Skeptical but Swayed: AI, Experts, and Household Economic Expectations

William Bazley Yosef Bonaparte Ryan Clark Carina Cuculiza*

June 22, 2026

Abstract

We experimentally examine whether artificial intelligence-attributed information about aggregate risk influences households' economic expectations and behavior. Randomly assigning recession forecasts attributed to an artificial intelligence (AI) model or a professional economist, we exogenously shift individuals' recession beliefs. Individuals revise their expectations in response to the AI-attributed forecast at a rate statistically indistinguishable from that for an identical professional forecast, even though they judge it as less accurate and less trustworthy. They extrapolate these beliefs to their inflation and labor market expectations, and we find evidence that their revised recession beliefs affect their stock trading decisions, particularly deterring market entry among non-investors. Our findings indicate that households incorporate AI-attributed macroeconomic information into their beliefs to a degree comparable to expert forecasts, even though they consider it to be less credible.

Keywords: Artificial Intelligence (AI), Stock Market Participation, Survey of Professional Forecasters, Subjective Beliefs.

JEL classification: D12, D14, D83, D84, D90, G50.

*Bazley (wbazley@ku.edu) and Cuculiza (carina.cuculiza@ku.edu) are at the University of Kansas, School of Business, 1654 Naismith Dr., Lawrence, KS 66045. Bonaparte (yosef.bonaparte@ucdenver.edu) is at the University of Colorado at Denver, 1475 Lawrence Street, Denver, CO 80202. Clark (rclark23@unl.edu) is at University of Nebraska–Lincoln, College of Business, 730 N 14th St., Lincoln, NE 68588. We appreciate helpful feedback from Elham Tabatabaei and participants at the 2026 Annual Meeting of the Southwesten Finance Association. We thank Jaroslav Cuculiza for excellent research assistance.

1. Introduction

Households' expectations about economic conditions shape their consumption, savings, and investment decisions. Their economic concerns, such as those about future recessions and unemployment, influence investment flows and move asset prices (Da, Engelberg, and Gao, 2015), and can affect aggregate outcomes (Bhandari, Borovička, and Ho, 2025). Yet households' economic beliefs are routinely biased, noisy, and based on limited information (Mankiw, Reis, and Wolfers, 2003; D'Acunto, Malmendier, and Weber, 2023).[1] Models of imperfect information attribute these patterns to infrequent exposure to rational macroeconomic signals or to difficulties interpreting them (Carroll, 2003; Sims, 2003). Accordingly, expanding households' access to credible and interpretable macroeconomic information should reduce belief distortions and facilitate the transmission of aggregate signals.

Advances in artificial intelligence (AI) are creating a new channel through which economic information can reach households. AI-powered tools are increasingly integrating into individuals' lives: more than half of U.S. adults report interacting with AI technologies at least several times per week (Kennedy, Tyson, and Saks, 2023), and nearly 40% use AI tools for help with financial decisions (Cottrill and Affairs, 2024). Unlike professional forecasts disseminated through conventional news media or surveys, AI-mediated economic information can be delivered to individuals directly, through channels distinct from traditional media. AI providers have also begun to proactively deliver news content to users. For example, in September 2025, OpenAI launched ChatGPT Pulse, which delivers personalized daily briefings that can include economic and financial information (OpenAI, 2025).[2]

Whether AI alters the transmission of macroeconomic information ultimately depends on a fundamental question: do households incorporate AI-attributed signals into their beliefs and decisions? Establishing that they do is a necessary condition for any AI-driven

[1] See also Armantier, Bruine de Bruin, Potter, Topa, Van Der Klaauw, and Zafar (2013).

