Belief Updating and AI Adoption by Entrepreneurs

Marcos Balmaceda* Juan Pedro Ronconi†

Abstract

Generative AI is free to use and can raise productivity on many tasks, yet entrepreneurs and small firms adopt it slowly and use it shallowly. A natural explanation is that entrepreneurs underestimate how much AI could improve their business, so correcting that belief should raise adoption. We test this in a pre-registered field experiment with 736 Chilean entrepreneurs, randomly assigned to vignettes that show a fictional entrepreneur applying ChatGPT to business tasks, or to an uninformative control. The vignettes shift beliefs: profit expectations rise in a Bayesian pattern, concentrated among the most pessimistic participants. Yet behavior does not follow. Willingness to pay for AI coaching and take-up of a free AI course are unchanged in the same session as the update—ruling out belief decay—and a four-month follow-up finds no effect on AI use, business practices, sales, or profits. The null holds even among the strongest updaters. Asked what holds them back, entrepreneurs name material constraints (money, time, and data privacy), not a failure to see AI's relevance. Correcting profit beliefs is not enough to move adoption. It leaves untouched the costs of acting and the firm-specific work of making a general-purpose tool work in one's own business.

Keywords: technology adoption, artificial intelligence, entrepreneurship, belief updating, micro-entrepreneurs

JEL codes: O33, L26, D83, O12

Disclosure: This research did not receive external funding. Neither author has any employment, consulting, or ownership relationship — active within the last three years — with any firm that could be substantially affected by this research, nor any other financial relationship or potential conflict of interest to disclose. This statement applies to both authors and to their spouses and partners.

*KU Leuven, Department of Management, Strategy and Innovation, Naamsestraat 69, 3000 Leuven, Belgium. Email: marcos.balmaceda@kuleuven.be.

†Universidad de los Andes, Chile, School of Business and Economics, Av. Monseñor Álvaro del Portillo 12455, Santiago, Chile. Email: jpronconi@uandes.cl.

1 Introduction

Conversational AI tools like ChatGPT are free to use, require no capital investment, and can raise measured productivity on a range of knowledge-work tasks (Brynjolfsson et al. 2025; Noy and Zhang 2023); yet among entrepreneurs and small firms, adoption remains low and use remains shallow, concentrated in larger firms (OECD et al. 2025; Fernández and Uriarte 2026). A natural explanation is that entrepreneurs underestimate how much AI could improve their own business, so that correcting this belief would be enough to move them to adopt. The premise is reasonable, and it is general: across development, finance, and management, a large class of information interventions presumes that correcting a biased belief is enough to change behavior. We designed, pre-registered, and implemented a field experiment with 736 Chilean entrepreneurs to test it, randomly assigning participants to informational vignettes depicting a fictional entrepreneur applying ChatGPT to her own business, or to an uninformative control. The vignettes do correct beliefs: profit expectations rise, and the most pessimistic entrepreneurs revise the most, the pattern Bayesian updating predicts. Behavior, however, does not follow. Neither incentivized willingness to pay for AI coaching nor take-up of the free AI course rises in the same session in which beliefs change, and a four-month follow-up shows no effect on AI use, business practices, or financial outcomes. Correcting beliefs about AI's profitability is not sufficient to produce adoption.

The prior evidence gives strong reason to expect otherwise. When Dominican students received accurate information about the returns to secondary education, enrollment rose substantially (Jensen 2010). When workers discovered that peers earned more than expected, they revised their wage expectations upward and firms adjusted compensation in response (Cullen and Perez-Truglia 2022). When households received information that shifted their inflation expectations, spending moved in the predicted direction (Coibion et al. 2020). In each case, credible information corrected a biased prior and behavior followed accordingly. Among entrepreneurs, overconfident beliefs drive entry and investment decisions (Åstebro et al. 2014) and ambiguity attitudes shape how they act on new information (Åstebro et al. 2026); below-market beliefs about technology returns suppress adoption among farmers and small firms (Foster and Rosenzweig 1995; Conley and Udry 2010). A vigorous line of work views entrepreneurship itself as disciplined belief revision under uncertainty, in which entrepreneurs act on beliefs they form and test against evidence, updating as they go (Camuffo et al. 2020; Koning et al. 2022; Gambardella and Messinese 2025; Agrawal et al. 2021). Our setting strips away the usual obstacles to this channel. Micro-entrepreneurs are sole decision-makers, with

no organizational layer between an updated belief and an adoption decision, and conversational AI tools are free, require no capital investment, and run on demand for everyday business tasks. These are about as favorable as conditions get for belief correction to move behavior, so if a corrected belief cannot move adoption here, it is hard to see where it would.

Why might so strong a prediction fail? Generative AI is not a domain-specific tool with a predefined use case, like a high-yield seed or a deworming pill; it is a general-purpose technology whose returns are task-contingent and realized only through firm-specific discovery and adaptation (Bresnahan and Trajtenberg 1995; Bresnahan 2024; Goldfarb et al. 2023). Learning that AI raises profits on average tells an entrepreneur that the technology is worth pursuing; it does not tell her which of her own tasks it improves, or how to rework them around it. That search through her own operations is one a generic signal cannot perform, because the signal conflates whether AI is valuable with how to apply it (Agrawal et al. 2021; Dell'Acqua et al. 2026; Aragonès et al. 2005). On this view, a corrected profit belief is not sufficient on its own: it settles whether AI is worth pursuing while leaving the firm-specific discovery and implementation problem intact. The setting thus poses a clean choice between two credible views—belief correction as the binding constraint, or belief correction as just one of several inputs that putting a general-purpose technology to work in one's own business requires.

One might object that the vignettes moved only an abstract belief that AI is useful in general, leaving entrepreneurs unable to see its relevance to their own firm. The design rules this out. The vignettes depict AI applied to specific business tasks, and the elicited belief is each entrepreneur's expectation for her *own* profits. Entrepreneurs do come to see AI as relevant; what they do not do is act. Asked in the follow-up to name their main barriers to adopting AI, they point overwhelmingly to material and practical constraints. Cost leads (39%, "no money to pay for AI tools"), followed by data privacy (19%) and time to learn (17%); only 10% select "I don't see how it could help my business specifically" and 9% distrust the quality of AI's output.[1] These are constraints on acting, not on understanding, and a corrected profit belief leaves each untouched. They echo a long literature on small firms in developing economies, held back less by what they know than by binding constraints on capital, managerial time, and cognitive bandwidth (Bruhn et al. 2010; McKenzie and Woodruff 2017; Mani et al. 2013). We read these as candidate explanations, not identified

[1] Respondents could select up to three barriers (mean 1.5 selections); we report shares among follow-up respondents who answered the item. A residual 33% selected "Other," which had no free-text follow-up and is therefore uninterpretable, and 5% selected none. While "Other" limits what we can read from the question, the low uptake of the interpretable *cognitive* barriers—seeing AI's relevance and trusting its quality—indicates these channels have little traction as self-reported primary barriers, which is itself consistent with our experimental result.

mechanisms (our design varies beliefs, not costs), and organize them in the discussion as the two conditions a corrected belief leaves unmet: the costs of acting and the firm-specific application problem.

We recruited entrepreneurs through ChileConverge, a non-profit digital platform that offers free online training courses and mentoring sessions to Chilean entrepreneurs. We randomly assigned them to one of two vignette treatments or an uninformative control. [2] Both treatment vignettes depict a fictional micro-entrepreneur using ChatGPT to improve her business outcomes. A second, independently randomized experiment invited half of all participants to a free AI training course roughly two months later. A follow-up survey, fielded approximately four months after the baseline, then measured AI adoption, business practices, sales, and profits. Baseline characteristics—sector, age, education, sales, profits, and prior AI use—were observed at intake; beliefs and willingness to pay were measured immediately after treatment. Prior AI use is broad but shallow. Most have tried AI for their business, yet roughly 60% do not use it regularly and only 16% use it daily. The design lets us exogenously shift beliefs about the profitability of using generative AI and then observe adoption from several vantage points, from incentivized willingness to pay in the same session to behavior four months out.

The vignettes raise profit expectations by 0.128 points on a seven-point scale ($p = 0.030$), with the revision concentrated among the most pessimistic entrepreneurs, who move roughly half a point (Treatment $\times$ Low Optimism $= 0.549$, $p = 0.002$), the textbook Bayesian pattern in which a larger prior-signal gap produces a larger update. Yet none of this moves behavior. Willingness to pay for personalized AI coaching is unchanged in the same session ($p = 0.253$), and take-up of the free AI course—a behaviorally revealed measure recorded in platform data—is equally flat. Four months later we find no effect on AI use, new-adopter status, business practices, sales, or profits. The null holds even among the strongest updaters, where the belief shift was largest.

These findings connect three literatures. First, they speak to the view of entrepreneurship as disciplined belief revision under uncertainty (Agrawal et al. 2026; Camuffo et al. 2020; Koning et al. 2022; Gambardella and Messinese 2025; Agrawal et al. 2021), in which entrepreneurs act on beliefs they form and test. We isolate the updating step, showing that exogenously correcting the profitability prior does not move action, even among the strongest updaters. Where prior AI-information experiments infer the constraint from a behavioral non-response, we measure the belief update and show the binding step lies downstream of it,

[2] The experiment is pre-registered at the AEA RCT Registry (AEARCTR-0016544) and received ethical approval from the Social and Societal Ethics Committee of KU Leuven (PRET G-2025-9674-R2(MIN)).

among conditions our design holds fixed rather than varies. Second, and most broadly, the result marks a limit of belief-based interventions for general-purpose technologies whose uses must be discovered. A belief update concentrated exactly as Bayesian theory predicts produces no response on any pre-registered behavioral measure, and because willingness to pay is elicited in the same session as the update, belief decay is ruled out by construction. The limit reaches any information campaign whose target belief is not the one that governs action (Haaland et al. 2023; Jensen 2010; Cullen and Perez-Truglia 2022; Coibion et al. 2020). Third, the paper speaks to why small firms fail to adopt profitable technologies and practices (Conley and Udry 2010; Åstebro 2002; McKenzie and Woodruff 2017; Gertler et al. 2025; Kim 2025; Ronconi 2025). This literature documents valuable, free information left unused but rarely observes the belief itself. By eliciting that belief immediately before and after information provision, we separate a failure to persuade from persuasion that is not sufficient, and show that the latter is the case. Concurrent experiments converge on a familiar pattern, yet the focus has not been on isolating the entrepreneur's own profit-belief channel. Cullen et al. (2025) vary information about competitors' adoption, which moves robotics investment but not AI; Menkhoff (2025) move productivity beliefs in established firms, where the response is gated by whether the respondent holds decision authority; and Otis et al. (2024) grant access to an AI assistant, which helps mainly those able to implement it. Our design isolates that channel by moving and *measuring* the profit belief among sole decision-makers, for whom no authority or access friction intervenes, and eliciting behavior in the same session. A belief shifted exactly as theory predicts does not move behavior; when correction is this complete and adoption is still unchanged, communicating AI's value more effectively will not raise it, because whatever binds adoption lies outside the belief our signal moves.

2 Theory and Hypothesis Development

2.1 Information and Belief Updating

Bayesian entrepreneurship treats entrepreneurial action as theory-laden belief revision under uncertainty: entrepreneurs hold heterogeneous and relatively strong priors about the opportunities they pursue, and update them via Bayes' rule as they learn (Agrawal et al. 2026; Gambardella et al. 2026; Camuffo et al. 2020). The relevant uncertainty is thus not Knightian (Gans 2026): beliefs about returns are probabilistic and revisable in response to a signal, not immeasurable. The framework distinguishes *simple* Bayesian learning, in which the posterior depends on the prior and the signal, from *deep* learning, which turns on how the prior was formed

and on grasping the causal structure that links a signal to one's own venture (Kim 2026). Our experiment isolates simple updating: a passively acquired signal about AI's profitability, and whether it revises the prior. A growing experimental literature studies such information provision (for a review, see Haaland et al. 2023), treating belief revision as Bayesian updating: an agent with prior belief μ_i^0 (precision τ_p) receives a signal s_i (precision τ_s) and forms a posterior:

$$\mu_i^1 = \alpha \cdot s_i + (1 - \alpha) \cdot \mu_i^0 \tag{1}$$

where $\alpha = \tau_s/(\tau_s + \tau_p)$ is the learning rate, representing the weight placed on new information relative to the prior, ranging from 0 (no updating) to 1 (complete updating). A higher ratio of signal precision to prior precision implies more weight on the new signal; for instance, if $\tau_s = 1$ and $\tau_p = 1$, then $\alpha = 0.50$ and the posterior places equal weight on the signal and the prior. Two predictions follow directly. First, treated agents should revise beliefs toward the signal. Second, agents whose priors are further from the signal should revise more. From Equation (1), the belief revision equals $\mu_i^1 - \mu_i^0 = \alpha(s_i - \mu_i^0)$: the magnitude of updating is proportional to the gap between the signal and the prior. We assume the vignette conveys a signal $s_i = \bar{s}$ above the prior for treated entrepreneurs ($\bar{s} > \mu_i^0$), constant across individuals; τ_s and τ_p are not separately identified, and only their ratio α is estimated. Under this assumption, entrepreneurs who are initially more pessimistic (lower μ_i^0) face a larger gap and therefore revise more.

We test these predictions by estimating:

$$\mu_i^1 = \gamma_0 + \gamma_1 T_i + \gamma_2 \mu_i^0 + \gamma_3 (T_i \times \mu_i^0) + \delta_s + \varepsilon_i \tag{2}$$

where $T_i \in \{0, 1\}$ pools T1 and T2 into a single treatment indicator, δ_s are strata fixed effects, and γ_2 captures how prior beliefs map into posterior beliefs for the control group, that is, the persistence of priors absent new information. The sum $\gamma_2 + \gamma_3$ captures this mapping for treated entrepreneurs. A negative γ_3 indicates that treatment compresses the prior–posterior relationship, as treated entrepreneurs place less weight on priors because they incorporate the experimental signal. We recover α by equating the Bayesian structural prediction with the OLS regression estimates, a method-of-moments approach that does not require the prior or signal to be expressed in monetary units. The control-group coefficient γ_2 captures the persistence of measured priors absent any signal; it falls below 1 to the extent that beliefs are measured with error or regress to the mean on a bounded scale. The treated-group coefficient $\gamma_2 + \gamma_3$ reflects that same persistence

scaled by the weight the prior retains once the signal is incorporated, $(1 - \alpha)$, so that $\gamma_2 + \gamma_3 = \gamma_2(1 - \alpha)$. Solving yields the ratio estimator $\alpha = -\gamma_3/\gamma_2$, which normalizes by the empirically estimated control-group persistence rather than imposing $\gamma_2 = 1$ (Cullen and Perez-Truglia 2022).[3]

The size of the updating response varies systematically with signal type. Cullen and Perez-Truglia (2022) estimate $\alpha = 0.69$ for salary information among employees, where the signal directly concerns the respondent's own expected earnings. Cullen et al. (2025) find $\alpha = 0.26$ for information about competitors' AI and robotics adoption, where the signal concerns aggregate technology diffusion rather than a personal financial outcome. The contrast reflects a general pattern: signals that are precise and personally diagnostic yield high learning rates, while signals about aggregate returns to a broadly applicable technology yield lower ones. Our vignettes fall in the second category, conveying evidence about how a fictional entrepreneur benefits from AI in general rather than about the respondent's own firm, and we therefore expect a learning rate in the range of Cullen et al. (2025).

Hypothesis 1. *Providing information about AI's business applications updates entrepreneurs' beliefs about AI's profitability, with larger revisions for entrepreneurs holding more pessimistic priors.*

2.2 From Beliefs to Technology Adoption

The technology adoption literature establishes that beliefs about a technology's expected returns are a central determinant of adoption (Foster and Rosenzweig 1995; Conley and Udry 2010; Suri 2011). In the standard framework, an agent who learns that a technology's returns exceed her prior expectation revises upward and, if the revised expectation exceeds the adoption threshold, adopts. The logic extends naturally from belief updating to action. If information corrects a downward-biased prior, the agent should be more willing to invest in the technology. This prediction has been confirmed across decisions as varied as schooling, wage-setting, and household spending (Jensen 2010; Cullen and Perez-Truglia 2022; Coibion et al. 2020): a credible piece of information corrects a biased prior, and behavior adjusts.

For micro-entrepreneurs, this link should be especially direct. The decision at stake is not building custom AI infrastructure but using a free conversational interface like ChatGPT for everyday tasks (clarifying documents, drafting marketing copy, basic financial analysis), low-cost integrations well within a micro-

[3]The alternative estimator $\alpha = 1 - (\hat{\gamma}_2 + \hat{\gamma}_3)$ assumes $\gamma_2 = 1$ exactly, whereas the ratio estimator does not impose this restriction and is therefore robust to attenuation in the control-group persistence coefficient γ_2, which can fall below 1 under measurement error or regression to the mean on a bounded scale. We adopt the ratio estimator for comparability with Cullen and Perez-Truglia (2022) and Cullen et al. (2025).

enterprise's reach. And unlike in larger organizations, where updated beliefs must travel through hierarchies, approvals, and budget cycles, a micro-entrepreneur is the sole decision-maker: the person whose beliefs shift is the one who decides whether to adopt. A downward-biased prior is therefore the natural candidate for the binding constraint, one that exogenously correcting it can test directly.

