen des technologischen Wandels für den Arbeitsmarkt: Auch komplexere Tätigkeiten könnten zunehmend automatisiert werden,” *IAB-Kurzbericht*, 13, 8, Zugriff am 13. Februar 2025. Cited on page 8.
- DENGLER, K., B. MATTHES, AND W. PAULUS (2014): “Occupational Tasks in the German Labour Market: An Alternative Measurement on the Basis of an Expert Database,” Discussion Paper FDZ-Methodenreport 12/2014 (en), Institute for Employment Research (IAB), Nürnberg. Cited on page 8.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2023): “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” Discussion Paper 2303.10130, arXiv. Cited on page 2.
- FELTEN, E. W., M. RAJ, AND R. SEAMANS (2023): “Occupational Heterogeneity in Exposure to Generative AI,” *Available at SSRN 4414065*. Cited on page 2.
- FRANCESCHELLI, I., S. GALIANI, AND E. GULMEZ (2010): “Performance Pay and Productivity of Low- and High-Ability Workers,” *Labour Economics*, 17(5), 886–898. Cited on page 23.
- GATHMANN, C., F. GRIMM, AND E. WINKLER (2024): “AI, Task Changes in Jobs, and Worker Reallocation,” Discussion paper, CESifo Working Paper No. 11585. Cited on page 3.
- GRIENBERGER, K., B. FITZENBERGER, C. KAGERL, AND B. MATTHES (2020): “Der IAB-Job-Futuromat: Beschäftigungsentwicklung und Fachkräfteengpässe variieren mit dem Substituierbarkeitspotenzial,” *IAB-Forum*. Cited on pages 8 and 22.

- JEFFREY, K. (2021): “Automation and the future of work: How rhetoric shapes the response in policy preferences,” *Journal of Economic Behavior & Organization*, 192, 417–433. Cited on page 4.
- KLINE, P., N. PETKOVA, H. WILLIAMS, AND O. ZIDAR (2019): “Who Profits from Patents? Rent-Sharing at Innovative Firms,” *The Quarterly Journal of Economics*, 134(3), 1343–1404. Cited on page 23.
- KOCH, M., I. MANUYLOV, AND M. SMOLKA (2021): “Robots and Firms,” *The Economic Journal*, 131(638), 2553–2584. Cited on page 3.
- KUZIEMKO, I., M. I. NORTON, E. SAEZ, AND S. STANTCHEVA (2015): “How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” *American Economic Review*, 105(4), 1478–1508. Cited on page 4.
- LERGETPORER, P., K. WEDEL, AND K. WERNER (2023): “Automatability of Occupations, Workers’ Labor-Market Expectations, and Willingness to Train,” IZA Discussion Papers 16687, Institute of Labor Economics (IZA). Cited on pages 4 and 8.
- NOY, S., AND W. ZHANG (2023): “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381(6654), 187–192. Cited on page 2.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of Monetary Economics*, 50(3), 665–690. Cited on page 18.
- WICKHAM, H. ET AL. (2019): “Welcome to the tidyverse,” *Journal of Open Source Software*, 4(43), 1686. Cited on page 18.
- WISWALL, M., AND B. ZAFAR (2015): “Determinants of College Major Choice: Identification using an Information Experiment,” *The Review of Economic Studies*, 82(2 (291)), 791–824. Cited on page 4.