[2] Economic information providers are likewise incorporating AI features into their offerings. Bloomberg, for instance, introduced AI-mediated news and earnings summaries on the Bloomberg Terminal beginning in January 2024 (Bloomberg, 2024, 2025).

change in how economic information propagates through the economy. However, there is limited evidence on whether and how AI-attributed macroeconomic information influences households' beliefs and behavior. Using a pre-registered, two-wave randomized controlled trial experiment with 3,796 participants, we provide evidence that AI-attributed aggregate risk forecasts causally influence individuals' macroeconomic beliefs and their consumption-savings decisions.[3]

Building on the experimental approach of Roth and Wohlfart (2020) and Coibion, Georgarakos, Gorodnichenko, Kenny, and Weber (2024), participants first report their prior belief about the probability of a U.S. recession in the first quarter of 2027. They are then randomly assigned to one of several key experimental conditions, which determines the recession probability forecast they receive: one attributed to a professional economist from the Survey of Professional Forecasters (SPF), an economically identical forecast attributed to an AI model, a more pessimistic AI-attributed forecast, or no forecast. Having the professional and the equivalent AI forecast both place the probability of a recession at 20% is the key design feature of our experiment. This structure allows us to test for a messenger effect, i.e., whether individuals respond differently to identical forecasts depending on whether the source is human or AI. The second AI forecast increases the recession probability to 32%, allowing us to assess whether individuals' responsiveness to AI-attributed information varies with the magnitude of the signal. Participants who receive no forecast serve as the control group. All participants then re-forecast the likelihood of a recession and report their expectations about a broad set of additional macroeconomic and personal economic outcomes, together with their consumption intentions and subsequent realized behaviors.

We assess participants' response to their aggregate risk signal using a Bayesian learning framework, where any adjustment between their posterior and initial beliefs reflects their degree of learning from the signal. Prior research offers mixed evidence on how individuals

[3]The experiment is pre-registered with the University of Pennsylvania's Wharton Credibility Lab (AsPredicted #280883 and #284669). The University of Kansas Institutional Review Board (IRB) approved the research methods used in the study (IRB Protocol #00151988).

are likely to respond to AI-attributed economic information. Studies documenting algorithm aversion find that individuals are reluctant to act on information from algorithmic or technological sources, including in financial decision-making contexts (Dietvorst, Simmons, and Massey, 2015; Greig, Ramadorai, Rossi, Utkus, and Walther, 2023). In contrast, other research shows that households respond to technology-driven financial information (D'Acunto, Prabhala, and Rossi, 2019; Capponi, Olafsson, and Zariphopoulou, 2022), and recent evidence documents that AI-generated forecasts can match or exceed the accuracy of professional forecasters (Hansen, Horton, Kazinnik, Puzzello, and Zarifhonarvar, 2024). The net effect on belief updating is therefore difficult to sign ex ante. Perceptions of AI's capacity to process economic data may enhance its credibility and increase households' responsiveness to AI-attributed forecasts, whereas persistent skepticism toward algorithmic sources may lead households to discount them relative to forecasts from professional economists. By randomly varying the source of an otherwise identical forecast, our experimental design allows us to assess how these competing forces shape macroeconomic belief formation.

We find that individuals initially hold pessimistic beliefs about future economic conditions, with the average participant placing the likelihood of a U.S. recession in the first quarter of 2027 at approximately 46%. In response to their aggregate risk forecast, participants revise their posterior recession beliefs toward the signal, with an average learning rate of about 24%. Notably, learning rates are statistically indistinguishable between participants who receive the 20% recession forecasts from the AI model and professional. This equivalence, however, masks a dissociation between how households judge AI-attributed information and how they act on it. Participants rate the AI forecast as significantly less accurate and less trustworthy than the economically identical professional forecast, yet revise their beliefs in response to it to a statistically equivalent degree.

To validate our findings, we conduct several robustness tests. First, we address the concern that participants' belief updating is driven by numerical anchoring on the stated probability in their aggregate risk signal. We find that providing participants with a placebo information provision, which incorporates a numeric value of 20%, leads to belief updating at less than half the rate among those who received a genuine recession forecast. Second, we examine whether participants' belief revisions are likely to be driven by effort to conform with perceived experimenter expectations. Using the approach of Roth and Wohlfart (2020), we find that providing participants with an information provision stating a directional hypothesis with respect to belief revisions leads to updating at rates statistically indistinguishable from those of control participants. Third, to address the concern that participants were inattentive to the source of their forecast, we find that about three-quarters of participants correctly identify the source of their aggregate risk signal on a post-experiment recall question. Restricting our empirical tests to these individuals leads to results that are economically and statistically similar to those based on the full sample. Collectively, the evidence indicates that numerical anchoring, experimenter demand effects, and inattention to the forecast source are unlikely to explain our primary findings.