The strength of this prediction depends on the nature of the technology. For a domain-specific innovation—a high-yield seed, a vaccine, a household appliance—the use case is predefined, and the path from "it is profitable" to "I will adopt it" is short. A general-purpose technology is different: its value is not intrinsic but is realized only by matching it to a specific application within the firm's own operations (Bresnahan and Trajtenberg 1995; Bresnahan 2024; Goldfarb et al. 2023; Grégoire and Shepherd 2012; Dushnitsky and Klueter 2016). For generative AI this matching is unusually hard, because the boundary between the tasks it improves and those it degrades is irregular and not intuitive even to experts (Dell'Acqua et al. 2026); it cannot be read off aggregate evidence about average returns but must be found by searching one's own operations, and the capacity to recognize and apply new knowledge itself builds on what one already knows (Aragonès et al. 2005; Cohen and Levinthal 1990).

In the Bayesian-entrepreneurship account, adoption follows only when the belief that governs action crosses a value–cost threshold, and two features of our setting separate that belief from the one an information signal moves. First, the account prises apart two beliefs the standard view runs together: the *aggregate* belief that the technology is worth pursuing, and the *application-specific* (or *operative*) belief that actually governs action, namely which of one's own tasks to change and how. A single signal about average returns conflates the quality of the technology with the efficacy of any particular way of applying it; only the entrepreneur's own testing separates the two (Agrawal et al. 2021). In this account, the same holds even when the signal *demonstrates* use through a worked example: seeing how the technology serves another business can raise the aggregate belief, and point to where it might apply, yet leave unconveyed the firm-specific knowledge the operative belief requires. That knowledge comes from *deep*, theory-laden learning, generated by interpreting one's own data or by advice the entrepreneur coproduces and tests, rather than from the *simple* updating a passively received signal supports (Kim 2026; Miller et al. 2024). Second, a profitability update does not relax the costs of acting on it. For the micro-entrepreneurs we study these are not capital or integration costs, since conversational AI tools are free and require no infrastructure; they are the *perceived* time and effort to learn to use AI well—to prompt effectively in one's own business context—together with the cognitive bandwidth that experimentation demands under financial strain (Mani et al. 2013).

We treat these as candidate constraints that a corrected profit belief leaves untouched, not as conditions we identify: our design varies beliefs alone, and we return to them in the discussion, anchored in the barriers the entrepreneurs themselves report. The noncapital, learning-related costs of adoption have long been found to predict technology adoption more powerfully than capital outlays (Åstebro 2002, 2004), and personalized coaching of the kind we price is precisely what would lower them. This yields a clean test between the two accounts. Under the standard adoption-threshold view, a downward-biased profit belief is the binding constraint—as it plausibly is for sole decision-makers facing a free, low-infrastructure tool—so correcting it should raise willingness to pay for personalized AI implementation coaching and intended adoption. The Bayesian-entrepreneurship account predicts the opposite: because the belief that governs action is the application-specific one, reached through the deep, firm-specific learning a passive signal does not supply, correcting the aggregate profit belief should leave willingness to pay and adoption unmoved. The two accounts coincide only where the application is obvious or the entrepreneur already holds the operative belief; for a general-purpose tool whose uses must be discovered, they diverge. We state the standard view's prediction as our pre-registered hypothesis, against which the Bayesian-entrepreneurship account predicts a null:

Hypothesis 2. *Information that updates beliefs about AI's profitability increases entrepreneurs' willingness to pay for a personalized AI implementation coaching session and their intended adoption of AI tools.*

3 Research Setting

Our experiment took place within ChileConverge, a non-profit impact corporation dedicated to strengthening entrepreneurship development in Chile. The organization operates a digital community platform providing business development resources, training, and peer networking support; roughly 4,800 entrepreneurs and micro-enterprises across the country actively used the platform in the year preceding the experiment.[4] Chilean micro-enterprises face constraints characteristic of Latin American developing contexts: limited access to formal credit, high informality, and limited exposure to structured business training (Martínez et al. 2018). Participants are digitally connected, accessing the platform online and reporting prior familiarity with AI tools, but their businesses operate under resource constraints qualitatively different from those of

[4] Platform administrative records shared by ChileConverge at the start of the project (April 2025 export): 4,808 of the registered users had logged into the platform within the preceding twelve months.

micro-entrepreneurs in advanced economies.

Participants are 736 micro-entrepreneurs drawn from ChileConverge's digital platform. Table 1 provides detailed summary statistics. The typical entrepreneur is 44 years old and female (63%); nearly two-thirds work five or more days per week on their business. Roughly 45% have higher education. Entrepreneurs are distributed across all 16 of Chile's regions, with 47% concentrated in the Santiago Metropolitan Region and the remainder spread across Valparaíso (12%), Biobío (7%), Araucanía (6%), and twelve other regions. Businesses span 13 sectors, the largest being retail and commerce (20%), artisanal and cultural products (11%), professional services (10%), and manufacturing (8%), and most operate at very small scale. The median business has one full-time employee (the entrepreneur), has operated for three years, and earns monthly sales of about 350,000 CLP (US$390) and profits of 200,000 CLP (US$220); means are higher and right-skewed (Table 1). At baseline, consistent AI use is limited: although 73% have used AI for their business at least once, roughly 60% do not use it regularly in the business (less often than several times per week), 46% make at most occasional use of AI even for non-business purposes, and only 16% use it daily to support their businesses. Expectations, by contrast, are optimistic: 84% of entrepreneurs expect their profits to increase if they were to adopt AI in their business, 14% expect no change, and fewer than 2% expect a decrease. Most entrepreneurs thus already hold positive beliefs about AI's potential, while the minority expecting no gain leaves room for upward revision.

[Insert Table 1 here.]

4 Research Design

Figure 1 summarizes the experimental sequence. We describe each stage in turn.

[Insert Figure 1 here.]

4.1 Experiment 1: Information Treatments

Eligible participants operated an active business and completed the baseline survey in full. We stratified the sample by gender, education level, and geographic macro-region (five zones: Santiago, North, Central, South-Central, South), then randomly assigned participants within strata to one of three conditions: Business Consultant (T1, $N = 245$), Creative Assistant (T2, $N = 253$), or Control ($N = 238$).

All three conditions are embedded in the baseline survey, administered via Qualtrics under a pseudonymous protocol. ChileConverge generated anonymized identifiers for registered platform members, which we used for stratified randomization and survey delivery without observing individual identities. After answering questions about their business characteristics, participants complete the initial belief elicitation: they reported their expected change in monthly profits (and separately in sales) if they were to use AI regularly in their business, on a seven-point scale anchored at "decrease by more than 20%" (1) and "increase by more than 20%" (7), with "no change" at the midpoint (4). These pre-treatment responses constitute each entrepreneur's prior belief. Respondents were then shown one of three texts depending on their assigned condition.

The two treatment vignettes describe a fictional micro-entrepreneur, a macramé bracelet seller on Instagram, who uses ChatGPT for business purposes, followed by an example ChatGPT conversation. Both share the same protagonist, visual format, length (T1: 124 words; T2: 116 words), number of task examples (three), actionable verbs (five), and minimum reading time (45 seconds). They differ in the domain of the examples and the format of the ChatGPT response. Business Consultant (T1) frames AI as an expert business advisor. The conversation shows the entrepreneur asking about pricing strategy and receiving a structured pricing formula with five action steps; task examples cover pricing optimization, cost management, and inventory planning. Creative Assistant (T2) frames AI as a practical creative tool. The conversation shows the entrepreneur requesting help with Instagram Reel content and receiving a production script with tips and hashtags; task examples cover social media posts, product naming, and visual differentiation. Across both arms, the task examples span domains the business-practices literature links to micro-enterprise performance (pricing, cost and inventory management, and marketing; McKenzie and Woodruff 2017), echoing the practices the training course later targets (Section 4.2); in each arm, the worked conversation *demonstrates* AI applied to a concrete task rather than merely asserting that AI is profitable. The Control group receives a text describing how ChatGPT works from a technical perspective, covering the transformer architecture, the training process, and recent improvements in response quality (citing the Stanford *AI Index* report), without any information about business applications or profitability. We pool the two treatment arms in our main specifications and report arm-specific estimates throughout.

Immediately after the vignette, we re-elicited beliefs about AI's impact on profits and sales, yielding within-session belief revision. We also measured willingness to pay for a personalized AI implementation coaching session with a business-and-AI faculty expert from Universidad de los Andes, tailored to the

entrepreneur's own business, using a multiple price list with five binary choices at escalating thresholds (5,000 to 115,000 CLP), following Cullen and Perez-Truglia (2022). To incentivize truthful revelation and survey completion, each wave includes two lottery prizes: a tablet drawn randomly among all completers, and the training session itself awarded to one randomly selected participant at a value corresponding to their stated WTP.[5]

4.2 Experiment 2: AI Training Course

Approximately two months after the baseline, we launched a second, independently randomized experiment. Approximately half of all baseline participants, across all Experiment 1 arms including control, received an invitation to enroll in a free online AI training course. The invitation is randomized independently of Experiment 1 treatment assignment, using the same stratification variables. Non-invited participants could not access the course; the platform infrastructure kept it hidden from their interface, preventing contamination across invitation arms. The course curriculum covers four core business practices (marketing, inventory management, record-keeping, and financial planning), building on the framework of McKenzie and Woodruff (2017), who identify these as the domains where management practices most directly affect micro-enterprise performance. Each module teaches prompt engineering techniques applied to the entrepreneur's own business context. Of the 240 participants in our follow-up survey, 98 (40.8%) have at least partial course exposure and 21 (8.8%) complete the full course. Because course completion is self-selected, all course-related analyses are interpreted as descriptive rather than causal.

A follow-up survey, administered approximately four months after the baseline (roughly two months after the course invitation), measured AI usage, business practices, financial outcomes, and prompt quality for all participants who could be re-contacted ($N = 240$ with complete data).

4.3 Empirical Strategy

The outcome Y_i is one of the following: post-treatment profit expectations (H1); WTP for a personalized AI implementation coaching session or course take-up (H2); or, at the follow-up, AI adoption frequency, new adopter status, business practices (the McKenzie-Woodruff 2017 19-item index), monthly sales, monthly profits, or PESI prompt sophistication score. Our primary estimating equation is:

[5] An internal rule, undisclosed to participants, awards the session winner a gift card valued at the maximum amount corresponding to their stated WTP, applied independently in the baseline and follow-up waves. Prize winners are contacted by ChileConverge rather than the research team, preserving double-blind anonymization.

$$Y_i = \beta_0 + \beta_1 T1_i + \beta_2 T2_i + \boldsymbol{\theta}' \mathbf{X}_i + \delta_s + \varepsilon_i \tag{3}$$

where $T1_i$ and $T2_i$ are treatment indicators, $\mathbf{X}_i$ is a vector of pre-specified ANCOVA controls (baseline profit expectations, business age, winsorized sales and profits, and prior AI use), δ_s are strata fixed effects, and standard errors are heteroskedasticity-robust (HC1). We also estimate a pooled specification replacing the two treatment dummies with a single indicator $T_i = \mathbf{1}[T1_i + T2_i > 0]$. The Experiment 1 pre-analysis plan designated AI adoption as the primary outcome; the Experiment 2 plan designated business practices (the McKenzie-Woodruff 2017 19-item index). Table 3 reports Equation (3) with full ANCOVA controls for the beliefs outcomes; Table 4 reports the adoption results with ANCOVA controls and a course invitation indicator, as pre-specified for the panel follow-up.

5 Data and Measurement

5.1 Baseline Measures

Our primary measures map directly onto the two hypotheses. For H1, we elicited profit and sales expectations using a seven-point scale (1 = decrease of more than 20%, 2 = decrease of 10–20%, 3 = decrease of up to 10%, 4 = no change, 5 = increase of up to 10%, 6 = increase of 10–20%, 7 = increase of more than 20%). Entrepreneurs reported these expectations before and after the vignette, yielding within-session belief revision; our main belief outcome is expected profit change. For H2, we use two behavioral measures. Willingness to pay for a personalized AI training session with a specialized faculty expert, tailored to the entrepreneur's own business context, is elicited both immediately before and immediately after the vignette, in the same session, via a multiple price list (five sequential binary choices at escalating thresholds, 5,000 to 115,000 CLP) from which we construct WTP as the lower bound of the switching interval (the last accepted price) (Cullen and Perez-Truglia 2022). The before-and-after design yields a within-person measure of how treatment shifts WTP, including an indicator for whether an entrepreneur's own WTP rose; because a respondent's budget is fixed across the minutes between the two elicitations, this differences out liquidity as an explanation for any non-response. WTP was elicited with real incentives. One randomly selected participant per wave received the training session as a prize at the monetary equivalent of their stated WTP, providing incentive for truthful revelation. Course take-up is a binary indicator, constructed from platform records,

of whether the entrepreneur started the free AI training course offered in Experiment 2; we also observe course completion. A complementary measure, expressed interest in AI training (elicited post-treatment in the baseline survey), turned out to be near-universal (97.9% in the control group) and therefore carries little usable variation; we report it in Appendix A for completeness.

For heterogeneity analysis, we classify entrepreneurs as having low optimism if their baseline profit expectation is at or below the neutral midpoint (≤ 4), and as optimistic otherwise; the continuous prior is used as a robustness check. Seven observations are excluded from baseline regressions due to missing stratification-cell identifiers ($N = 729$).

5.2 Follow-Up Measures

At the follow-up, we measured AI adoption through two complementary indicators: self-reported usage frequency on a five-point scale (1 = never, 5 = always) and a binary "new adopter" indicator for entrepreneurs who did not use AI at baseline but reported using it at follow-up. Business practices were assessed using the McKenzie and Woodruff (2017) 19-item index, covering marketing, inventory management, record-keeping, and financial planning practices, each scored 0/1 for a total ranging from 0 to 19. Monthly sales and profits are self-reported by participants at both the baseline and follow-up surveys, and are winsorized at the 99th percentile to limit the influence of outliers. For robustness, we also examine these outcomes on a log scale using log(1 + sales) and log(1 + profits), a standard transformation for monetary outcomes; the log specification confirms the null (Table A14, Appendix A).[6]

We also elicit prompt engineering sophistication. Entrepreneurs are asked: "Imagine you want to ask ChatGPT for advice to increase your sales. How would you phrase the question to get the best possible answer?" Open-ended responses are scored on the Prompt Engineering Sophistication Index (PESI), which captures two dimensions: technique (role assignment, context provision, and structured output requests; scored 0–3) and specificity (references to the entrepreneur's own product, numerical targets, timeframes, and audience; scored 0–4). PESI ranges from 0 to 7 and serves as an intermediate outcome linking vignette exposure to the quality of AI engagement. The PESI is administered at the follow-up only; it is not part of the baseline or treatment session.

[6]The one exception is a marginal negative coefficient on T2 profits in the log specification ($p = 0.051$), which does not appear in the levels results ($p = 0.916$) and is consistent with noise in a small follow-up sample ($N = 231$).

5.3 Sample Characteristics and Balance

Table 1 presents summary statistics for the baseline sample. Table 2 reports balance on baseline characteristics across treatment arms. No covariate differs significantly at or below the 5% level; the largest unadjusted imbalance is in baseline profit expectations (T1 vs. control, $p = 0.054$, unadjusted), which reflects sampling variation rather than a failure of randomization. Because treatment is assigned after baseline measurement, treatment indicators cannot causally affect pre-treatment outcomes; Table A8 in Appendix A confirms this directly: regressing pre-treatment beliefs on treatment indicators yields near-zero coefficients for T2 and for sales expectations, while the T1 profit coefficient (-0.208, $p = 0.054$) mirrors the balance table finding and has the opposite sign to the belief-updating result, making omitted-variable bias an unlikely explanation for our main findings.

[Insert Table 2 here.]

6 Results

6.1 Information Successfully Updates Beliefs

Table 3 reports treatment effects on profit expectations.[7] The dependent variable is the profit expectation measure described in Section 5. The full-sample baseline profit expectation averages 5.77 (control group: 5.82), meaning the average entrepreneur already expected a moderate profit increase, leaving limited room for upward revision, particularly for already-optimistic entrepreneurs. The update is accordingly concentrated where priors had room to move: among the pessimistic minority, treatment raises profit expectations by roughly half a scale point (Treatment $\times$ Low Optimism = 0.549, $p = 0.002$; Column 4 below), so the pooled average—diluted by optimists already near the top of the scale—understates the revision among those whose beliefs could move. In the main specification with strata fixed effects only (Column 1), the Business Consultant vignette (T1) increases profit expectations by 0.164 scale points ($p = 0.020$), while the Creative Assistant vignette (T2) produces a positive but individually insignificant effect (0.096 scale points, $p = 0.146$). Adding ANCOVA controls (Column 2) yields similar magnitudes (0.158 scale points, $p = 0.027$). The pooled treatment effect is 0.128 scale points ($p = 0.030$, Column 3). The two treatment

[7] Sample sizes vary by outcome and specification. The baseline experiment has $N = 736$; belief regressions with strata fixed effects report $N = 729$ after dropping fixed-effect singletons, and WTP regressions report $N = 703$ after 33 baseline-WTP non-responses. The follow-up barriers item is answered by 308 respondents. The four-month panel has 240 respondents, of whom 231 enter regressions with non-missing baseline covariates.

arms are not statistically distinguishable from each other ($p > 0.10$). Within-session belief revision in the control arm is near-zero, confirming the control vignette is uninformative about business applications.