A Additional figures and tables


Table A.1: Summary Statistics

	Mean	SD	Min	P25	Median	P75	Max	Histogram
Female (Dummy)	0.32	0.47	0.00	0.00	0.00	1.00	1.00	
Self-Employed (Dummy)	0.75	0.43	0.00	1.00	1.00	1.00	1.00	
Personell Responsibility (Dummy)	0.88	0.33	0.00	1.00	1.00	1.00	1.00	
Tax Advisor (Dummy)	0.98	0.12	0.00	1.00	1.00	1.00	1.00	
Legal Form GmbH (Dummy)	0.26	0.44	0.00	0.00	0.00	1.00	1.00	
Legal Form Partner Association (Dummy)	0.22	0.42	0.00	0.00	0.00	0.00	1.00	
Age	51.01	10.82	28.00	44.00	52.00	58.00	80.00	
Total Employees	70.39	625.15	0.00	4.00	9.00	20.00	13300.00	
Revenue (thousand EUR)	32694.13	263154.29	25.00	430.00	1000.00	2500.00	3000000.00	
Firing share (next three years)	0.22	0.24	0.00	0.06	0.15	0.33	1.00	
Hiring share (next three years)	0.14	0.16	0.00	0.03	0.10	0.17	1.00	
Expected Revenue Change (Pct.)	12.70	12.77	-15.00	0.00	10.00	20.00	100.00	
Expected Profit Change (Pct.)	10.75	19.08	-15.00	1.00	10.00	15.00	400.00	
Expected Cost Change (Pct.)	8.89	19.25	-30.00	0.00	9.00	10.00	400.00	
Expected Wage Change (Pct.)	5.82	5.06	-20.00	3.00	5.00	10.00	35.00	

Note: This table presents summary statistics for the full survey data including big firms with more than 150 employees or more than 10 Mio. Euro of revenue.

Source: German Business Panel Tax Advisor and Auditor Survey 2025


Figure A.1: Distribution of Firm Revenue and Employment



Note: This figure presents the distribution of firm revenue (left) and total employees (right) among survey respondents. The revenue distribution is displayed in thousand euros and excludes firms with revenues above 10 million euros. The employee distribution is truncated at 150 employees. Histograms illustrate the relative frequency of firms within each range, while boxplots provide additional insight into the spread and presence of outliers. The distributions confirm the presence of a highly skewed firm size distribution, with most firms being relatively small but a subset of large firms contributing to long right tails.

Source: German Business Panel Tax Advisor and Auditor Survey 2025


Figure A.2: Covariate Balance across Treatment Arms



Note: This figure displays the mean differences in covariates between each treatment arm and the control group, with whiskers representing 95% confidence intervals. The results indicate no systematic imbalances across key firm characteristics, including firm size, revenue, regional composition, and workforce demographics, confirming that the randomization was successful.

Source: German Business Panel Tax Advisor and Auditor Survey 2025


Figure A.3: Current Generative AI Use by Firm Size



Note: This figure shows the frequency of generative AI use across firms of different sizes: 03, 47, 810, and 11144 employees. Categories include "Never," "Rarely," "Sometimes," "Often," and "Always."

Source: German Business Panel Tax Advisor and Auditor Survey 2025

Figure A.4: Belief Upadting by Occupation



Note: This figure illustrates belief updating about automation rates for all four elicited occupations. The x-axis represents respondents' prior beliefs, and the y-axis shows their posterior beliefs about automation rates. Light blue dots represent individual values in the control group, while red triangles denote those who receive any treatment. Larger, darker markers indicate averages for 10 quantile bins within each group. The dashed 45° line represents no belief updating.