We next look beneath the population-level equivalence to examine what influences the weight individuals place on their recession signal. We find that belief revision is responsive to the perceived credibility of the forecast, measured along the dimensions of accuracy and trustworthiness. The credibility–updating slopes for the AI forecasts are at least as steep as for the professional forecast, indicating that households do not categorically discount AI-attributed recession signals. Responsiveness also varies systematically across individuals. More frequent users of AI tend to update more strongly from the AI signal, while participants expressing greater trust in the Federal Reserve update more strongly from the professional forecast. Because we elicit these characteristics after the information treatment, we interpret them as descriptive patterns of heterogeneity rather than identified mechanisms. Collectively, this evidence indicates that the population-level equivalence conceals source-specific differences in how individuals respond to economic information from AI and professional sources.

Recession beliefs formed from conventional sources of information, e.g., professionals' forecasts, are known to propagate to other dimensions of households' economic outlook (Roth and Wohlfart, 2020). Establishing whether AI-influenced beliefs propagate beyond the initial recession forecast bolsters the economic meaningfulness of our source equivalence evidence, as it suggests that beliefs shaped by an AI signal carry the same downstream implications as beliefs shaped by a professional forecast. We find that AI-influenced recession beliefs have downstream consequences. Higher posterior recession beliefs increase individuals' expectations about future unemployment, at both the national and county-of-residence levels, and induce expectations of higher inflation, consistent with a stagflationary view of the economy. The evidence on microeconomic spillovers is mixed. We find no significant effects on participants' perceptions of their personal financial prospects, income growth, or job loss probability. However, consistent with their pessimistic beliefs about future unemployment rates, we find that, as perceived recession likelihood rises, individuals are less optimistic about finding a new job should they become unemployed.

Since canonical models predict that households' consumption-savings decisions should respond to innovations in their economic beliefs (Carroll, 1997; McKay, Nakamura, and Steinsson, 2016; Challe and Ragot, 2016; Kuchler and Zafar, 2019), our final analysis focuses on participants' intended and realized consumption of nondurable and durable goods, as well as their stock trading activity. We find that more pessimistic recession belief updating significantly increases planned spending on necessary nondurable goods, i.e., food. Effects on discretionary spending intentions and realized nondurable expenditure are not statistically significant, suggesting that the primary behavioral response may operate through the composition of spending rather than the level. We find no evidence that recession belief updating affects participants' perceived climate for durable goods purchases or their realized durable goods consumption over the two-week follow-up period. With respect to financial market behavior, higher recession beliefs significantly deter stock market participation among non-investors, as they become less likely to initiate stock positions as their recession expectations rise. Finally, the evidence indicates that AI- and professionally-attributed recession signals produce statistically equivalent downstream behavioral responses. Overall, the evidence highlights that, by altering individuals' behavior, AI-attributed macroeconomic belief updating has economic implications.

Our paper contributes to several strands of literature. First, we add to a body of work on how households form and revise macroeconomic expectations and on the use of information-provision experiments to study them (Enders, Trautmann, and Vinogradov, 2026; Coibion and Gorodnichenko, 2015; D'Acunto and Weber, 2023). This literature has shown that households update their beliefs in response to professional and institutional forecasts (Roth and Wohlfart, 2020; Coibion, Gorodnichenko, and Weber, 2022; Coibion et al., 2024), but that the sporadic transmission of such information through traditional channels can sustain noisy and biased beliefs driven by local information signals (Carroll, 2003; D'Acunto, Malmendier, Ospina, and Weber, 2021). Initial evidence suggests that the precise wording of a message matters less than its informational content (Coibion et al., 2022). Whether the message's source, as opposed to its wording, affects belief updating is a distinct question, particularly for algorithmically-attributed information, where perceived credibility may diverge from that of professional sources. We extend this literature by examining a new and rapidly proliferating source of economic information: generative AI. Our central finding, that individuals revise their recession expectations as strongly from an AI-attributed forecast as from an economically identical professional forecast, shows that the belief-formation mechanism documented in this line of research extends to algorithmically-attributed information.