Column 4 introduces heterogeneity by baseline optimism using a discretized version of Equation (2), where Low Optimism is a binary indicator for baseline profit expectations ≤ 4 ("no change" or below on the seven-point scale), replacing the continuous prior μ_i^0.[8] Among pessimistic entrepreneurs, those who expected AI to have no effect or a negative effect on their profits, treatment raises profit expectations by roughly half a scale point. Treatment × Low Optimism = 0.549 ($p = 0.002$, Figure 2), while the direct treatment effect among already-optimistic entrepreneurs is small and insignificant (0.043 scale points, $p = 0.479$). Belief revision is concentrated entirely among entrepreneurs with pessimistic priors, consistent with the textbook Bayesian learning pattern where skeptics revise more because their prior-signal gap is larger. The causal forest underlying Figure 2 returns an average effect of the same sign but smaller and not individually significant (0.090, SE 0.058), reflecting the method's regularization toward zero in a sample of this size rather than a weaker first stage; the OLS pooled estimate (0.128, $p = 0.030$) is the appropriate average-effect benchmark. The implied learning rate is low to moderate. Our preferred (ratio) estimator gives $\alpha \approx 0.25$ (iid bootstrap 95% CI $[0.09, 0.38]$, bounded away from zero); an alternative estimator that does not normalize by the control group's prior persistence gives $\alpha \approx 0.43$. Both are comparable to or somewhat above Cullen et al. (2025)'s estimate for information about competitors' technology adoption ($\alpha = 0.26$) and well below Cullen and Perez-Truglia (2022)'s for personally relevant salary information ($\alpha = 0.69$). This is consistent with our signal concerning the returns to a broadly applicable technology rather than a directly actionable personal financial variable.[9]

Whether the belief shift persists at the follow-up is less clear. The pooled treatment effect on follow-up profit expectations is near zero ($\beta = 0.030$, $p = 0.833$; Table A10, Appendix A): beliefs do not remain detectably elevated relative to control four months later, although the Treatment × Low Optimism interaction at follow-up (0.668, $p = 0.163$) is comparable in magnitude to the within-session estimate (0.549), leaving

[8] Table A6 in Appendix A reports the continuous-prior specification as written in Equation (2), yielding Treatment × Prior = −0.187 ($p = 0.002$) and confirming the same pattern. In Column 4, the continuous baseline belief is additionally included as a covariate alongside the Low Optimism binary indicator to partial out residual variation in priors within each optimism group.

[9] Following Cullen and Perez-Truglia (2022), α is estimated as the ratio $-\hat{\gamma}_3/\hat{\gamma}_2$, where $\hat{\gamma}_3 = -0.187$ is the coefficient on the Treatment × Prior interaction and $\hat{\gamma}_2 = 0.753$ is the main effect of prior belief in Equation (2). Equating the Bayesian structural prediction with the OLS estimates yields this expression without requiring beliefs to be expressed in monetary units (see Section 2.1). The ratio normalizes by the control group's estimated prior persistence; $\hat{\gamma}_2 = 0.753$ likely reflects measurement error or regression-to-the-mean on the bounded 1–7 scale, causing the alternative estimator $\alpha = 1 - (\hat{\gamma}_2 + \hat{\gamma}_3) \approx 0.43$ to diverge. We retain the ratio estimator for comparability with Cullen and Perez-Truglia (2022) and Cullen et al. (2025); its standard error is 0.07 from a nonparametric bootstrap ($B = 2{,}000$, resampling individuals and recomputing $\hat{\gamma}_2$ and $\hat{\gamma}_3$ jointly).

open persistence precisely in the subgroup that updated. Two considerations govern interpretation. First, these estimates come from the attrition-reduced panel ($N = 231$, roughly one-third of the baseline sample) and follow-up beliefs were not a pre-registered outcome, so they are imprecise and subject to composition change. Second, and more important, the paper's central behavioral test does not depend on persistence: WTP is elicited in the same session as the belief update, so the contemporaneous null cannot be explained by decay. Transience of light-touch informational effects is a documented pattern for short interventions with micro-entrepreneurs (Davies et al. 2024), and it is the reason the design anchors its primary test within the treatment session.

[Insert Table 3 and Figure 2 here.]

6.2 Updated Beliefs Do Not Change Within-Session Behavior

If updated beliefs drive adoption, the behavioral measures should respond. They do not.

Willingness to pay for a personalized AI implementation coaching session is unchanged by treatment (Table A1, Appendix A). The pooled treatment effect on WTP level is 1,705 CLP ($p = 0.253$, Column 2); the effect on WTP Increased is 0.029 ($p = 0.162$, Column 4), both economically small and statistically indistinguishable from zero. Neither arm individually produces a significant effect on WTP level (T1: $p = 0.590$; T2: $p = 0.146$, Column 1). Because WTP is measured in the same baseline session as beliefs, immediately after treatment, belief decay cannot account for this null. Despite believing AI is more profitable than they previously thought, entrepreneurs are no more willing to invest in personalized training that would teach them specifically how to apply AI in their own business. An alternative interpretation is that entrepreneurs plan to explore the technology independently rather than invest in formal training immediately. The follow-up adoption null, however, covers actual behavior and is inconsistent with planned self-exploration translating into adoption. The null on WTP is particularly informative because the coaching session is precisely the kind of firm-specific, hands-on support that would help an entrepreneur put AI to work in her own business, yet a favorable belief update about AI's own-business profit potential does not increase demand for it. Figure A1 (Appendix A) shows the full WTP distribution by treatment arm across all five price thresholds; no threshold shows a significant shift, confirming the null is not an artifact of distributional changes that cancel in the mean. Equivalence tests reinforce that this null is informative rather than merely imprecise: on the full baseline sample we can reject treatment effects on WTP level larger than

0.20 SD (two one-sided tests; Table A15, Appendix A), though the noisier binary WTP-Increased margin remains inconclusive at that bound.

This contemporaneous null holds even among the entrepreneurs who revise their beliefs the most. Among low-optimism entrepreneurs, the subgroup whose profit expectations move roughly half a scale point (Table 3, Column 4), treatment leaves the *level* of willingness to pay unchanged (Table A2, Appendix A): the Treatment × Low Optimism interaction is small and insignificant ($p = 0.806$), and within the low-optimism subsample the treatment effect on WTP level is null ($p = 0.599$). Because these comparisons use the within-session WTP measure on the full baseline sample, they are free of follow-up attrition. The one margin on which a response surfaces is the extensive one: among low-optimism entrepreneurs, treatment marginally raises the probability that WTP rises at all relative to baseline (interaction $p = 0.069$; subsample $p = 0.205$). This flicker sits on a low base—a 2.9% control rate—and does not move the WTP level, so even the strongest belief-updaters do not convert a corrected profit belief into economically meaningful demand for implementation support.

How informative is this null? A natural worry is that the experimental belief shift was simply too small to expect a behavioral response. We benchmark it directly. Across entrepreneurs at baseline, willingness to pay rises steeply with profit expectations—a one-point higher profit expectation is associated with 8,880 CLP higher WTP ($p < 0.001$)—so WTP is far from inelastic to beliefs about AI's returns. Applying this observational gradient to the treatment-induced belief shift implies a WTP increase of 0.024 SD on average and 0.103 SD among the pessimists who update most. The within-session equivalence tests reject WTP-level effects larger than 0.088 SD: the implied average effect lies below what the design can resolve, while the implied pessimist effect sits at the edge of it. We therefore rest the behavioral null not on the precision of any single within-session measure but on the convergence of the willingness-to-pay, free-course take-up, and four-month adoption nulls. It rests, too, on a sharp contrast: the steep *observational* belief–WTP gradient disappears under *exogenous* belief variation. Moving beliefs experimentally does not move behavior, even though, across individuals, beliefs and behavior move together.

Course take-up tells the same story. Roughly two months after the vignettes, Experiment 2 invited a random half of participants to a free AI training course (Section 4.2). Experiment 1 treatment assignment does not predict take-up (Table B1, Appendix B): the effect on starting the course is 3.3 percentage points for T1 ($p = 0.245$) and 0.3 percentage points for T2 ($p = 0.911$), on a control-group base of 36.1% (which pools the invited and non-invited halves; since non-invited participants could not access the course, take-up

among the invited is roughly twice that), and course completion is equally unaffected. Expressed interest in AI training, measured immediately post-treatment, is near-universal (control mean = 0.979; Table A9, Appendix A)[10] and likewise unaffected, which rules out low motivation as an explanation. Entrepreneurs overwhelmingly wanted to learn about AI regardless of treatment, yet updated beliefs moved neither that disposition nor the decision to start the course. Together with the same-session WTP null, these results constitute the paper's primary behavioral evidence.

Figure 3 summarizes the belief-behavior disconnect visually. The standardized treatment effect on profit expectations is positive and significant, while no behavioral outcome, whether measured contemporaneously or at follow-up, reaches significance.

[Insert Figure 3 here.]

6.3 Follow-Up Outcomes

Does the disconnect extend beyond the baseline session? We measure adoption and business outcomes approximately four months after treatment, among the $N = 231$ follow-up respondents.

Table 4 presents treatment effects on adoption and business outcomes at the follow-up ($N = 231$; 9 observations with missing baseline covariates are excluded from regressions). Given that 73% of our sample reports at least occasional prior AI use at baseline (roughly half use it at least weekly for business purposes), these outcomes capture both first-time adoption among non-users and intensified use among existing users. We examine five outcomes: AI usage frequency (five-point scale), new AI adopter status, total business practices (McKenzie-Woodruff index; McKenzie and Woodruff 2017), monthly sales, and monthly profits. Neither treatment arm produces a significant effect on any outcome (Columns 1–5). The pooled estimates confirm the pattern: AI usage frequency ($p = 0.525$), new adopter ($p = 0.651$), business practices ($p = 0.882$), sales ($p = 0.577$), and profits ($p = 0.937$) are all null (Columns 6–10).

The entrepreneurs who revise their beliefs most dramatically show no more adoption than anyone else. The Treatment × Low Optimism interaction on AI usage frequency is insignificant ($p = 0.111$), as is the interaction on new adopter status ($p = 0.550$).[11]

[10] The near-universal expressed interest reflects that the course was offered free of charge with a credentialed faculty expert. Variation in actual take-up, not in expressed interest, is the meaningful behavioral measure.

[11] As an exploratory exception to the pre-registered subgroup plan, we also examine heterogeneity by baseline profit level (above vs. below median) and by sex. Both interactions are uniformly null across belief updating, WTP, and follow-up adoption outcomes (all $p > 0.10$). These analyses were not pre-registered and are clearly labeled as exploratory.

The follow-up sample ($N = 231$) limits precision for these medium-term outcomes; the design's power calculations are available in the pre-registration. We therefore read them alongside the within-session results of Section 6.2, which are estimated on the full baseline sample without attrition and constitute the paper's primary behavioral evidence.

[Insert Table 4 here.]

6.4 Exploratory: Prompt Quality and the AI Training Course

Although the informational vignettes produce no adoption effects, the Experiment 2 training course offers suggestive evidence consistent with the framework. Of the 240 follow-up respondents, 21 (8.8%) complete the course and 77 (32.1%) start without completing. Course completers score higher on the PESI prompt-sophistication index (+0.726, $p = 0.008$; Appendix B), a pattern consistent with hands-on engagement building the firm-specific prompting skill that passive information does not, and one concentrated in prompt quality rather than in the adoption outcomes that are our primary hypotheses (Table B4, Appendix B). Because completion is self-selected, we read this as suggestive rather than causal; it is directionally consistent with the training literature, in which sustained, hands-on engagement builds capabilities a single informational session cannot (Camuffo et al. 2020; McKenzie 2021). The intention-to-treat effect of the course *invitation* on prompt sophistication is, if anything, negative—consistent with invitations alone, absent completion, not replicating the gains of engagement—and we report the full course analysis in Appendix B, treating Experiment 2 as exploratory throughout.

6.5 Robustness

The main results hold across a series of checks reported in Appendix A. Table A13 maps each pre-registered hypothesis to its estimated effect and paper section.

We address potential attrition bias through three complementary approaches. The follow-up panel ($N = 240$) represents approximately one-third of the baseline sample. We verify that stayers and attritors are balanced on baseline characteristics (Table A3), and bound worst-case selection using Lee (2009) bounds (Table A4) and inverse probability weighting (Table A5). A falsification test confirms that treatment is orthogonal to pre-treatment beliefs. Regression of baseline profit and sales expectations on treatment assignment yields null coefficients for both arms (Table A8, Appendix A), with a marginally significant T1

coefficient ($p = 0.054$) attributable to sampling variation and addressed by controlling for baseline beliefs in all outcome regressions. Lee bounds on AI usage frequency [−0.027, 0.496] and total business practices [−0.400, 1.130] both include zero, confirming that the null adoption results are not artifacts of selective attrition. For the new adopter indicator, both bounds are negative [−0.123, −0.019], indicating that the treatment effect on new adoption is non-positive, and potentially modestly negative, even under worst-case attrition trimming. We interpret this cautiously given the small cell sizes involved in the bounding procedure, and note that the inverse-probability-weighted estimate for the same outcome is null (Table A5).

Follow-up response rates are similar across treatment arms (Control = 29.8%, T1 = 33.5%, T2 = 34.4%), and the differential is not statistically significant (see Table A3). These are follow-up response rates, representing the share of the 736 baseline participants who completed the follow-up survey, and should not be confused with baseline attrition, which is zero by design (the experiment runs inside the baseline survey).

An IV strategy instrumenting post-treatment beliefs with treatment assignment is infeasible on the follow-up panel, as the first-stage F-statistic falls below 10, reflecting the modest belief-updating coefficient ($\alpha \approx 0.25$) combined with reduced sample precision from attrition. The paper therefore relies on the ITT estimand throughout.[12]

We probe the robustness of the belief-updating result along two dimensions. Table A6 reports the continuous-prior Bayesian learning specification from Equation (2), yielding Treatment × Prior = −0.187 ($p = 0.002$) and confirming the heterogeneity pattern with additional statistical power. Table A7 extends the analysis to sales expectations and an inverse-variance-weighted belief index (Anderson 2008). The direction of treatment effects is consistent across all three belief measures, though statistical significance is concentrated in profit expectations ($p = 0.030$), reflecting the analytical framing of the vignettes.

To account for multiple testing, we apply Benjamini-Hochberg corrections (Benjamini and Hochberg 1995) within the family of four pre-registered *primary* outcomes (hypotheses H1–H4 in Table A13): beliefs, AI usage frequency, new-adopter status, and business practices. The correction preserves the belief result ($p_{\text{BH}} < 0.05$) and leaves the adoption null intact. Willingness to pay is a separately pre-registered hypothesis (H5–H6) and is therefore not part of this family; its within-session null is already insignificant before any correction.

[12] Beyond the weak first stage, the instrumental-variables logic adds little here. The compliers whose beliefs respond to treatment are the pessimistic updaters (Table 3, Column 4), and they are precisely the entrepreneurs for whom we observe no differential change in adoption (Table 4) or willingness to pay (Table A2). A reduced-form null concentrated among the strongest updaters already implies the absence of a belief-driven behavioral response; rescaling it by a weak first stage would only widen the confidence interval, not uncover a hidden effect.

The null adoption result also holds across subgroups. Table A11 reports adoption effects interacted with prior AI use. The Treatment × Prior AI User interaction is insignificant for all outcomes ($p > 0.16$), indicating that the belief-behavior disconnect is not concentrated among prior users (intensive margin) or non-users (extensive margin). Table A12 tests whether the treatment arm producing the strongest belief revision (T1) generates differential adoption among low-optimism entrepreneurs; the interaction is positive but statistically indistinguishable from zero ($p = 0.218$ for AI frequency), consistent with the main finding.