Source: German Business Panel Tax Advisor and Auditor Survey 2025

Table A.2: Regression: Updating of Beliefs by Occupation and Treatment Type

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Any treatment arm				Lower skilled and combined treatment				Higher skilled and combined treatment			
	Dep. Var: Posterior Automatability				Dep. Var: Posterior Automatability				Dep. Var: Posterior Automatability			
	Tax Clerk	Tax Assistant	Tax Advisor	Auditor	Tax Clerk	Tax Assistant	Tax Advisor	Auditor	Tax Clerk	Tax Assistant	Tax Advisor	Auditor
Treatment \times Prior	-0.143*** (0.042)	-0.180*** (0.038)	-0.145** (0.066)	-0.044 (0.044)	-0.113** (0.047)	-0.150*** (0.043)	-0.103 (0.072)	-0.040 (0.046)	-0.162*** (0.049)	-0.197*** (0.042)	-0.184** (0.078)	-0.105** (0.049)
Treatment	15.417*** (2.433)	15.483*** (1.864)	9.579*** (1.423)	5.464*** (1.853)	16.675*** (2.793)	16.038*** (2.138)	8.353*** (1.577)	5.502*** (2.000)	15.268*** (2.761)	15.634*** (2.104)	11.823*** (1.699)	8.446*** (2.062)
Prior	0.913*** (0.025)	0.947*** (0.021)	0.949*** (0.043)	0.883*** (0.036)	0.913*** (0.025)	0.947*** (0.021)	0.949*** (0.043)	0.883*** (0.036)	0.913*** (0.025)	0.947*** (0.021)	0.949*** (0.043)	0.883*** (0.036)
Constant	6.473*** (1.530)	3.716*** (1.077)	2.424*** (0.853)	6.125*** (1.535)	6.473*** (1.532)	3.716*** (1.078)	2.424*** (0.853)	6.125*** (1.536)	6.473*** (1.532)	3.716*** (1.078)	2.424*** (0.853)	6.125*** (1.536)
<i>N</i>	857	803	873	729	647	614	661	556	641	598	650	547
F-Stat												
Adjusted R^2	0.64	0.65	0.63	0.67	0.66	0.67	0.68	0.69	0.67	0.70	0.65	0.65

Note: This table presents the results of multiple regressions examining how respondents update their beliefs about automation rates across different tax occupations in response to the information treatment. Each column corresponds to a separate regression model for a specific occupation. The treatment indicator is defined in three different ways: (i) receiving any treatment, (ii) receiving the treatment for lower-skilled occupations or the combined treatment, and (iii) receiving the treatment higher-skilled occupations or the combined treatment. Respondents in the treatment groups are compared to those in the control group.

Table A.3: Employment and Revenue Plans for different Posterior/Prior Specifications

Using the Prior and Posterior for Tax Clerks						
	(1)	(2)	(3)	(4)	(5)	(6)
	Firing	Hiring	Revenue	Profit	Costs	Wages
Predicted Posterior	-0.0001 (0.0005)	0.0003 (0.0004)	0.0014*** (0.0003)	0.0013*** (0.0003)	0.0005 (0.0003)	0.0001 (0.0001)
Constant	0.2335*** (0.0310)	0.1121*** (0.0210)	0.0489*** (0.0141)	0.0243* (0.0145)	0.0608*** (0.0159)	0.0513*** (0.0078)
First Stage F-Stat	179.40	189.99	130.92	130.92	130.92	130.92
Prediction at Mean	-0.54 p.p.	1.74 p.p.	7.82 p.p.	7.44 p.p.	2.66 p.p.	0.8 p.p.
<i>N</i>	663	667	533	533	533	533
Adjusted R^2	-0.002	0.002	0.045	0.054	0.004	-0.002
Using the Prior and Posterior for Tax Advisors						
	Firing	Hiring	Revenue	Profit	Costs	Wages
Predicted Posterior	0.0006 (0.0006)	0.0004 (0.0004)	0.0016*** (0.0004)	0.0012*** (0.0003)	0.0008* (0.0004)	0.0001 (0.0002)
Constant	0.2105*** (0.0184)	0.1197*** (0.0122)	0.0800*** (0.0102)	0.0646*** (0.0088)	0.0660*** (0.0112)	0.0560*** (0.0049)
First Stage F-Stat	179.95	191.28	124.92	124.92	124.92	124.92
Prediction at Mean	1.79 p.p.	1.02 p.p.	4.63 p.p.	3.37 p.p.	2.2 p.p.	0.25 p.p.
<i>N</i>	663	669	538	538	538	538
Adjusted R^2	0.002	-0.000	0.050	0.040	0.014	-0.000
Using the Prior and Posterior for Auditors						
	Firing	Hiring	Revenue	Profit	Costs	Wages
Predicted Posterior	0.0001 (0.0005)	0.0006 (0.0004)	0.0013*** (0.0004)	0.0013*** (0.0004)	0.0007* (0.0004)	0.0002 (0.0001)
Constant	0.2223*** (0.0175)	0.1129*** (0.0117)	0.0835*** (0.0109)	0.0606*** (0.0094)	0.0677*** (0.0112)	0.0552*** (0.0044)
First Stage F-Stat	167.57	172.98	125.80	125.80	125.80	125.80
Prediction at Mean	0.18 p.p.	1.95 p.p.	3.97 p.p.	3.73 p.p.	1.97 p.p.	0.45 p.p.
<i>N</i>	589	593	463	463	463	463
Adjusted R^2	-0.001	-0.002	0.029	0.051	0.014	0.002