Second, we contribute to a growing literature on the role of AI and financial technology in household decision-making. Much of this literature studies robo-advising and algorithmic delegation, documenting that technology-driven tools can debias choices and improve portfolio outcomes (D'Acunto et al., 2019; D'Acunto and Rossi, 2023a,b). For instance, Capponi et al. (2022) show that algorithmic advice reshapes investors' portfolio choices by supplying less biased expectations about equities, and Choukhmane, de Silva, Lin, and Akuzawa (2026) examine the supply of and demand for AI-based financial advice over the life cycle.

An emerging strand of research asks how generative AI affects the way individuals process financial information (Blankespoor, Croom, and Grant, 2024). Our study expands this nascent literature. While prior research largely focuses on AI tools that execute or recommend financial actions, we study AI as a source of macroeconomic information that shapes the beliefs underlying households' consumption-savings actions. By tracing AI-influenced recession beliefs through to expectations across economic domains and to downstream consumption and investment behavior, we show that AI-attributed information shapes behavior by moving a fundamental driver of economic decisions.

Third, our findings speak to the evolving literature on algorithm aversion and appreciation (Jussupow, Benbasat, and Heinzl, 2024). Prior research documents that individuals discount advice and information from algorithmic sources, even when it is accurate. That is, people tend to resist algorithms after seeing them err (Dietvorst et al., 2015), and they can be reluctant to rely on AI even when it matches or exceeds human performance (Longoni, Bonezzi, and Morewedge, 2019). These behavioral patterns can extend to financial settings (Burton, Stein, and Jensen, 2020; Greig et al., 2023; Stradi and Verdickt, 2025). However, another strand of research identifies algorithm appreciation, i.e., a tendency to weight algorithmic advice more heavily than human advice (Logg, Minson, and Moore, 2019). Our setting offers a novel test of these competing forces in the domain of macroeconomic information, and our evidence indicates that the two operate on different margins. Consistent with algorithm aversion, households rate an AI-attributed recession forecast as significantly less accurate and less trustworthy than an identical professional forecast. This aversion, however, arises in how households assess the signal rather than in how they respond to it. Overall, in the domain of macroeconomic news, algorithm aversion appears to operate on households' assessments of AI-attributed information rather than on the weight they place on it when forming beliefs.

2. Experimental Design

Since traditional datasets of households' expectations and economic behaviors do not include the necessary elements to examine our research questions, we conduct an experiment using the Prolific online platform. The platform provides access to a heterogeneous pool of high-quality participants (Douglas, Ewell, and Brauer, 2023). To further ensure data quality, we require participants to have at least a 90% successful completion rating on their prior tasks. Moreover, we use detection functionality in our Qualtrics survey tool to identify and remove responses that have a high likelihood of being submitted by bots.[4] As a further quality check, we incorporate a question in the experiment wherein participants are directed to select a particular answer option in order to assess their diligence during the task. We construct an indicator, *Attentive*, that equals one if a participant correctly answers the question, and zero otherwise. We use it as a control variable in our empirical tests.

2.1. *Prior Belief Elicitation*

Upon agreeing to participate in the first wave of our experiment, individuals respond to a series of questions which elicit their demographic characteristics and attitudes toward artificial intelligence.[5] After completing the preliminary information questions, participants receive a set of instructions that characterizes how they should express their economic expectations using probabilities. We also describe our key economic concept, a recession, as a quarter in which the economy produces less goods and services than in the preceding quarter, i.e., a fall in real GDP.[6] Subsequently, we ask all participants to forecast the likelihood of a recession in the U.S. during the first quarter of 2027. Since we conducted our experiment in

[4]The bot filter results in omitting 37 responses. Our analysis yields similar results when including these observations.