7 Discussion

7.1 Why Do Updated Beliefs Not Produce Adoption?

Our central finding is a disconnect that the standard information provision framework does not predict. Information that successfully updates entrepreneurs' beliefs about AI's profitability produces no statistically detectable change in within-session behavioral outcomes; medium-term follow-up outcomes are estimated with limited precision due to attrition. The information was not the problem: generative AI can raise productivity where it is put to effective use, even as its gains are uneven across tasks (Brynjolfsson et al. 2025; Noy and Zhang 2023; Dell'Acqua et al. 2026). To interpret the disconnect, we begin from what the entrepreneurs themselves report holds them back. The evidence points first to the costs of acting—money, time, and the effort to learn the tool—rather than to any failure to grasp AI's value. The general-purpose nature of AI is what makes a corrected profit belief insufficient to relax those costs: knowing that AI pays on average does not fund a paid tier, free up time, or do the firm-specific work of learning to apply it. Two conditions a corrected belief leaves unmet organize what follows (Figure 4). The first is the costs of acting, which the entrepreneurs themselves report; the second is the firm-specific application problem, which of one's own tasks AI improves and how. Our design, which varies beliefs alone, does not isolate them individually.

Adoption costs are the constraint the entrepreneurs themselves point to. Asked in the follow-up to name their main barriers to adopting AI (up to three; Figure 5), they named material and practical constraints far more often than cognitive ones: cost led at 39%, followed by data privacy (19%) and the time to learn the tool (17%), with technical complexity (12%) behind them. The barriers that would implicate AI's value or the entrepreneur's grasp of it were named least: only 10% did not see how AI could help their own business, 10% preferred to work without it, and 9% distrusted the quality of its output—a form of algorithm aversion

(Liu et al. 2026; Bockstedt and Buckman 2026). Cost is also the hub of the responses, co-occurring most often with privacy and time, so the typical respondent describes a bundle of resource constraints rather than doubt about the technology. These are constraints on acting, not on understanding, and a corrected profit belief leaves each untouched.[13]

[Insert Figure 5 here.]

Two features of these barriers deserve emphasis. First, they are largely *perceived* rather than literal. Conversational AI is free and runs in a browser, so "no money to pay for AI tools" reflects an anticipated cost of useful adoption (paid tiers, or, more plausibly, the time and effort to learn to prompt well for one's own business) rather than a price of access; the relevant cost is noncapital and learning-related, the kind long found to dominate technology adoption decisions (Åstebro 2002, 2004). Second, privacy concerns about business data are real and shape technology use independently of beliefs about returns (Acquisti et al. 2015). None of these is something a message about AI's average profitability can dislodge. This ordering matches what the small-firm literature would predict: micro-enterprises in developing economies are held back by limited access to capital (Bruhn et al. 2010) and shortfalls of managerial time (McKenzie and Woodruff 2017), and scarcity itself taxes the attention that experimentation requires (Mani et al. 2013). The course take-up pattern makes the time constraint concrete, with only 8.8% of follow-up respondents completing the free AI course even when invited. Since most of our sample has tried AI but few use it consistently, the binding constraint appears to be not awareness of the tool but the cost of using it well. Humlum and Vestergaard (2025) document a parallel pattern at the macro level, where AI spreads quickly yet measurable gains emerge only where firms invest in training and integration. Fragmented markets compound the problem. Without peer adoption to lower uncertainty, each micro-entrepreneur bears the full cost of early adoption alone (Dimitriadis and Koning 2025; Fafchamps and Quinn 2018).

Could entrepreneurs simply fail to see where AI fits their own business? Only 10% name lack of perceived relevance as a main barrier, so task identification appears secondary in our data, but self-reports may understate discovery frictions. Learning that AI is profitable on average and identifying which tasks in one's own business benefit are different operations (Aragonès et al. 2005; Agrawal et al. 2021), and the boundary between tasks AI helps and tasks it harms is irregular even for elite professionals (Dell'Acqua et

[13] Respondents could select up to three barriers; 33% chose an uninterpretable "other" option (which carried no free-text follow-up) and 5% named none, so shares are computed over the 308 follow-up respondents who answered the item. The low uptake of the interpretable cognitive barriers is itself informative and is robust to how "other" is treated.

al. 2026). Hanna et al. (2014) show that agents can fail to learn not from lack of data but from lack of attention to the relevant dimensions, and Chatterji et al. (2025) find that even frequent ChatGPT users default to general advice-seeking rather than embedding the tool in specific task workflows. Information about average profitability does not perform this firm-specific search, and neither does a worked demonstration of AI applied to another firm's tasks, as in our vignettes. In the Bayesian-entrepreneurship account, that search is the entrepreneur's own to do; a vicarious signal cannot substitute for it (Kim 2026).

Willingness to pay for personalized coaching, a measure of demand for implementation support, is consistent with this reading: a positive shift in expected value should, under standard expected utility, raise the price an entrepreneur will pay for the very help designed to realize that value, yet it does not. Because WTP is elicited in the same session as the belief shift, belief decay cannot explain the null. Nor can liquidity, the obvious worry when the priced object is coaching: WTP is elicited both before and immediately after the vignette, so the treatment effect is identified from within-person change over an interval of minutes, across which an entrepreneur's budget is fixed and therefore differenced out. The share of entrepreneurs whose own WTP rises after treatment is flat (0.029, $p = 0.162$; Table A1), as is WTP among above-median-profit entrepreneurs, who are least constrained (Section 6.3); and take-up of the free course, which carries no price at all, is equally unmoved. The contemporaneous WTP null (robust to belief decay) and the free take-up null together close off the two leading alternatives: behavior does not move whether or not money is at stake or time has elapsed. We do not, however, read the flat WTP as decisive about mechanism. Coaching is precisely the service that would lower the learning cost we emphasize, so its flatness is striking. But coaching carries its own time cost, and a within-session belief shift of 0.128 points is small relative to the price increments at stake (5,000 to 115,000 CLP). Indifference is therefore consistent with two readings: the constraint lies downstream of beliefs, or the nudge is too small to clear the price of acting. Our design does not adjudicate between these readings.

The heterogeneity across the belief distribution points the same way. Entrepreneurs with pessimistic priors revise most—moving roughly half a scale point (Treatment × Low Optimism = 0.549, $p = 0.002$)—yet even these strong updaters show no differential adoption (Treatment × Low Optimism on AI frequency: $p = 0.111$; on new adoption: $p = 0.550$). Were the aggregate profit belief sufficient for action, the largest revisions should produce the largest responses; they do not. Our elicitation is telling. We ask how profits would change *if the entrepreneur used AI regularly and effectively*, so the item moves the return expected *conditional on* effective use, the aggregate belief. It does not move the application-specific belief, which of

her own tasks AI improves and how, that governs whether she acts; raising the former need not move the latter. That this null holds even among entrepreneurs who, as unincorporated sole proprietors, are if anything more responsive to new information than employees (Åstebro et al. 2026) only deepens the puzzle.

For the follow-up outcomes, simple belief transience is a further possibility—brief, passive vignettes may produce effects that fade (Davies et al. 2024; Jensen 2007)—but it cannot explain the contemporaneous nulls on willingness to pay and course take-up, measured in the same session as treatment. The same gap separates initial adoption from productive use (Comin and Hobijn 2010): the task-specific mental models that turn adoption into performance are built by hands-on testing in one's own business, not delivered by passive information (Camuffo et al. 2020; Koning et al. 2022; Gambardella and Messinese 2025). Consistent with this, sustained engagement among course completers produces higher prompt quality (Section 6.4 and Appendix B).

[Insert Figure 4 here.]

7.2 Contributions and Implications

Three contributions follow. The first is to the view of entrepreneurship as disciplined belief revision under uncertainty (Agrawal et al. 2026; Camuffo et al. 2020; Koning et al. 2022; Gambardella and Messinese 2025; Agrawal et al. 2021), in which entrepreneurs act on beliefs they form and test. We isolate the simple-updating channel: most prior work infers updating from a behavioral non-response, whereas we *measure* the update and show that exogenously correcting a generic profitability belief does not move action, even among the strongest updaters. Correcting that belief is therefore not sufficient on its own. The binding step lies in what such a signal cannot reach: the application-specific belief built through firm-specific learning, and the costs of acting it does not relax. It does not lie in the aggregate belief itself, and our design isolates that belief channel without separating the conditions downstream of it. Second, and more broadly, the result marks where information provision reaches its limit (Jensen 2010; Cullen and Perez-Truglia 2022; Coibion et al. 2020; Haaland et al. 2023): when the belief that governs action is application-specific—as it is for a general-purpose technology whose uses must be discovered—a credible update that moves the aggregate belief need not move behavior, in the same session or months later. Third, contributing to the literature on technology adoption and management practices in small firms (Foster and Rosenzweig 1995; Conley and Udry 2010; Åstebro 2002; McKenzie and Woodruff 2017; Bloom et al. 2013), the experiment isolates the

belief channel while holding information quality, tool access, and implementation support constant; training and advice interventions here often fail to change firm behavior (McKenzie 2021) and small firms leave even valuable, free information unused (Kim 2025), and our design separates failure to persuade from persuasion that is not enough, ruling out the first. The constraint lies not in what entrepreneurs know but in what they can do with that knowledge.

Three concurrent studies test related channels from distinct experimental angles. Cullen et al. (2025) provide Italian firms with information about competitors' technology adoption and find that investment in robotics responds while AI adoption does not. Menkhoff (2025), in a survey experiment with established German firms, finds that productivity information raises beliefs and that self-reported adoption responds mainly among respondents with decision authority; our micro-entrepreneurs are sole decision-makers, so the absence of a response here cannot be attributed to authority frictions. Otis et al. (2024) randomize access to an AI business assistant among Kenyan micro-entrepreneurs and find that gains concentrate among those with the capacity to screen and implement the advice. Across these designs no single lever moves adoption on its own; ours isolates the profit-belief channel itself, measured among sole decision-makers facing no authority or access friction. In this study, the first-stage result is not in dispute: the vignettes do exactly what information provision is supposed to do, yet adoption is unmoved even within the same survey session, making the evidence particularly difficult to reconcile with beliefs as the key adoption barrier.

For program designers and technology managers, the diagnostic question is which constraint binds. Where pessimistic beliefs are the barrier, information campaigns are the right tool; in our setting, beliefs respond but behavior does not, so the implication is to move resources from communicating AI's value toward implementation support that addresses the constraints entrepreneurs actually report: the monetary cost of tools, the time cost of learning them, and assurance about data privacy, alongside hands-on guidance that builds business-specific applications. Sustained engagement across multiple sessions, and structural support that lowers the time cost of early experimentation, are likely necessary components of effective capability-building interventions. The SME management practices literature, from intensive management consulting (Bloom et al. 2013; Bruhn et al. 2018) to entrepreneurship training programs (Anderson et al. 2018; McKenzie 2021), has established that changing firm behavior in developing economies requires more than information. Evidence from accelerator programs and training interventions finds that combining basic services with entrepreneurship education generates larger venture improvements than services alone (González-Uribe and Leatherbee 2018; Campos et al. 2025), consistent with the view that entrepreneurial

capital, similar to managerial capital, is a type of capital that is missing among certain populations (Bruhn et al. 2010). The non-adoption of good managerial practices documented in Bloom et al. (2013) is itself partly attributable to informational frictions, as firms are often unaware of better practices or do not believe those practices will improve their performance (González-Uribe and Leatherbee 2018). Our experiment targets exactly this friction for AI, providing credible evidence that AI improves performance, yet resolving this informational barrier is not sufficient. Our results add that the barriers that matter for AI (its monetary and time costs, privacy concerns, and the firm-specific capabilities needed to put it to work) remain in place even after beliefs are corrected. This mirrors findings from the broader technology adoption literature, where exporting firms that receive direct production support, not just market access, are the ones that improve quality and productivity (Atkin et al. 2017). For micro-enterprises worldwide (Sotiriou et al. 2025), AI represents a productivity opportunity that will not be unlocked by showing its value alone; it requires interventions that build the firm-specific capabilities needed to realize that value.

7.3 Limitations

Four limitations bound the interpretation of our findings. First, the design establishes that exogenous variation in beliefs does not produce adoption, but it does not directly measure the two candidate constraints (the costs of acting and the firm-specific application problem) that we advance to interpret it. Future work with designs that independently manipulate each condition would allow causal decomposition of which constraint dominates and in which populations. Second, the reduced-form ITT estimates do not trace the causal pathway from belief revision to adoption; they establish only that the pathway is not operative at the experimental margins studied. Designs using instrumental variables or sequential randomization can decompose the channels more precisely. Third, follow-up attrition (approximately one-third of the baseline sample was re-contacted) limits the precision of medium-term estimates relative to comparable information-provision experiments (Jensen 2010; Cullen and Perez-Truglia 2022; Cullen et al. 2025). This caveat applies only to the follow-up outcomes. The within-session WTP measure and course take-up, which is observed from platform records rather than survey responses, are estimated on the full baseline sample without attrition and form the primary evidentiary basis for the belief-action gap. The experiment characterizes adoption dynamics at an early stage of AI diffusion in a digitally engaged Latin American sample; whether the belief channel binds differently in lower-access populations or at later stages of diffusion remains an open question for future research. Fourth, the design cannot rule out the option value of waiting as an alternative explanation for the

follow-up adoption null. When a technology evolves rapidly, rational agents may delay adoption to avoid committing to a version that will quickly be superseded (Jovanovic and Lach 1989). This explanation cannot account for the contemporaneous null on WTP, measured within the same session as treatment before any delay could occur, but it remains a candidate for the absence of adoption gains at the follow-up horizon.

8 Conclusion

Informational vignettes shift beliefs about AI's profit impact in a pattern consistent with Bayesian updating, yet produce no statistically detectable change in adoption across the pre-registered behavioral measures. This disconnect between updated beliefs and unchanged behavior is the paper's central finding, and it marks a limit of information provision: where adoption turns on more than a favorable assessment of value, correcting the belief is not enough. What entrepreneurs report standing in their way is not doubt about AI but the cost of acting on it—money, the time to learn to use it well, and concern for the privacy of their data. Because AI is general-purpose, information about its average profitability leaves each of these exactly where it was: the price of the tool, the hours to learn it, and the discovery of which tasks to hand over are untouched by knowing that AI pays on average. Information settles the first-order question of whether AI is worth pursuing and leaves the second-order problem of how to act on it intact.

The implication for policy and practice is that information campaigns are not sufficient for AI adoption. This shifts the diagnostic question—from how to communicate AI's value more effectively to what keeps adoption low even among entrepreneurs who already believe in it. Small firms worldwide have yet to realize AI's productivity gains, and our evidence suggests this owes less to a want of awareness than to the resources and firm-specific know-how needed to put the tool to work. Closing the gap requires interventions that go beyond telling entrepreneurs that AI works—lowering the cost and effort of adoption, and helping them identify and implement the uses that fit their own business.

Several directions follow. First, the two candidate constraints (the costs of acting and the firm-specific application problem) are theorized but not directly measured here; designs that manipulate each independently—subsidized experimentation time, or business-specific AI task maps—would identify which dominates, and where. Second, the belief channel may bind differently at earlier stages of diffusion, among lower-access populations without prior AI exposure, or where peer learning is stronger; cross-setting replication would establish whether the null generalizes beyond an already-exposed, digitally connected sample.

Third, the course-completer evidence shows that sustained, hands-on engagement builds the prompting skill that passive information cannot. That points to interventions targeting implementation capacity rather than the first-order belief alone.

References

Agrawal A, Gans JS, Stern S (2021) Enabling entrepreneurial choice. *Management Sci.* 67(9):5510–5524.

Agrawal A, Camuffo A, Gambardella A, Gans J, Scott EL, Stern S, eds (2026) *Bayesian Entrepreneurship* (MIT Press, Cambridge, MA).

Acquisti A, Brandimarte L, Loewenstein G (2015) Privacy and human behavior in the age of information. *Science* 347(6221):509–514.

Anderson ML (2008) Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *J. Amer. Statist. Assoc.* 103(484):1481–1495.

Anderson SJ, Chandy R, Zia B (2018) Pathways to profits: The impact of marketing vs. finance skills on business performance. *Management Sci.* 64(12):5559–5583.

Aragonès E, Gilboa I, Postlewaite A, Schmeidler D (2005) Fact-free learning. *Amer. Econom. Rev.* 95(5):1355–1368.

Åstebro T (2002) Noncapital investment costs and the adoption of CAD and CNC in U.S. metalworking industries. *RAND J. Econom.* 33(4):672–688.

Åstebro T (2004) Sunk costs and the depth and probability of technology adoption. *J. Indust. Econom.* 52(3):381–399.

Åstebro T, Fossen FM, Gutierrez C (2026) Attitudes toward ambiguity among self-employed and incorporated entrepreneurs. *J. Bus. Venturing* 41(4):106601.

Åstebro T, Herz H, Nanda R, Weber RA (2014) Seeking the roots of entrepreneurship: Insights from behavioral economics. *J. Econom. Perspect.* 28(3):49–70.

Athey S, Wager S (2019) Estimating treatment effects with causal forests: An application. *Observ. Stud.* 5(2):37–51.

Atkin D, Khandelwal AK, Osman A (2017) Exporting and firm performance: Evidence from a randomized experiment. *Quart. J. Econom.* 132(2):551–615.

Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* 57(1):289–300.

Bloom N, Eifert B, Mahajan A, McKenzie D, Roberts J (2013) Does management matter? Evidence from India. *Quart. J. Econom.* 128(1):1–51.