Note: This table presents robustness checks for the results shown in Table 2, which estimate the impact of updated automation beliefs on firm-level employment and revenue expectations. Specifically, we explore alternative specifications using prior and posterior beliefs for different occupational groups: Tax Clerks, Tax Advisors, and Auditors as instruments. Each column corresponds to a different outcome variable: firing and hiring plans (Columns 1 and 2), as well as expected future relative revenue, profit, cost, and wage growth (Columns 3-6). The predicted posterior represents the instrumented automation belief, using treatment-induced variation.

B Survey and experimental design

Figure B.1: Screenshot of the Survey Introduction



Note: This figure presents a screenshot of the opening of the German Business Panel Tax Advisor and Auditor Survey 2025 survey.

Source: German Business Panel Tax Advisor and Auditor Survey 2025

Table B.1: Relevant questions from the GBP Tax Advisor survey

Question Number	Question	Answer Options
Q1	What is your current employment status?	- Employed - Self-employed
Q2	Which of the following positions best describes your role?	- Board Member/Executive Management - Senior Partner - Partner - Director - Senior Manager - Manager - Senior Consultant - Consultant, expert, analyst - student, intern
Q3	In your role as selected role : Do you have personnel responsibility?	- Yes - No
Q4	In your professional role: How often do you work with AI-powered tools that generate text independently? For example, ChatGPT, Claude, etc.	- Always - Often - Sometimes - Rarely - Never
Q5	In your company: How many employees are working in the following professions? Please provide the number in full-time equivalents.	- Tax consultant [0,100000] - Chartered accountant [0,100000] - Tax specialist [0,100000] - Tax clerk [0,100000]
Q6	What do you estimate: How much of the core activities in the following professions can be automated by 2024? Please provide a percentage.	- Tax consultant [0,100] - Chartered accountant [0,100] - Tax specialist [0,100] - Tax clerk [0,100]
Q7	If you think again: What do you think now? Would you like to adjust your information?	- Tax consultant [0,100] - Chartered accountant [0,100] - Tax specialist [0,100] - Tax clerk [0,100]

Table B.2: Relevant questions from the GBP Tax Advisor survey

Question Number	Question	Answer Options
Q8	From your company’s perspective: Which of the following areas of responsibility have emerged due to automation in tax consulting?	<ul style="list-style-type: none"> - Quality control of automation results - Data protection and compliance monitoring - Prompt engineering - Application and support of legal tech/large language models (LLMs) - Other areas of responsibility
Q9	From your company’s perspective: How many new employees do you plan to hire in the coming years? How many of them will be for new areas of responsibility created by automation? Note: Please indicate the number of new hires in each year in full-time equivalents.	<div>2025 2026 2027</div> <ul style="list-style-type: none"> Total new hires Of which employees for areas of responsibility created by automation
Q10	Regarding your personnel planning: How would you proceed if tasks could be replaced by automation? Please provide the number of affected staff.	<div>2025 2026 2027</div> <ul style="list-style-type: none"> Assign employees new tasks Dismiss employees
Q11	How do you perceive the changes in your profession due to automation?	<ul style="list-style-type: none"> - As a threat - As an opportunity for professional development - Neither a threat nor an opportunity
Q12	Which profession would you most likely switch to?	<ul style="list-style-type: none"> - Public Accounting - Tax Consulting - Tax Technology Expert - Prompt Engineer - Data Scientist - No change
Q13	Given the level of automation in your occupation, how likely is it that you would change occupation? Note: 0% (no career change) - 100% (career change)	[0,100]