[5]We focus on demographic characteristics known to affect expectations and financial decision-making (e.g., Campbell (2006); Gomes, Haliassos, and Ramadorai (2021)), such as participants' gender, age, household income, education level, employment status, political preference, marital status, and race. We define all variables in Appendix I.

[6]Following Roth and Wohlfart (2020), we adopt this definition of a recession to align participants' forecast horizon with that of the Survey of Professional Forecasters.

March of 2026, participants are forecasting economic conditions approximately ten to twelve months into the future.

2.2. Information Provisions

Following the elicitation of participants' priors, they are randomly assigned to one of six conditions, which determines the information provision they receive. Below, we discuss features of the information provisions for our primary experimental conditions. We characterize the features of the robustness conditions, i.e., the Placebo and Demand conditions, in Section 3.1.1.

Professional$_{20}$. Individuals in this condition receive a forecast submitted to the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF). The SPF routinely collects predictions from professional forecasters related to a variety of economic elements, including the likelihood of a decline in real GDP in the quarter of the survey and in the four subsequent quarters. We characterize the forecaster as knowledgeable about current macroeconomic data and economic theories, and present a forecast of a 20% recession likelihood for the first quarter of 2027 to participants assigned to this condition.[7] Additionally, we contrast each participant's prior belief with the prediction of the forecaster using a bar chart.[8] The advantage of including this professional condition in the experimental design is that it can serve as a useful reference point against which we can assess AI conditions, since prior research establishes that households update their beliefs in response to receiving recession forecasts from professionals (Roth and Wohlfart, 2020).

AI$_{20}$ and AI$_{32}$. Participants in these conditions receive a recession probability prediction attributed to a publicly available AI model. To generate the forecasts, we provide the model with a prompt that: (i) describes a recession as a fall in real GDP, which is reflected in

[7]This forecast is close to the average expectation, which is about 24%, among professionals who submit predictions to the SPF during the particular survey wave.

[8]Examples of the information provision screens are shown in Figures IA1 and IA2.

a decline of real income, employment, production, and sales; (ii) instructs the model to consider relevant macroeconomic theories and currently available data; and (iii) asks it to provide its single best numeric forecast of the probability there will be a fall in real GDP in the United States during the first quarter of 2027.[9] We present the forecast to participants by stating it reflects the recession forecast of an artificial intelligence model with access to current macroeconomic data and relevant economic theories. To avoid activating potential biases or participants' preferences with respect to particular AI models, we do not inform them which specific AI model provides the forecast.

Control. Participants in this condition do not receive a recession forecast. Rather, we present them with a screen which directs them to click the button to proceed. Accordingly, participants in this condition serve as a no-information counterfactual group.

2.3. Posterior Beliefs Elicitation

Following the information treatment, we elicit treated participants' degree of trust in and perceived accuracy with respect to the recession forecast signal they receive. We then ask a series of questions designed to measure a variety of expectations with respect to macroeconomic outcomes. In particular, we ask participants for their expectations about national unemployment in the U.S. over the next twelve months using both point estimate and categorical response questions. Participants also provide their beliefs with respect to the trend in unemployment in their county of residence.

We also elicit density forecasts of inflation over the next twelve months by having participants assign probabilities to ten brackets of potential inflation rates. This approach allows us to calculate measures of each participant's mean inflation rate expectation and inflation rate uncertainty (i.e., standard deviation). The last dimension of aggregate beliefs that we elicit is participants' outlook with respect to firms' profits. Specifically, we ask them to re-

[9] The specific prompt used to generate the recession forecasts is shown in Section IA1. We obtained the forecasts from Claude and Copilot in March of 2026.

spond to the question: Over the next twelve months, do you think the profitability of firms in the U.S. economy will increase or decrease? Participants respond on a five-point Likert scale that ranges from strongly decrease to strongly increase. We use their response to create *Firms' Profits*, which is increasing in expected future profitability.