Bockstedt JC, Buckman JR (2026) Humans' use of AI assistance: The effect of loss aversion on willingness to delegate decisions. *Management Sci.* 72(1):323–342.

Bresnahan T (2024) What innovation paths for AI to become a GPT? *J. Econom. Management Strategy* 33(2):305–316.

Bresnahan TF, Trajtenberg M (1995) General purpose technologies: "Engines of growth"? *J. Econometrics* 65(1):83–108.

Bruhn M, Karlan D, Schoar A (2010) What capital is missing in developing countries? *Amer. Econom. Rev.* 100(2):629–633.

Bruhn M, Karlan D, Schoar A (2018) The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in Mexico. *J. Political Econom.* 126(2):635–687.

Brynjolfsson E, Li D, Raymond LR (2025) Generative AI at work. *Quart. J. Econom.* 140(2):889–942.

Campos F, Frese M, Iacovone L, Johnson HC, McKenzie D, Mensmann M (2025) Long-term and lasting impacts of personal initiative training on entrepreneurial success. *Amer. Econom. Rev. Insights* 7(4):463–479.

Camuffo A, Cordova A, Gambardella A, Spina C (2020) A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Sci.* 66(2):564–586.

Chatterji A, Cunningham T, Deming DJ, Hitzig Z, Ong C, Shan C, Wadman K (2025) How people use ChatGPT. NBER Working Paper 34255, National Bureau of Economic Research, Cambridge, MA.

Cohen WM, Levinthal DA (1990) Absorptive capacity: A new perspective on learning and innovation. *Admin. Sci. Quart.* 35(1):128–152.

Coibion O, Gorodnichenko Y, Ropele T (2020) Inflation expectations and firm decisions: New causal evidence. *Quart. J. Econom.* 135(1):165–219.

Comin D, Hobijn B (2010) An exploration of technology diffusion. *Amer. Econom. Rev.* 100(5):2031–2059.

Conley TG, Udry CR (2010) Learning about a new technology: Pineapple in Ghana. *Amer. Econom. Rev.* 100(1):35–69.

Cullen Z, Perez-Truglia R (2022) How much does your boss make? The effects of salary comparisons. *J. Political Econom.* 130(3):766–822.

Cullen Z, Faia E, Guglielminetti E, Perez-Truglia R, Rondinelli C (2025) The innovation race: Experimental evidence on advanced technologies. NBER Working Paper 34532, National Bureau of Economic Research, Cambridge, MA.

Davies E, Deffebach P, Iacovone L, McKenzie D (2024) Training microentrepreneurs over Zoom: Experimental evidence from Mexico. *J. Development Econom.* 167:103244.

Dell'Acqua F, McFowland E III, Mollick E, Lifshitz H, Kellogg K, Rajendran S, Krayer L, Candelon F, Lakhani K (2026) Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Organ. Sci.* 37(2):403–423.

Dimitriadis S, Koning R (2025) Networking frictions and entrepreneurial learning in developing economies. *Management Sci.* 71(4):3127–3145.

Dushnitsky G, Klueter T (2016) Linking technologies to applications. Folta TB, Helfat CE, Karim S, eds. *Advances in Strategic Management*, vol. 35 (Emerald Publishing, Bingley, UK), 285–317.

Fafchamps M, Quinn S (2018) Networks and manufacturing firms in Africa: Results from a randomized field experiment. *World Bank Econom. Rev.* 32(3):656–675.

Fernández V, Uriarte S (2026) IA en microemprendimientos chilenos: oportunidades, brechas y desafíos éticos. Report, Doble Click Económico, UAI Escuela de Negocios/ICARE, Santiago, Chile.

Foster AD, Rosenzweig MR (1995) Learning by doing and learning from others: Human capital and technical change in agriculture. *J. Political Econom.* 103(6):1176–1209.

Gambardella A, Messinese D (2025) Design- and theory-based approaches to strategic decisions. *Organ. Sci.* 36(4):1271–1287.

Gambardella A, Gius L, Stern S (2026) Homo Entrepreneuricus. Agrawal A, Camuffo A, Gambardella A, Gans J, Scott EL, Stern S, eds. *Bayesian Entrepreneurship* (MIT Press, Cambridge, MA).

Gans J (2026) Knightian uncertainty and Bayesian entrepreneurship. Agrawal A, Camuffo A, Gambardella A, Gans J, Scott EL, Stern S, eds. *Bayesian Entrepreneurship* (MIT Press, Cambridge, MA).

Gertler P, Higgins S, Malmendier U, Ojeda W (2025) Do behavioral frictions prevent firms from adopting profitable opportunities? Working paper, National Bureau of Economic Research, Cambridge, MA.

Goldfarb A, Taska B, Teodoridis F (2023) Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings. *Res. Policy* 52(1):104653.

González-Uribe J, Leatherbee M (2018) The effects of business accelerators on venture performance: Evidence from Start-Up Chile. *Rev. Financial Stud.* 31(4):1566–1603.

Grégoire DA, Shepherd DA (2012) Technology-market combinations and the identification of entrepreneurial opportunities. *Acad. Management J.* 55(4):753–785.

Haaland I, Roth C, Wohlfart J (2023) Designing information provision experiments. *J. Econom. Literature* 61(1):3–40.

Hanna R, Mullainathan S, Schwartzstein J (2014) Learning through noticing: Theory and evidence from a field experiment. *Quart. J. Econom.* 129(3):1311–1353.

Humlum A, Vestergaard E (2025) Still waters, rapid currents: AI adoption and work reorganization. NBER Working Paper 33777, National Bureau of Economic Research, Cambridge, MA.

Jensen R (2007) The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *Quart. J. Econom.* 122(3):879–924.

Jensen R (2010) The (perceived) returns to education and the demand for schooling. *Quart. J. Econom.* 125(2):515–548.

Jovanovic B, Lach S (1989) Entry, exit, and diffusion with learning by doing. *Amer. Econom. Rev.* 79(4):690–699.

Kim H (2025) The value of competitor information: Evidence from a field experiment. *Management Sci.* 71(4):3600–3621.

Kim H (2026) Learning from data in entrepreneurial experimentation. Agrawal A, Camuffo A, Gambardella A, Gans J, Scott EL, Stern S, eds. *Bayesian Entrepreneurship* (MIT Press, Cambridge, MA).

Koning R, Hasan S, Chatterji A (2022) Experimentation and start-up performance: Evidence from A/B testing. *Management Sci.* 68(9):6434–6453.

Lee DS (2009) Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Rev. Econom. Stud.* 76(3):1071–1102.

Liu M, Tang X, Xia S, Zhang S, Zhu Y, Meng Q (2026) Algorithm aversion: Evidence from ridesharing drivers. *Management Sci.* 72(1):193–203.

Mani A, Mullainathan S, Shafir E, Zhao J (2013) Poverty impedes cognitive function. *Science* 341(6149):976–980.

Martínez A. C, Puentes E, Ruiz-Tagle J (2018) The effects of micro-entrepreneurship programs on labor market performance: Experimental evidence from Chile. *Amer. Econom. J. Appl. Econom.* 10(2):101–124.

McKenzie D (2021) Small business training to improve management practices in developing countries: Re-assessing the evidence for 'training doesn't work'. *Oxford Rev. Econom. Policy* 37(2):276–301.

McKenzie D, Woodruff C (2017) Business practices in small firms in developing countries. *Management Sci.* 63(9):2967–2981.

Menkhoff M (2025) Belief updating and AI adoption: Experimental evidence from firms. CESifo Working Paper 12291, CESifo, Munich, Germany.

Miller A, O'Mahony S, Cohen SL (2024) Opening the aperture: Explaining the complementary roles of advice and testing when forming entrepreneurial strategy. *Organ. Sci.* 35(1):1–26.

Noy S, Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654):187–192.

OECD, BCG, INSEAD (2025) The adoption of artificial intelligence in firms: New evidence for policymaking. Report, OECD Publishing, Paris.

Otis NG, Clarke R, Delecourt S, Holtz D, Koning R (2024) The uneven impact of generative AI on entrepreneurial performance. Working paper, SSRN, Rochester, NY.

Ronconi JP (2025) Planning and performance: Experimental evidence from small firms in e-commerce. Working paper, SSRN, Rochester, NY.

Sotiriou A, Sawhney S, Kruijff D, Kimani SW (2025) Big futures for the smallest firms: Finance that works for nano, micro and small enterprises. Working paper, CGAP, Washington, DC.

Suri T (2011) Selection and comparative advantage in technology adoption. *Econometrica* 79(1):159–209.

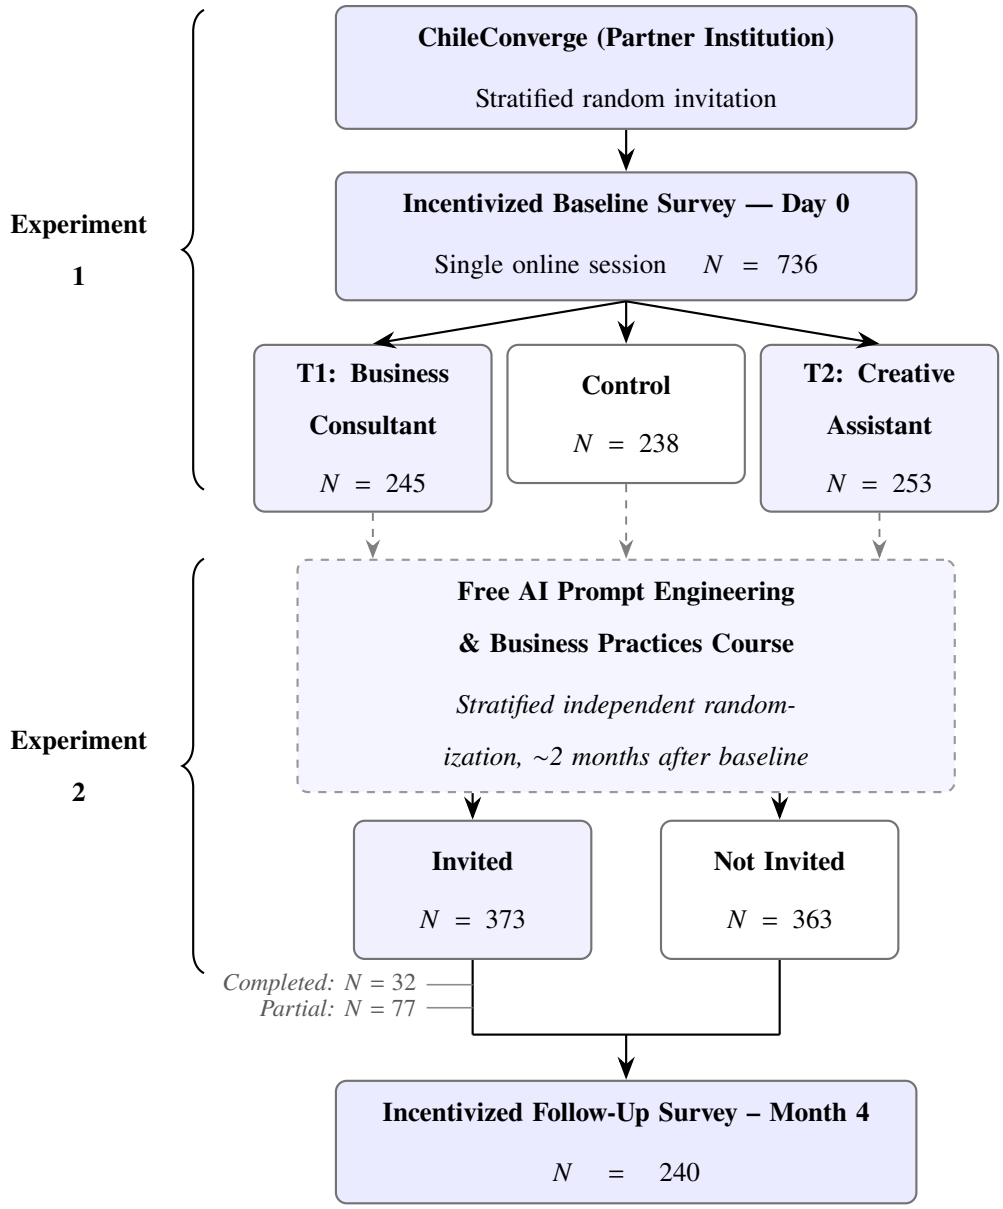


Figure 1: Experimental Design. Within the baseline survey (N = 736, Day 0), participants were randomly assigned within strata to one of three vignette conditions (T1, T2, Control). Approximately two months later, a stratified 50/50 randomization, independent of vignette assignment, invited half to a free AI training course (N_{invited} = 373); the platform concealed the course entirely from non-invited participants, ensuring zero contamination. The follow-up survey, fielded approximately four months after the baseline (N = 240), measures AI adoption, business practices, sales, and profits; of the 373 course-invited, 32 completed the full course, 21 of whom are in this sample.

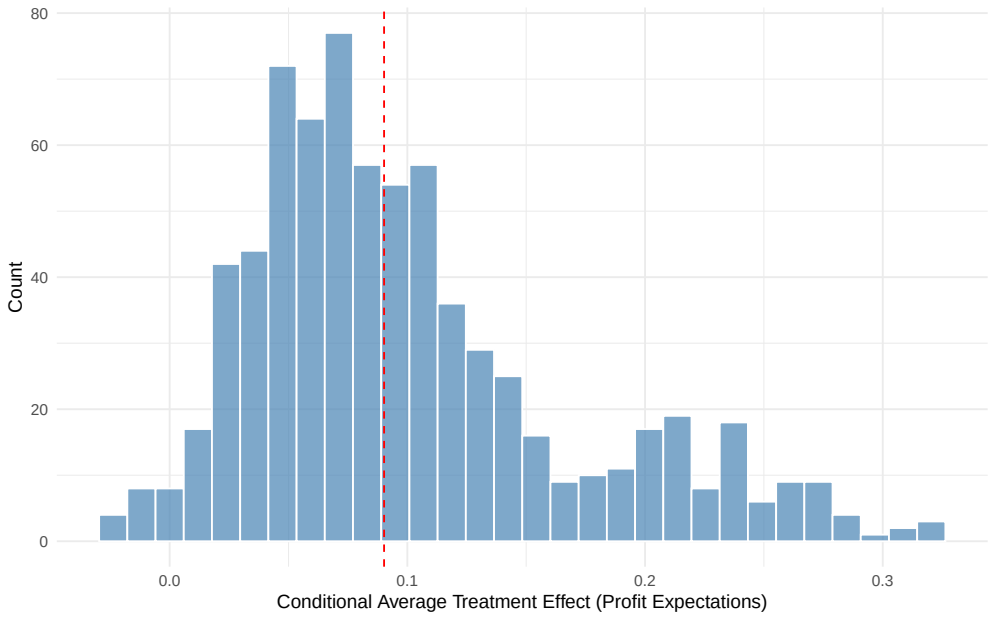


Figure 2: Distribution of Conditional Average Treatment Effects on Profit Expectations ($N = 736$). Individual-level CATEs from a causal forest (Athey and Wager 2019) on post-treatment profit expectations (pooled treatment indicator); the red dashed line marks the forest ATE (0.090, SE = 0.058). All CATEs are positive, and the right-skewed distribution (most entrepreneurs revise modestly, a tail substantially) is consistent with the Bayesian pattern that lower priors produce larger updates (Table 3). The forest ATE differs from the OLS estimate (0.128) due to regularization.

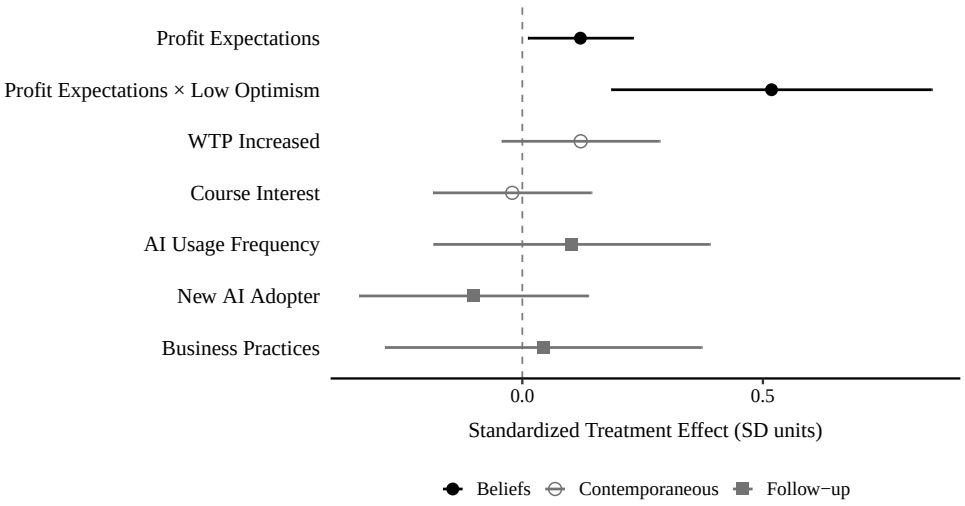


Figure 3: Standardized Treatment Effects on Beliefs and Behavioral Outcomes. Each point represents the pooled treatment effect (Any Treatment vs. Control) standardized by the control-group standard deviation. Horizontal bars show 95% confidence intervals. Belief outcomes (blue, $N = 729$) and contemporaneous behavioral outcomes (open circles, $N = 703$–729) are measured in the baseline session; follow-up behavioral outcomes (filled circles, $N = 231$) are measured at the follow-up survey, approximately four months after the baseline. The belief effect on profit expectations is statistically significant ($p = 0.030$); no behavioral outcome reaches significance at conventional levels.