Table B.3: Relevant questions from the GBP Tax Advisor survey

Question Number	Question	Answer Options
Q14	How many new employees do you plan to hire in the following occupations in total by 2027? Note: Please indicate the number of new employees in each year in full-time equivalents.	<ul style="list-style-type: none"> - Tax consultant [0,1000] - Chartered accountant [0,1000] - Tax specialist [0,1000] - Tax clerk [0,1000]
Q15	In which area do you see automation potential in your company?	<ul style="list-style-type: none"> - Business consulting - Financial accounting - International tax law - Payroll accounting - Succession planning - Tax consulting - Tax declaration
Q16	There are now several new AI solutions for tax advisors. Would you like to learn more about examples of such AI solutions?	<ul style="list-style-type: none"> - Yes - No
Q17	Have you ever heard of or actively used one of these AI solutions for tax advisors? - Taxy.io: A platform that develops AI solutions specifically for tax advisors. This tool analyzes tax questions and provides precise answers based on specialized literature. - DATEV LexInform AI Assistant (LEA): An AI solution that supports tax advisors in researching legal documents by providing relevant information and sources (e.g., UStAE, BMF letters).	<ul style="list-style-type: none"> - Yes - No
Q18	How frequently do you use these AI solutions?	<ul style="list-style-type: none"> - Use them regularly - Use them irregularly - Do not use them

Table B.4: Relevant questions from the GBP Tax Advisor survey

Question Number	Question	Answer Options
Q19	Do you plan to use AI solutions in the future?	- Yes - No
Q20	What increase in turnover per working hour do you expect for the following professions? Note: Please indicate the expected percentage change (positive or negative values).	- Tax Advisor - Auditor - Tax Clerk - Tax Assistant
Q21	Compared to today: How does your company plan to adjust the average hourly wage for all employees in the next 12 months? Note: Please enter the change in per cent. You can enter positive or negative values.	- Change in hourly wage in per cent
Q22	How much time do you plan to spend on your own digital training in an average week in the future? Note: Please enter the value in hours.	
Q23	Is your company planning investments or further training on automation topics for employees?	- Yes - No
Q24	What kind of investments or further training on automation topics is your company planning?	- General further training (e.g. part-time study) - Specialized further training (e.g. certified fibutronics) - Investments in hardware and software (e.g. ChatGPT, computer) - Other
Q25	What do you estimate for your company? By what percentage will the following variables change through the use of AI? Note: Please enter a value in percent. You can enter positive or negative values.	- Profit Change - Revenue change - Cost change

Table B.5: Relevant questions from the GBP Tax Advisor survey

Question Number	Question	Answer Options
Q26	When were you appointed as a tax consultant or auditor? Note: Please enter the year of your appointment.	<ul style="list-style-type: none"> - year of appointment - (not yet) appointed
Q27	How would you like to be addressed in a greeting?	<ul style="list-style-type: none"> - Mr - Ms - Not specified
Q28	When were you born? Note: Please enter your year of birth.	
Q29	What is the legal form of your company?	<ul style="list-style-type: none"> - Sole proprietorship - GmbH - GmbH and Co. KG - UG - AG - oHG - GbR - PartG - KG - SE - Verein - KGaA - Genossenschaft - Public-law company - Other
Q30	Please enter the annual revenue (in EUR) of your company in the previous calendar year. Note: Please enter a whole number without using thousands or decimal separators.	
Q31	If you could not or did not want to answer our question on revenue, do you think you could at least give us a range in which your revenue lies. Which of the following intervals most closely corresponds to your company's annual revenue in the previous calendar year?	Intervals from less than 50,000 EUR to more than 60,000,000 EUR
Q32	Do you have any comments or questions? Your opinion is important to us!	