Next, participants answer several questions designed to elicit their expectations about their personal circumstances in the future. We start by asking all participants to respond to the question: Looking ahead, do you think you (and any family living with you) will be financially better or worse off twelve months from now than you are these days? Individuals respond on a seven-point Likert scale that ranges from much worse off to much better off. We construct the variable *Financial Prospects*, based on their responses, where higher values indicate more optimistic expectations about future financial well-being.

We then elicit participants' beliefs about their personal employment prospects. Employed participants are asked to provide density forecasts of their expected earnings growth over the next twelve months, under the assumption that they will remain in their main job. Using their forecasts, we calculate each individual's expected income growth and income growth uncertainty. Participants also report the perceived probability, in percent terms, that they will lose their main job within the next year. We construct *Personal Unemployment* to capture these responses, where higher values indicate a greater perceived likelihood of job loss. The final element in this module asks individuals to provide the likelihood they would find an acceptable replacement job if they happened to lose their main position.

2.4. Consumption Intentions

Subsequently, participants provide their consumption intentions. With respect to nondurable consumption, they indicate their spending plans for necessities (food at home and in restaurants) and discretionary (leisure activities and clothing) items over the next two weeks on five-point Likert scales ranging from spend much less to spend much more, relative to the past two weeks. Participants also estimate the dollar value of spending on nondurable consumption goods over the week before the survey. With respect to durable consumption, we elicit individuals' perceived climate for purchasing durable goods.

Following the battery of macro- and microeconomic questions, participants again provide a point forecast related to a recession in the U.S. in the first quarter of 2027. The concluding section of the survey includes questions related to: the frequency with which the individual uses an AI system during a typical day, financial literacy (Lusardi and Mitchell, 2008, 2014), whether they follow news about the national economy, if they own stocks, and to what extent they trust the Federal Reserve. Finally, we ask participants if they remember the source of their information provision. For completing the experiment, participants are compensated with a fixed payment of \$2.75 and a variable payment of \$0.25 for answering the attention check question correctly. Based on the median participant's time to complete the experiment, the compensation rate is about \$13.85 per hour.

2.5. Design Rationale

Two features of the experimental design are important to our identification strategy. First, both the Professional_{20} and AI_{20} conditions provide 20% recession probability signals with the only difference being the source of the forecast. By comparing both groups to Control participants, we can attribute any differences in belief updating or downstream behavioral responses to a pure messenger effect. That is, any differences in posterior beliefs and behaviors are driven by the source label of the information rather than the content. Accordingly, Control participants serve two roles in our empirical tests. They provide the no-information benchmark for estimating causal effects within each experimental condition and serve as the common baseline against which potential variation in the processing of AI-attributed versus human expert information signals are assessed.

Second, the AI_{32} condition extends the experimental design by including a recession likelihood signal of 32%. By including this condition, we can examine whether participants' responsiveness to AI-attributed information varies with the magnitude of the signal, holding the source constant. Additionally, when incorporating it in empirical tests with the Professional_{20} condition, we can assess whether any source-specific heterogeneity in responsiveness varies across levels of recession signals.

2.6. Participants' Characteristics

We recruit 3,796 participants for our baseline experiment. Of these, 2,908 are randomly assigned to our primary experimental conditions and constitute our main analysis sample. The remaining participants are assigned to the Placebo and Demand conditions used in our robustness tests. Table 1 reports descriptive statistics for our main analysis sample. About 48% of participants identify as male and the average participant is in their mid-30s. The typical participant reports a gross household income between $50,000 and $75,000 in 2025 and holds a bachelor's degree. Approximately 76% identify as White and 41% are married. These characteristics broadly align with the U.S. population in terms of gender, race, and marital status, though our participants are somewhat more educated and moderately younger. Moreover, our typical participant's household income is slightly below the national median reported in the Current Population Survey Annual Social and Economic Supplement conducted by the Census Bureau.