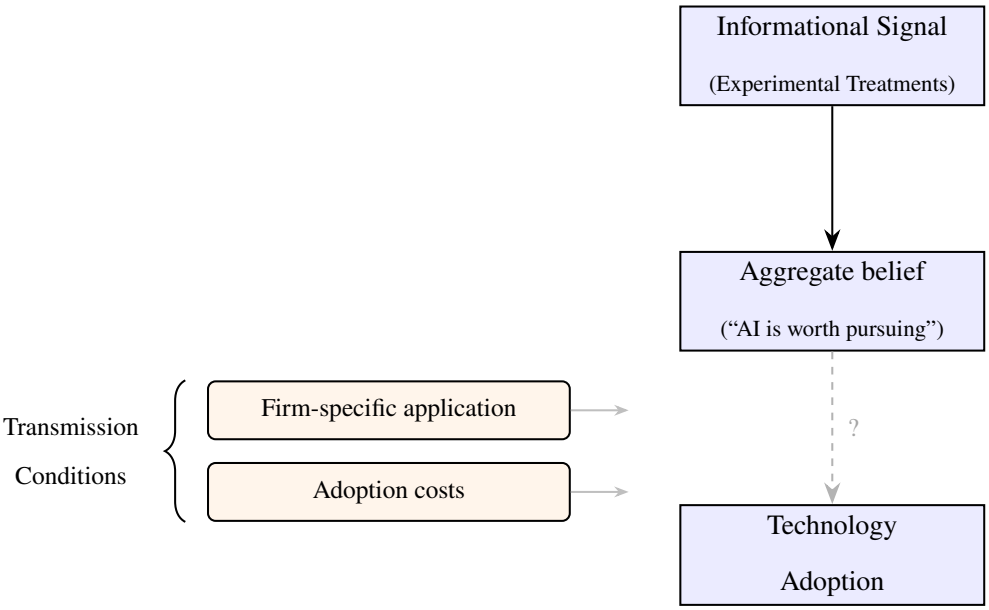


Figure 4: The Conditional Transmission Framework. An informational signal can move the *aggregate* belief that a technology is worth pursuing, but adoption turns on the *operative* belief—which of the entrepreneur's own tasks it improves, and how—which the signal does not supply. Two transmission conditions stand between belief and action: the firm-specific application problem and the costs of acting. Each is a distinct constraint that correcting the aggregate belief alone cannot resolve. Arrows represent conceptual relationships, not estimated causal paths.

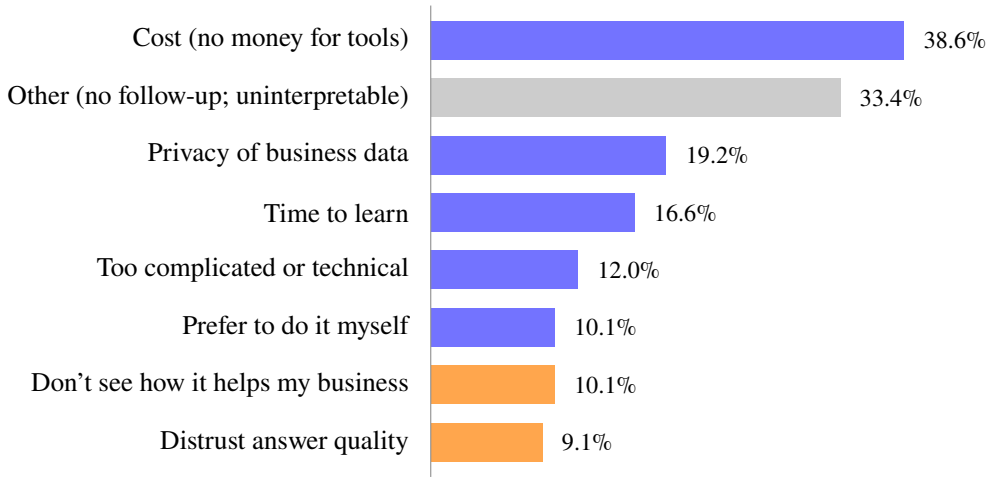


Figure 5: Self-reported barriers to adopting AI (N = 308 follow-up respondents). Respondents could select up to three of eight options (mean 1.5 selections), so shares sum to more than 100%. Material and practical barriers (blue)—cost, data privacy, time, and complexity—dominate, while the cognitive barriers that would implicate AI's value or the entrepreneur's grasp of it (orange)—not seeing how AI helps one's own business, and distrust of output quality—are selected least. "Other" (grey) carried no free-text follow-up and is not interpretable; a further 5% of respondents selected no barrier. Cost co-occurs most often with privacy and time (Section 7).

Table 1: Summary Statistics by Treatment Group

	Control (N = 238)			T1: Business (N = 245)			T2: Creative (N = 253)		
	Mean	SD	Median	Mean	SD	Median	Mean	SD	Median
Panel A: Entrepreneur Demographics									
Age (years)	44.8	9.8	45.0	44.1	10.1	43.0	44.1	10.1	44.0
Female	0.65	0.48	1.00	0.65	0.48	1.00	0.60	0.49	1.00
Higher Education	0.44	0.50	0.00	0.43	0.50	0.00	0.48	0.50	0.00
Santiago Metro Region	0.45	0.50	0.00	0.49	0.50	0.00	0.47	0.50	0.00
Panel B: Business Characteristics									
Business Age (years)	4.31	4.31	3.00	4.24	4.16	3.00	4.13	3.79	3.00
Employees (full-time)	1.53	0.91	1.00	1.79	3.21	1.00	1.58	0.94	1.00
Works 5+ Days per Week	0.65	0.48	1.00	0.66	0.48	1.00	0.60	0.49	1.00
Monthly Sales (CLP)	1,155,086	1,980,653	340,000	999,255	1,596,162	350,000	1,100,289	1,866,873	350,000
Monthly Profits (CLP)	541,230	779,411	200,000	526,709	719,313	200,000	487,779	712,186	200,000
Panel C: AI Use and Beliefs									
Any AI Use	0.71	0.45	1.00	0.73	0.44	1.00	0.75	0.44	1.00
Profit Expectations (1–7)	5.82	1.06	6.00	5.68	1.17	6.00	5.81	1.12	6.00
Sales Expectations (1–7)	5.79	1.14	6.00	5.72	1.14	6.00	5.82	1.07	6.00
WTP Baseline (CLP)	54,957	47,588	35,000	57,094	44,453	65,000	57,744	46,252	65,000

Notes: Summary statistics for the baseline sample (N = 736). Age is self-reported. Santiago Metro Region indicates residence in the Metropolitana de Santiago region; the remaining 53% are distributed across 14 other regions. Employees (full-time) includes the entrepreneur. Monthly sales and profits are in Chilean pesos (CLP; 1 USD ≈ 900 CLP), winsorized at the 99th percentile. Profit and sales expectations are on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). Any AI Use is a binary indicator for prior use of any AI tool for business purposes. WTP Baseline is the lower-bound willingness to pay for a personalized AI implementation coaching session, elicited via a multiple price list as the last accepted price threshold; N varies from 230–246 across groups due to non-response.

Table 2: Baseline Balance Across Treatment Arms

	Biz Age (years)	Sales (CLP)	Profits (CLP)	Any AI Use	Profit Exp. (1–7)	Sales Exp. (1–7)
T1: Business Consultant	−0.167	−109 791	5716	0.029	−0.208*	−0.105
	(0.395)	(168 493)	(71 991)	(0.043)	(0.108)	(0.108)
	[0.673]	[0.515]	[0.937]	[0.501]	[0.054]	[0.330]
T2: Creative Assistant	−0.210	−109 364	−84 112	0.044	−0.004	0.056
	(0.365)	(170 012)	(68 030)	(0.040)	(0.101)	(0.100)
	[0.564]	[0.520]	[0.217]	[0.273]	[0.972]	[0.572]
Observations	729	729	729	729	729	729
Adj. R^2	0.020	0.058	0.044	0.037	0.009	0.030
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports coefficients from a regression of the baseline characteristic on treatment indicators with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Sales and profits are monthly values in Chilean pesos (CLP), winsorized at the 99th percentile. Profit Exp. and Sales Exp. are expected changes on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). Any AI Use is a binary indicator for prior use of any AI tool. Female and education are absorbed by strata fixed effects (stratification on gender × education × macro-region). Profit Exp. shows a marginally significant difference for T1 ($p = 0.054$); we control for baseline beliefs in all outcome regressions. Sample: $N = 729$; seven observations excluded due to missing stratification-cell identifiers. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Treatment Effects on Profit Expectations

T1: Business Consultant	0.164**	0.158**		
	(0.070)	(0.071)		
	[0.020]	[0.027]		
T2: Creative Assistant	0.096	0.100		
	(0.066)	(0.067)		
	[0.146]	[0.137]		
H_0: T1 = T2	[>0.10]	[>0.10]		
Any Treatment			0.128**	0.043
			(0.059)	(0.061)
			[0.030]	[0.479]
Low Optimism				−0.526***
				(0.155)
				[<0.001]
Treatment x Low Optimism				0.549***
				(0.180)
				[0.002]
Observations	729	729	729	729
Adj. R^2	0.507	0.507	0.507	0.515
Strata FE	Yes	Yes	Yes	Yes
ANCOVA Controls	No	Yes	No	No
Control Mean	5.941	5.941	5.941	5.941

Notes: OLS estimates. Dependent variable: post-treatment expected profit change on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). All specifications include the baseline profit expectation and strata fixed effects. ANCOVA controls (Column2): business age, winsorized sales and profits, and prior AI use. Column(3) pools both treatment arms. Column (4) introduces heterogeneity by baseline optimism: Low Optimism is a binary indicator for baseline profit expectations ≤ 4 ("no change" or below); the positive Treatment × Low Optimism interaction is consistent with Bayesian learning. See Table A6 for the continuous prior specification and implied learning rate ($\alpha \approx 0.25$). Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. $N = 729$; seven observations excluded due to missing stratification-cell identifiers. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table 4: Treatment Effects on Adoption and Business Outcomes

	By Treatment Arm					Pooled				
	(1) AI Freq	(2) New Adopter	(3) BP Total	(4) Sales (CLP)	(5) Profits (CLP)	(6) AI Freq	(7) New Adopter	(8) BP Total	(9) Sales (CLP)	(10) Profits (CLP)
T1: Business Consultant	0.213	−0.037	−0.399	−382 504.846	17 788.394					
	(0.271)	(0.050)	(0.777)	(338 832.454)	(118 538.449)					
	[0.434]	[0.461]	[0.608]	[0.260]	[0.881]					
T2: Creative Assistant	0.089	−0.004	0.473	−39 525.617	−27 295.724					
	(0.237)	(0.043)	(0.705)	(386 971.299)	(108 009.267)					
	[0.709]	[0.927]	[0.503]	[0.919]	[0.801]					
Any Treatment						0.142	−0.018	0.094	−188 603.736	−7699.623
						(0.223)	(0.040)	(0.635)	(337 963.130)	(97 116.741)
						[0.525]	[0.651]	[0.882]	[0.577]	[0.937]
Course Invited	−0.088	0.009	−0.347	−146 978.512	−18 403.436	−0.080	0.007	−0.408	−170 769.953	−15 276.086
	(0.205)	(0.039)	(0.633)	(266 392.464)	(89 582.864)	(0.201)	(0.039)	(0.633)	(268 011.848)	(89 205.105)
	[0.668]	[0.821]	[0.584]	[0.582]	[0.837]	[0.693]	[0.866]	[0.520]	[0.525]	[0.864]
Observations	231	231	231	231	231	231	231	231	231	231
Adj. R^2	0.419	0.460	−0.024	0.224	0.312	0.421	0.461	−0.026	0.224	0.315
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	3.831	0.169	12.197	1450879	502568	3.831	0.169	12.197	1450879	502568

Notes: OLS estimates with ANCOVA controls and strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. AI Freq is self-reported AI usage frequency for business on a five-point scale (1 = never, 5 = always). New Adopter is a binary indicator for entrepreneurs who did not use AI at baseline but report using it at follow-up. BP Total is the McKenzie and Woodruff (2017) business practices index (0–19). Sales and profits are monthly values in CLP, winsorized at the 99th percentile. ANCOVA controls: baseline profit expectations, business age, winsorized baseline sales and profits, prior AI use. Columns (1)–(5) report separate treatment arm effects; columns (6)–(10) pool both arms. Course Invited is a binary indicator for the Experiment 2 randomized invitation to a free AI training course, cross-randomized independently of vignette assignment. Negative Adj. R^2 values (columns 3, 8) indicate the model fits no better than the within-stratum mean. All outcomes measured at two-month follow-up ($N = 231$; nine observations excluded due to missing baseline covariates). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A1: Treatment Effects on Willingness to Pay

	WTP Level (CLP)		WTP Increased	
	(1) By Arm	(2) Pooled	(3) By Arm	(4) Pooled
T1: Business Consultant	982		0.039	
	(1821)		(0.025)	
	[0.590]		[0.117]	
T2: Creative Assistant	2354		0.019	
	(1617)		(0.023)	
	[0.146]		[0.416]	
Any Treatment		1705		0.029
		(1491)		(0.020)
		[0.253]		[0.162]
WTP Baseline (CLP)	0.895***	0.895***	0.000	0.000
	(0.020)	(0.020)	(0.000)	(0.000)
	[<0.001]	[<0.001]	[0.283]	[0.284]
Observations	703	703	703	703
Adj. R^2	0.845	0.845	0.031	0.031
Strata FE	Yes	Yes	Yes	Yes
Control Mean	54326	54326	0.061	0.061

Notes: OLS estimates with ANCOVA controls and strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. WTP is the post-treatment lower-bound willingness to pay (last accepted price) in CLP for a personalized AI implementation coaching session, elicited via a multiple price list with five binary choices (5,000–115,000 CLP); WTP equals zero if all offers are rejected. WTP Increased = 1 if post-treatment WTP strictly exceeds own baseline WTP. Both outcomes are elicited within the same survey session, immediately following vignette exposure. $N = 703$; 33 observations excluded due to non-response in the baseline WTP elicitation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A2: Within-Session Willingness to Pay, by Baseline Optimism

	WTP Level (CLP)		WTP Increased	
	(1) Full	(2) Low-Opt.	(3) Full	(4) Low-Opt.
Any Treatment	1,865	2,672	0.015	0.065
	(1,644)	(5,051)	(0.022)	(0.051)
	[0.257]	[0.599]	[0.500]	[0.205]
Low Optimism	3,064		-0.116**	
	(3,023)		(0.047)	
	[0.311]		[0.014]	
Treatment × Low Optimism	-811		0.083*	
	(3,305)		(0.046)	
	[0.806]		[0.069]	
Observations	703	91	703	91
Control Mean	54,326	28,143	0.061	0.029
WTP Baseline Control	Yes	Yes	Yes	Yes
ANCOVA Controls	Yes	Yes	Yes	Yes
Strata FE	Yes	Yes	Yes	Yes

Notes: OLS estimates with ANCOVA controls (baseline profit expectations, business age, winsorized sales and profits, and prior AI use) and strata fixed effects; heteroskedasticity-robust (HC1) standard errors in parentheses and exact p-values in brackets. *Low Optimism* equals one for entrepreneurs whose baseline profit expectation is at or below the "no change" midpoint (≤ 4 on the 1–7 scale) – the subgroup that revises beliefs most. Columns (1) and (3) use the full baseline sample with a treatment × low-optimism interaction; columns (2) and (4) restrict to the low-optimism subsample. WTP Level is the post-treatment lower-bound willingness to pay (CLP) for an AI implementation coaching session; WTP Increased equals one if post-treatment WTP strictly exceeds own baseline WTP. Both are elicited within the same survey session, so the comparison is free of follow-up attrition. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A3: Baseline Characteristics: Stayers vs. Attritors

	Biz Age (years)	Sales (CLP)	Profits (CLP)	Any AI Use	Profit Exp. (1–7)	Sales Exp. (1–7)
Panel Stayer	0.241	−157 458.016	−45 067.590	0.035	0.068	0.081
	(0.343)	(137 463.042)	(57 302.741)	(0.036)	(0.094)	(0.092)
	[0.482]	[0.252]	[0.432]	[0.325]	[0.470]	[0.376]
Observations	729	729	729	729	729	729
Adj. R^2	0.022	0.061	0.043	0.038	0.004	0.029
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column regresses a baseline characteristic on a binary indicator for appearing in the two-month follow-up, with strata fixed effects. $N = 729$; seven observations excluded due to missing stratification-cell identifiers. Of the full baseline sample of 736 participants, 240 (32.6%) completed the follow-up. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. No characteristic differs significantly, suggesting attrition is orthogonal to baseline observables. Female and education are absorbed by strata fixed effects. Treatment assignment does not predict follow-up response: T1 coefficient = 0.009 ($p = 0.837$); T2 coefficient = 0.046 ($p = 0.288$). $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A4: Lee (2009) Bounds on Treatment Effects (Pooled)

	Lower Bound	Upper Bound	Point Est.	Control Mean	N
BP Total	-0.566	1.306	0.258	12.197	240
BP Inventory	-0.021	0.358	0.129	1.725	233
AI Freq	-0.086	0.585	0.181	3.831	240
New Adopter	-0.149	-0.015	-0.033	0.169	240

Notes: Lee (2009) trimming bounds on pooled treatment effects for the two-month follow-up. Bounds are constructed by trimming the higher-response group by a fraction $q = |p_1 - p_0|/\max(p_1, p_0) = 12.09\%$, where p_1 and p_0 are follow-up response rates in the treatment and control groups. The lower bound trims the top q-fraction of outcomes from the higher-response group; the upper bound trims the bottom q-fraction. "Point Est." is the unadjusted OLS estimate without trimming. BP Inventory is the inventory management sub-index of the McKenzie and Woodruff (2017) business practices index ($N = 233$ due to item non-response). For New Adopter, both bounds lie below zero; we interpret this cautiously given the low base rate of new adoption (control mean: 16.9%), where trimming few observations can shift both bounds negative. No confidence intervals are reported; zero inclusion or exclusion within the bounds is the relevant criterion under worst-case attrition.