Our experimental conditions are well balanced based on normalized differences relative to the Control group (Imbens and Wooldridge, 2009; Imbens and Rubin, 2015). That is, the absolute normalized differences, reported in brackets, are below conventional threshold levels and no systematic imbalances are evident. Importantly, participants' prior recession beliefs are closely balanced across all arms, with normalized differences below 0.04.

2.7. Follow-Up Survey

Two weeks after the conclusion of our baseline experiment, we conduct a follow-up survey of our participants. Specifically, we post a survey advertisement on Prolific that is accessible to individuals who successfully completed the initial experiment and were assigned to one of the primary treatment conditions (i.e., Professional$_{20}$, AI$_{20}$, or AI$_{32}$) or the Control group. About 1,291 individuals complete the follow-up questionnaire, representing about 44% of initial participants in the primary experimental conditions. Participation is well-balanced across the experimental arms, with 325 individuals in AI$_{20}$, 327 in AI$_{32}$, 322 in Professional$_{20}$, and 317 in the Control condition.

Our follow-up survey serves two purposes. First, following the obfuscated-recontact approach recommended by Haaland, Roth, and Wohlfart (2023), it serves as a safeguard against experimenter demand effects. While our baseline experiment minimizes such concerns through its design features, e.g., eliciting beliefs using both probabilistic and qualitative questions, embedding the key questions among a battery of unrelated items, and showing no discernible effects among Demand condition participants, the follow-up survey administers no information treatment and is separated from the original forecast by two weeks. Any belief differences that persist across this interval of time are unlikely due to demand effects, as they typically manifest within the immediate experimental context. Second, the follow-up survey allows us to examine the persistence of both the professional and AI information treatments on participants' beliefs and behaviors. Moreover, we can assess whether the durability of any shifts in beliefs differ by the source of the forecast.

Upon accessing the follow-up survey, individuals provide their Prolific ID, which enables us to link their responses to their answers in the baseline experiment. Participants then provide information related to some of the central elements of our study, such as the likelihood that national and county-level unemployment rates increase, and their expectations about firms' profitability, inflation, job outcomes, and personal financial prospects. We also re-elicit individuals' expectations about the probability of a recession in the U.S. in the first quarter of 2027. Finally, we collect information on participants' consumption and financial behaviors during the two week interval. Specifically, individuals provide their combined spending on nondurable consumption during the week before the follow-up survey, whether they purchased nondurable goods, and whether they actively increased or decreased their stock holdings in the prior two weeks.

3. Initial Recession Expectations and Belief Updating

Prior to receiving an information treatment, the average participant places the likelihood of a U.S. recession in the first quarter of 2027 at approximately 46%. Consistent with successful randomization, pre-treatment recession beliefs are similar across experimental conditions. Specifically, the mean values are 45.65%, 45.85%, 45.99%, and 46.42% for the Professional$_{20}$, AI$_{20}$, AI$_{32}$, and Control conditions, respectively.

3.1. Belief Updating in Response to Recession Forecasts

Consistent with models of incomplete information and inattention to macroeconomic news (e.g., Carroll (2003)), if households are not fully informed about the economy, then their beliefs should adjust in response to receiving relevant information signals. Prior research documents that households consider professionals' expectations, including those contained in the SPF, to be informative signals about the macroeconomy (Carroll, 2003; Roth and Wohlfart, 2020). Additionally, recent research by Hansen et al. (2024) suggests that forecasts generated by large language models are similar in terms of accuracy to those of human experts. Therefore, we expect participants to revise their recession expectations towards their aggregate risk signal if they are not fully informed prior to the experiment.

Figure 1 provides initial graphical evidence that is consistent with our hypothesis. The figure shows, for each of our primary experimental conditions, the distribution of participants' prior and posterior recession beliefs. Across all three treatment conditions, participants' beliefs visibly shift toward their forecast signal, indicating that participants were not fully informed about their signal's content prior to the experiment and that it was relevant to their macroeconomic outlook. In contrast, participants in the Control condition show no discernible systematic change between their prior and posterior recession expectations.