Table A5: IPW-Weighted Robustness: Key Outcomes

	BP Total	BP Inventory	AI Freq	New Adopter	BP Total (IPW)	BP Inventory (IPW)	AI Freq (IPW)	New Adopter (IPW)
	Unweighted				IPW-Weighted			
	(1) BP Total	(2) BP Inv.	(3) AI Freq	(4) New Ad.	(5) BP Total	(6) BP Inv.	(7) AI Freq	(8) New Ad.
T1: Business Consultant	−0.399	−0.060	0.063	−0.037	−0.494	−0.062	0.076	−0.042
	(0.777)	(0.186)	(0.288)	(0.050)	(0.759)	(0.187)	(0.291)	(0.052)
	[0.608]	[0.748]	[0.827]	[0.461]	[0.516]	[0.740]	[0.795]	[0.428]
T2: Creative Assistant	0.473	0.098	0.082	−0.004	0.511	0.112	0.084	−0.007
	(0.705)	(0.180)	(0.253)	(0.043)	(0.695)	(0.179)	(0.252)	(0.046)
	[0.503]	[0.588]	[0.746]	[0.927]	[0.463]	[0.535]	[0.740]	[0.879]
Course Invited	−0.347	0.130	−0.101	0.009	−0.426	0.115	−0.076	0.006
	(0.633)	(0.149)	(0.220)	(0.039)	(0.615)	(0.148)	(0.222)	(0.042)
	[0.584]	[0.386]	[0.647]	[0.821]	[0.489]	[0.436]	[0.732]	[0.884]
Observations	231	224	231	231	231	224	231	231
Adj. R^2	−0.024	−0.070	0.306	0.460	−0.007	−0.066	0.318	0.445
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: OLS estimates with ANCOVA controls and strata fixed effects. Columns (1)–(4) replicate the main panel specifications (unweighted). Columns (5)–(8) apply inverse probability weights estimated from a probit model of panel response on baseline characteristics (business age, winsorized sales and profits, prior AI use, baseline profit and sales expectations, gender, education, and treatment indicators). Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. BP Total = business practice index (0–19); BP Inv. = inventory sub-index; AI Freq = AI usage frequency (1–5); New Ad. = new AI adopter indicator. Results are virtually identical across weighted and unweighted specifications, indicating that attrition does not drive the null adoption effects. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Bayesian Learning: Treatment x Prior Belief Interaction

	Pooled (1) Pooled	By Arm (2) By Arm
Any Treatment	1.213***	
	(0.372)	
	[0.001]	
T1: Business Consultant		1.279***
		(0.446)
		[0.004]
T2: Creative Assistant		1.123**
		(0.472)
		[0.018]
Prior Belief	0.753***	0.753***
	(0.042)	(0.042)
	[<0.001]	[<0.001]
Treatment x Prior	−0.187***	
	(0.061)	
	[0.002]	
T1 x Prior		−0.194***
		(0.073)
		[0.008]
T2 x Prior		−0.175**
		(0.077)
		[0.022]
Observations	729	729
Adj. R^2	0.515	0.514
Strata FE	Yes	Yes

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Dependent variable: post-treatment profit expectations (1–7 scale). Prior Belief is the continuous baseline profit expectation (1–7); lower values indicate more pessimistic priors. The negative Treatment × Prior coefficient is consistent with Bayesian learning (Equation **??**): entrepreneurs with lower priors update more, as the prior-signal gap is larger. The implied learning rate $\alpha = -\hat{\gamma}_3/\hat{\gamma}_2 = 0.187/0.753 \approx 0.25$. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A7: Treatment Effects on Alternative Belief Measures

	By Treatment Arm			Pooled		
	(1) Profit Exp.	(2) Sales Exp.	(3) Belief Index	(4) Profit Exp.	(5) Sales Exp.	(6) Belief Index
T1: Business Consultant	0.164**	0.107	0.177**			
	(0.070)	(0.073)	(0.087)			
	[0.020]	[0.143]	[0.041]			
T2: Creative Assistant	0.096	0.044	0.083			
	(0.066)	(0.071)	(0.083)			
	[0.146]	[0.529]	[0.316]			
Any Treatment				0.128**	0.074	0.128*
				(0.059)	(0.063)	(0.074)
				[0.030]	[0.241]	[0.084]
Observations	736	736	736	736	736	736
Adj. R^2	0.507	0.499	0.219	0.507	0.500	0.219
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	5.941	5.975	$-0.101^{\dagger}$	5.941	5.975	$-0.101^{\dagger}$

Notes: OLS estimates with ANCOVA controls and strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Profit Exp. and Sales Exp. are post-treatment expectations on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). Belief Index is an Anderson (2008) inverse-variance-weighted composite of profit and sales expectations, standardized to the full-sample mean and covariance.[†] The belief revision is concentrated in profit expectations, consistent with the Business Consultant vignette (T1) emphasizing analytical and financial applications. The pooled Belief Index effect is marginally significant ($p = 0.084$), consistent with a positive direction across both components. [†]The control-group Belief Index mean of −0.101 reflects full-sample standardization: when treatment groups shift beliefs upward, the control group sits slightly below the full-sample mean. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A8: Falsification: Treatment on Pre-Treatment Beliefs

	Baseline Profit Exp.	Baseline Sales Exp.
T1: Business Consultant	−0.208*	−0.105
	(0.108)	(0.108)
	[0.054]	[0.330]
T2: Creative Assistant	−0.004	0.056
	(0.101)	(0.100)
	[0.972]	[0.572]
Observations	729	729
Adj. R^2	0.009	0.030
Strata FE	Yes	Yes
Control Mean	5.824	5.790

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Dependent variables are pre-treatment (baseline) profit and sales expectations on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). Because treatment is assigned after baseline measurement, coefficients should be zero. The T1 coefficient on Profit Exp. is marginally significant at the 10% level ($p = 0.054$); we interpret this as sampling variation given that the effect is absent for T2 and for Sales Exp., and baseline beliefs enter as ANCOVA controls in all outcome regressions. $N = 729$; seven observations excluded due to missing stratification-cell identifiers. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A9: Treatment Effects on Course Interest (Baseline, $N = 736$)

	(1) By Arm	(2) Pooled	(3) ANCOVA
T1: Business Consultant	0.006		
	(0.015)		
	[0.674]		
T2: Creative Assistant	−0.012		
	(0.015)		
	[0.421]		
Any Treatment		−0.003	−0.004
		(0.012)	(0.012)
		[0.780]	[0.758]
Observations	729	729	729
Adj. R^2	0.007	0.006	0.004
Strata FE	Yes	Yes	Yes
Control Mean	0.979	0.979	0.979

Notes: OLS estimates with strata fixed effects. Column (3) adds ANCOVA controls (business age, winsorized sales and profits, prior AI use). Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Course Interest is a binary indicator equal to one if the respondent expressed interest in enrolling in the AI training course, elicited immediately after vignette exposure in the same session as the baseline. The near-universal rate (control mean = 0.979) reflects that entrepreneurs overwhelmingly wanted to learn about AI regardless of treatment assignment. The null treatment effects indicate that updated profit beliefs do not translate into increased stated course interest. $N = 729$; seven observations excluded due to missing stratification-cell identifiers. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A10: Treatment Effects on Follow-Up Beliefs (Two-Month Panel)

	Profit Exp. (t3)	Sales Exp. (t3)	Profit Exp. (t3)	Sales Exp. (t3)	Profit (Interact)
T1: Business Consultant	0.032	0.206			
	(0.170)	(0.216)			
	[0.851]	[0.342]			
T2: Creative Assistant	0.028	−0.035			
	(0.156)	(0.197)			
	[0.859]	[0.860]			
Any Treatment			0.030	0.072	−0.064
			(0.140)	(0.180)	(0.152)
			[0.833]	[0.691]	[0.672]
Low Optimism					−0.513
					(0.473)
					[0.279]
Treatment x Low Optimism					0.668
					(0.477)
					[0.163]
Observations	231	231	231	231	231
Adj. R^2	0.124	0.166	0.129	0.163	0.130
Strata FE	Yes	Yes	Yes	Yes	Yes
Control Mean	5.915	5.831	5.915	5.831	5.915

Notes: OLS estimates with strata fixed effects. Dependent variables are profit and sales expectations measured at the two-month follow-up on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). All specifications control for the corresponding baseline expectation and the Experiment 2 course invitation indicator. Low Optimism is a binary indicator equal to one if baseline profit expectations are at or below the sample median. Columns (1)–(2) report separate treatment arm effects; columns (3)–(4) pool both arms; column (5) adds a Treatment × Low Optimism interaction to examine heterogeneity by baseline optimism. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. $N = 231$; nine observations excluded due to missing baseline covariates. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A11: Heterogeneity by Prior AI Use (Interaction Model)

	(1) Beliefs	(2) AI Freq	(3) New Adopter	(4) BP Total
Any Treatment	0.182	−0.102	−0.053	−1.476
	(0.114)	(0.511)	(0.161)	(1.224)
	[0.110]	[0.842]	[0.743]	[0.230]
Prior AI User	0.074	2.086***	−0.576***	0.077
	(0.103)	(0.489)	(0.139)	(1.147)
	[0.474]	[<0.001]	[<0.001]	[0.947]
Treatment × Prior AI User	−0.076	0.232	0.046	2.078
	(0.134)	(0.573)	(0.165)	(1.481)
	[0.570]	[0.685]	[0.782]	[0.162]
Course Invited		−0.092	0.008	−0.321
		(0.217)	(0.040)	(0.635)
		[0.670]	[0.831]	[0.614]
Observations	729	231	231	231
Adj. R^2	0.506	0.307	0.459	−0.023
Strata FE	Yes	Yes	Yes	Yes

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Column (1): post-treatment profit expectations measured in the baseline session ($N = 729$; seven observations excluded due to missing stratification-cell identifiers). Columns (2)–(4): two-month follow-up outcomes ($N = 231$; nine observations excluded due to missing baseline covariates). Course Invited is a binary indicator for the Experiment 2 randomized course invitation, included as a control in follow-up specifications. Prior AI User is a binary indicator for any prior AI use reported at baseline. The Treatment × Prior AI User interaction tests whether the belief-action gap differs by prior AI experience; all interactions are statistically insignificant ($p > 0.16$), indicating that the null adoption result holds for both prior users (intensive margin) and non-users (extensive margin). $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table A12: T1 vs. T2 Adoption: Heterogeneity by Baseline Optimism (Interaction Model)

	AI Freq	New Adopter	BP Total
T1: Business Consultant	−0.058	−0.052	−0.225
	(0.312)	(0.051)	(0.881)
	[0.853]	[0.307]	[0.799]
T2: Creative Assistant	0.067	−0.016	0.506
	(0.270)	(0.043)	(0.750)
	[0.806]	[0.711]	[0.501]
Low Optimism	0.099	−0.055	0.587
	(0.618)	(0.165)	(1.740)
	[0.873]	[0.738]	[0.736]
T1 × Low Optimism	0.855	0.107	−0.995
	(0.691)	(0.175)	(1.981)
	[0.218]	[0.539]	[0.616]
Course Invited	−0.134	0.006	−0.352
	(0.223)	(0.039)	(0.645)
	[0.550]	[0.879]	[0.586]
Observations	231	231	231
Adj. R^2	0.306	0.453	−0.039
Strata FE	Yes	Yes	Yes

Notes: OLS estimates with ANCOVA controls, strata fixed effects, and interaction terms. T1 = Business Consultant vignette; T2 = Creative Assistant vignette; omitted category is the control group. Low Optimism is a binary indicator equal to one if baseline profit expectations are at or below the neutral midpoint (≤ 4 on the 1–7 scale), consistent with the definition used throughout the paper. The interaction T1 × Low Optimism tests whether entrepreneurs who were initially more pessimistic exhibit differential adoption responses to the Business Consultant vignette (the arm most directly targeting the profit channel). T2 × Low Optimism is not reported as T2 targets creative rather than financial applications; the coefficient is similar in magnitude and insignificant. Course Invited is a binary indicator for the cross-randomized Experiment 2 invitation. All outcomes measured at two-month follow-up. Negative Adj. R^2 (BP Total) indicates the model fits no better than the within-stratum mean. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. $N = 231$; nine observations excluded due to missing baseline covariates. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A12 tests whether the null adoption result differs by baseline optimism level. Low-optimism entrepreneurs (those who update beliefs most substantially by construction) show no differential adoption response to either treatment arm ($p > 0.10$ for all T1 × Low Optimism interactions), indicating that structural barriers to adoption bind regardless of the magnitude of belief revision.

Table A13: Pre-Registration to Analysis Mapping

Hyp.	Outcome	Units	Timing	N	ITT Estimate	p-value
Panel A: Within-session outcomes — full sample, no attrition (primary behavioral evidence)						
H1	Profit expectations	7-pt scale	Session	729	+0.128	0.030**
H5	WTP for AI training	CLP	Session	703	+1,705	0.253
H6	WTP Increased	0/1	Session	703	+0.029	0.162
H7	Course interest	0/1	Session	729	−0.003	0.780
Panel B: Follow-up outcomes — limited precision due to 68% attrition (confirmatory)						
H2	AI usage frequency	5-pt scale	Follow-up	231	+0.154	0.485
H3	New AI adopter	0/1	Follow-up	231	−0.045	0.410
H4	Business practices index	0–19 pts	Follow-up	231	+0.164	0.794
H8	Monthly sales	CLP	Follow-up	231	−166,701	0.622
H9	Monthly profits	CLP	Follow-up	231	−10,161	0.916

Notes: Each row corresponds to a pre-registered hypothesis. All ITT estimates are pooled (Any Treatment vs. Control) from OLS specifications with strata fixed effects (plus ANCOVA controls for H5–H6). Exact p-values from heteroskedasticity-robust (HC1) standard errors. Panel A outcomes are measured within the baseline survey session, immediately after vignette exposure; they are estimated at the full sample with no attrition and constitute the paper's primary behavioral evidence (Section 6.2). Panel B outcomes are measured at the follow-up (approximately four months after baseline); 68% attrition from baseline yields a follow-up sample of $N = 231$, limiting precision for these outcomes. The Benjamini–Hochberg multiple-testing correction applies to H1–H4 (the four pre-registered primary outcomes): after correction, H1 remains significant at 5% and H2–H4 remain null. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A14: Treatment Effects on Financial Outcomes: log(1+x) Specification

T1: Business Consultant	−0.789	−0.785	−0.290	−0.288
	(0.515)	(0.507)	(0.608)	(0.613)
	[0.127]	[0.123]	[0.634]	[0.640]
T2: Creative Assistant	−0.509	−0.490	−1.164*	−1.163*
	(0.482)	(0.486)	(0.591)	(0.591)
	[0.293]	[0.314]	[0.051]	[0.051]
Course Invited (Experiment 2)	−0.308	−0.325	−0.196	−0.216
	(0.438)	(0.441)	(0.514)	(0.519)
	[0.483]	[0.463]	[0.703]	[0.678]
Observations	231	231	231	231
Adj. R^2	0.225	0.216	0.168	0.152
Strata FE	Yes	Yes	Yes	Yes
ANCOVA Controls	No	Yes	No	Yes
Control Mean	12.755	12.755	11.677	11.677

Notes: OLS estimates with strata fixed effects. Dependent variables are log(1 + monthly sales) and log(1 + monthly profits) at the two-month follow-up, winsorized at the 99th percentile before transformation. Columns (1) and (3) include only the log-transformed own baseline and strata fixed effects; columns (2) and (4) add ANCOVA controls (baseline profit expectations, business age, log baseline sales, log baseline profits, any prior AI use). Course Invited is a cross-randomized indicator for invitation to a free AI training course; invited participants had no access to the course during the control period. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. $N = 231$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A15: Equivalence Tests (TOST) for Within-Session Outcomes

	Estimate (SD)	SE (SD)	90% CI Lower	90% CI Upper	Equivalent? (±0.20 SD)
WTP Level	0.036	0.032	-0.016	0.088	Yes
WTP Increased	0.119	0.085	-0.021	0.259	No

Notes: Two one-sided tests (TOST) for equivalence against a ±0.20 SD bound. Estimates and standard errors are standardized by the control-group standard deviation, from pooled OLS (Any Treatment vs. Control) with the own baseline value, ANCOVA controls, and strata fixed effects; heteroskedasticity-robust (HC1) standard errors. The 90% confidence interval corresponds to two one-sided tests at the 5% level. "Equivalent? = Yes" indicates the 90% CI lies entirely within $(-0.20, 0.20)$ SD, rejecting effects larger than 0.20 SD. Both outcomes are elicited within the baseline survey session on the full sample ($N = 703$), free of follow-up attrition. $\delta = 0.20$ SD of the control-group outcome distribution.

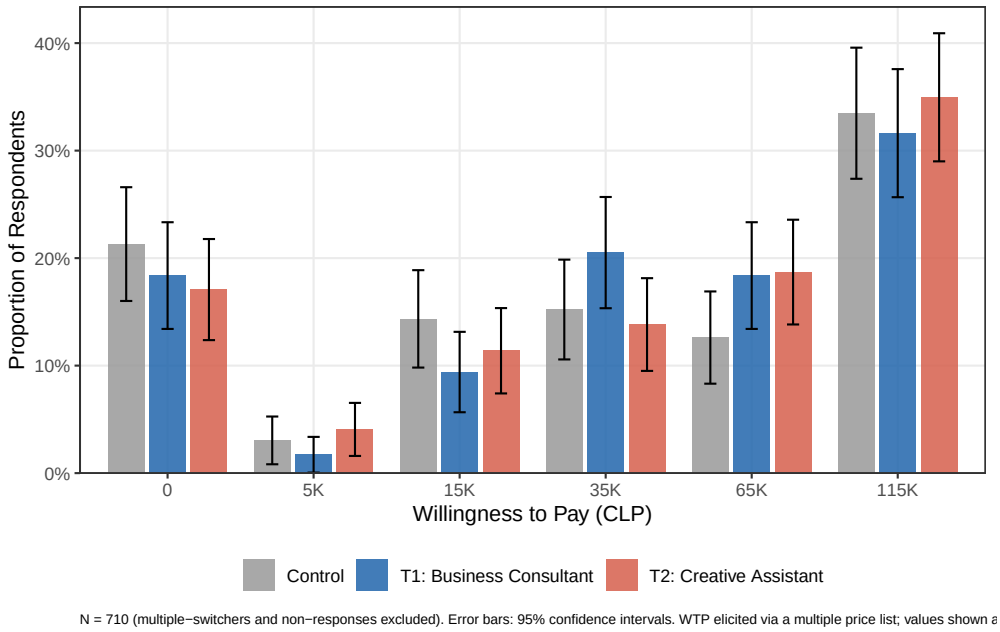


Figure A1: WTP Distribution by Treatment Arm ($N \approx 703$). Bars show the proportion of respondents in each arm whose stated WTP equals each threshold (0, 5K, 15K, 35K, 65K, or 115K CLP); error bars are 95% confidence intervals. No threshold shows a statistically significant difference across arms, confirming the pooled null in Table A1 is not an artifact of distributional shifts that cancel in the mean.

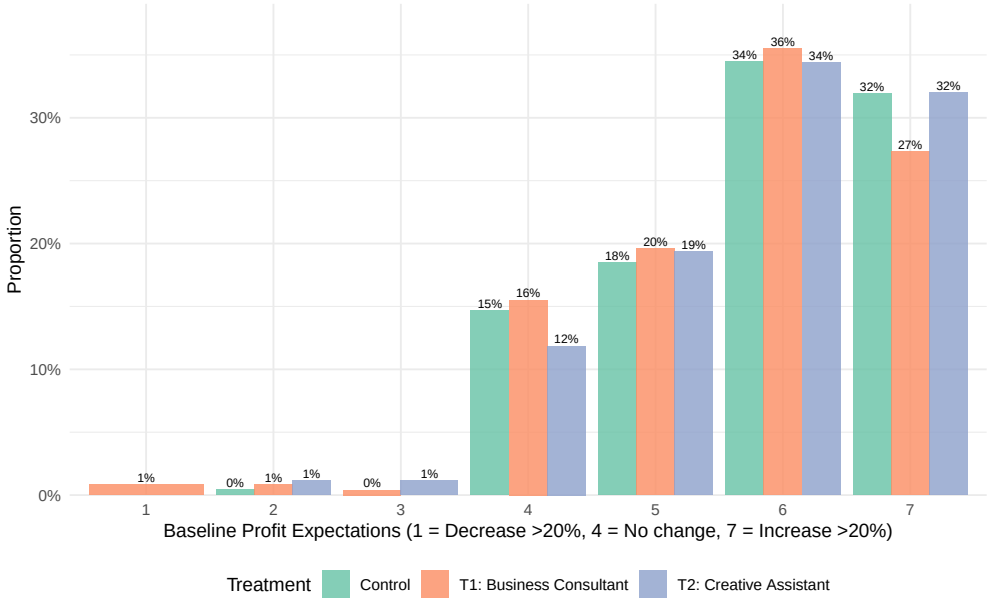


Figure A2: Distribution of Baseline Beliefs about AI's Impact on Profits ($N = 736$). Pre-treatment profit expectations on a seven-point scale (1 = decrease >20%, 4 = no change, 7 = increase >20%). The modal response is 6, with roughly 30% at the ceiling (category 7) and fewer than 2% reporting expectations below *no change* (categories 1–3), leaving limited upward revision room for already-optimistic entrepreneurs.

Table B1: Treatment Effects on Course Engagement

	(1) Completed	(2) Started	(3) In Progress
T1: Business Consultant	−0.011	0.033	0.044
	(0.020)	(0.028)	(0.032)
	[0.583]	[0.245]	[0.165]
T2: Creative Assistant	−0.011	0.003	0.014
	(0.020)	(0.028)	(0.032)
	[0.595]	[0.911]	[0.662]
Course Invited	0.090***	0.773***	0.683***
	(0.016)	(0.022)	(0.025)
	[<0.001]	[<0.001]	[<0.001]
Observations	729	729	729
Adj. R^2	0.016	0.633	0.522
Strata FE	Yes	Yes	Yes
Control Mean	0.046	0.361	0.315

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. Course engagement measured over the full baseline sample ($N = 736$). Completed = finished all four modules of the AI training course. Started = began the course (includes both completers and in-progress). In Progress = started but did not complete. Course Invited is the Experiment 2 randomized treatment. The control group completion rate (4.6%) reflects entrepreneurs who enrolled independently of the experimental invitation. Experiment 1 treatment assignment does not predict differential course engagement. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B2: Prompt Engineering Sophistication (PESI): Course Effects

	ITT (Invitation)			With Completion Control		
	(1) PESI Sum	(2) PESI SW	(3) PESI Extended	(4) PESI Sum	(5) PESI SW	(6) PESI Extended
Course Invited (ITT)	−0.390***	−0.543***	−0.409***	−0.550***	−0.327**	−0.471***
	(0.139)	(0.138)	(0.139)	(0.137)	(0.139)	(0.139)
	[0.005]	[<0.001]	[0.004]	[<0.001]	[0.020]	[<0.001]
Course Completed		0.726***		0.667**		0.682**
		(0.271)		(0.267)		(0.269)
		[0.008]		[0.013]		[0.012]
Observations	231	231	231	231	231	231
Adj. R^2	0.110	0.141	0.073	0.099	0.117	0.144
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	0.014	0.014	0.032	0.032	0.008	0.008

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. PESI Sum is the raw Prompt Engineering Sophistication Index (0–7), combining technique (0–3) and specificity (0–4). PESI SW is an inverse-variance-weighted composite (Anderson 2008). PESI Extended is a standardized version of PESI Sum. Near-zero control means reflect that the indices are standardized to the full sample. PESI is measured at two-month follow-up ($N = 240$). Columns (1)–(3) report the intention-to-treat effect of course invitation. Columns (4)–(6) add a course completion indicator. The negative ITT reflects that most invitees did not engage substantively; course completers score significantly higher. Course completion is endogenous; the coefficient should not be interpreted causally. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B3: Treatment Effects on Prompt Engineering Sophistication (PESI)

	(1) PESI (T1/T2)	(2) PESI (+Compl.)	(3) PESI (No Course)	(4) PESI (Invited)	(5) SW (T1/T2)	(6) SW (+Compl.)	(7) Extended	(8) BP Coverage	(9) Instructions	(10) Words
T1: Business Consultant	−0.046	−0.049	0.038	−0.063	−0.085	−0.087	−0.017	0.035	−0.186	−9.887
	(0.178)	(0.174)	(0.236)	(0.292)	(0.188)	(0.184)	(0.174)	(0.108)	(0.178)	(6.413)
	[0.795]	[0.778]	[0.873]	[0.828]	[0.653]	[0.636]	[0.922]	[0.745]	[0.298]	[0.125]
T2: Creative Assistant	0.120	0.114	0.363	−0.226	0.112	0.107	0.128	0.064	−0.027	−1.815
	(0.166)	(0.164)	(0.248)	(0.290)	(0.172)	(0.170)	(0.160)	(0.099)	(0.133)	(5.205)
	[0.471]	[0.486]	[0.148]	[0.440]	[0.516]	[0.531]	[0.425]	[0.514]	[0.839]	[0.728]
Course Invited	−0.385***	−0.537***			−0.399***	−0.539***	−0.469***	−0.084	−0.003	−4.531
	(0.141)	(0.140)			(0.143)	(0.141)	(0.142)	(0.088)	(0.116)	(4.559)
	[0.007]	[<0.001]			[0.006]	[<0.001]	[0.001]	[0.338]	[0.976]	[0.322]
Course Completed		0.723***		0.983***		0.664**	0.679**	0.194	0.245	10.196
		(0.275)		(0.311)		(0.272)	(0.274)	(0.178)	(0.213)	(7.678)
		[0.009]		[0.002]		[0.015]	[0.014]	[0.279]	[0.253]	[0.186]
Observations	231	231	103	105	231	231	231	231	231	231
Adj. R^2	0.106	0.137	0.207	0.120	0.070	0.096	0.139	0.028	0.088	0.113
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. PESI is the Prompt Engineering Sophistication Index (higher = more sophisticated prompts). SW = Specificity-Wordiness sub-index. Columns (1) and (5): main ITT specifications without course completion. Columns (2) and (6): add Course Completed indicator (endogenous; completers are not randomly selected). Column (3): subsample not invited to the course. Column (4): subsample invited to the course. Columns (7)–(10): PESI sub-components. The negative ITT effect of course invitation (cols 1, 5) is consistent with a discouragement mechanism: formal training access may suppress informal experimentation among non-completers. The positive Course Completed coefficient reflects the completer premium, interpreted as suggestive rather than causal. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B4: Course Completion Effects on Outcomes ($N = 231$)

	(1) Profit Bel. (Compl.)	(2) Profit Bel. (Started)	(3) Belief Idx (Compl.)	(4) Belief Idx (Started)	(5) AI Freq (Compl.)	(6) AI Freq (Started)	(7) New Adopter (Compl.)	(8) New Adopter (Started)	(9) PESI (Compl.)	(10) PESI (Started)
Course Completed	0.011		0.016		0.595*		−0.052		0.725***	
	(0.192)		(0.271)		(0.356)		(0.077)		(0.272)	
	[0.954]		[0.953]		[0.096]		[0.505]		[0.008]	
Course Started		0.022		−0.063		0.321		0.149**		0.103
		(0.216)		(0.273)		(0.345)		(0.073)		(0.268)
		[0.920]		[0.819]		[0.354]		[0.043]		[0.702]
Course Invited	0.023	0.008	0.067	0.122	−0.205	−0.345	0.017	−0.117	−0.549***	−0.481*
	(0.128)	(0.200)	(0.140)	(0.248)	(0.220)	(0.334)	(0.043)	(0.072)	(0.141)	(0.257)
	[0.857]	[0.970]	[0.633]	[0.623]	[0.353]	[0.303]	[0.690]	[0.104]	[<0.001]	[0.063]
Any Treatment	0.152	0.152	0.155	0.155	0.139	0.139	−0.018	−0.020	0.043	0.047
	(0.117)	(0.116)	(0.131)	(0.131)	(0.223)	(0.224)	(0.041)	(0.040)	(0.148)	(0.152)
	[0.194]	[0.193]	[0.240]	[0.235]	[0.534]	[0.533]	[0.658]	[0.618]	[0.771]	[0.760]
Observations	231	231	231	231	231	231	231	231	231	231
Adj. R^2	0.503	0.503	0.266	0.266	0.426	0.420	0.460	0.471	0.137	0.101
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	5.986	5.986	−0.086	−0.086	3.831	3.831	0.169	0.169	0.014	0.014

Notes: OLS estimates with strata fixed effects. Heteroskedasticity-robust (HC1) standard errors in parentheses; exact p-values in brackets. All outcomes measured at two-month follow-up ($N = 231$). Odd columns (Compl.) use the Course Completed indicator; even columns (Started) use the Course Started indicator (completed or in progress). Course Completed and Course Started are endogenous; these estimates are descriptive, not causal. The completion premium on PESI (+0.725, $p = 0.008$, col 9) and AI usage frequency (+0.595, $p = 0.096$, col 5) are consistent with sustained hands-on engagement building firm-specific capabilities. Belief Index (cols 3, 4) is an Anderson (2008) inverse-variance-weighted composite; control-group mean of −0.086 reflects standardization relative to the full-sample mean. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We report results from Experiment 2, in which approximately half of all baseline participants were invited to a free online AI training course. Course invitation was randomized independently of Experiment 1 treatment assignment. Of the 240 participants in our analysis sample, 21 completed the full course (8.8%) and 77 started without completing (32.1%). The low completion rate limits the statistical power of these estimates; we therefore present the course results as suggestive evidence on the patterns predicted by the framework, without making causal claims about the effect of course completion on outcomes. Figure B1 shows the distribution of PESI scores by treatment arm and course completion status.

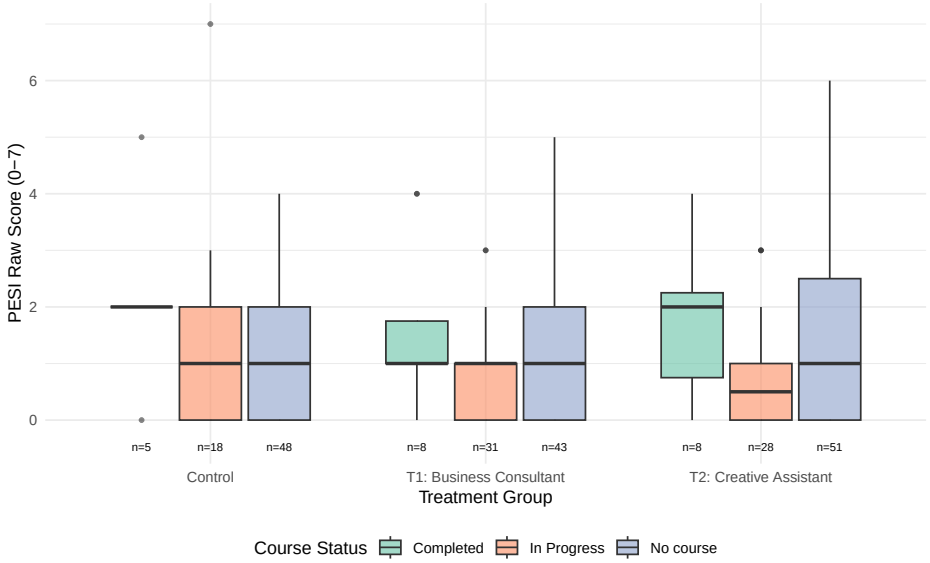


Figure B1: Prompt Engineering Sophistication (PESI) by Treatment Arm and Course Completion ($N = 240$). The figure shows the distribution of PESI scores (scale 0–7) across Experiment 1 treatment arms and course completion status. Course completers exhibit higher prompt sophistication regardless of Experiment 1 treatment assignment, consistent with sustained engagement building capabilities that passive information does not. Course completion is endogenous; this figure is descriptive only.