To measure participants' degree of responsiveness to their recession signal, we follow Roth and Wohlfart (2020) and assume that individuals rely on a Bayesian learning framework when interpreting information signals. The framework is discussed in Section IA2 of the Internet Appendix. The key benefit of the framework is that it allows us to quantify an individual's rate of learning in response to receiving their particular aggregate risk signal as the difference between their posterior recession forecast and their initial forecast. We term this key element *Belief Revision*.

We also create the variable *Shock*, which is the difference between the information signal, i.e., the forecasted likelihood of a recession contained in the information provision, and the participant's prior belief. As in Roth and Wohlfart (2020), individual i is expected to apply a linear learning rule that is approximated by: $Belief\ Revision_i = \alpha_1 Shock_i$. The learning rate is captured by α_1, which is bounded between zero and one.

To estimate α_1, we conduct ordinary least squares regressions specified as:

$$Belief\ Revision_i = \alpha_0 + \alpha_1 Shock_i + \theta' \mathbf{X}_i + \varepsilon_i. \tag{1}$$

In our baseline regressions, we include a vector of additional variables, $\mathbf{X}$, to account for heterogeneity in individuals' socioeconomic characteristics and increase the precision with which we estimate treatment effects. While the learning rule expresses belief revision as strictly proportional to the shock, our estimating equation additionally includes an intercept, α_0, and indicators for the participant's experimental condition, i.e., Professional_{20}, AI_{20}, and AI_{32}, with Control participants serving as the reference group. These condition indicators absorb any baseline differences in belief revision across arms that arise independently of the information signal. Accordingly, α_1 captures the updating response to the shock conditional on any baseline differences. Finally, we adjust the standard errors for heteroskedasticity (White, 1980).

We report the empirical results in Table 2. Column (1) shows that the univariate estimate of the learning rate is highly statistically significant and corresponds to about 24% of the recession shock to participants' beliefs. In Column (2), including controls to account for variation in participants' age, household income, gender, education, political affiliation, marital status, race, and financial literacy level has minimal impact on the estimated rate. In Column (3), we incorporate our control variable that captures the participant's diligence during the experimental task. We find that the estimate on *Shock* remains positive at 0.235 and is strongly statistically significant (p-value < 0.001). With respect to the magnitude of the effect, our participants display incomplete learning behavior, as they do not revise their posterior beliefs to fully align with their aggregate risk signal. This behavior is consistent with prior research that documents similar incomplete learning among households in response to economic information (Cavallo, Cruces, and Perez-Truglia, 2017; Armantier, Nelson, Topa, Van der Klaauw, and Zafar, 2016). For example, Roth and Wohlfart (2020) find that households update their recession beliefs in response to professionals' recession forecasts by about one-third of the shock.

Next, we examine whether participants' changes in beliefs are consistent with Bayesian updating. Bayesian updating predicts that individuals who are more confident in their initial beliefs should update less. In Column (4), we find that participants who report higher confidence in their initial recession beliefs update less in response to their aggregate risk signal. In Column (5), we assess whether prior exposure to economic news reduces participants' belief updating. We find that individuals who are initially informed about the state of the national economy place lower weight on their recession signal. Overall, the estimates imply that our participants were not, ex ante, fully informed about the aggregate risk signals, that the signals contained some relevant information for their recession expectations, and that participants updated their beliefs in a manner consistent with a Bayesian learning process.

3.1.1. Alternative Explanations

We take steps to address two common concerns related to information provision experiments. First, to test whether the observed belief updating reflects genuine economic information processing rather than numerical anchoring on the signal value, we incorporate a placebo information provision. Participants assigned to the Placebo condition receive a statistic from the U.S. Centers for Disease Control and Prevention (CDC) related to tobacco usage among adults, rather than an economic forecast. Specifically, they are informed that the CDC found that 20% of U.S. adults use tobacco products (CDC, 2024).

While the domain of the information is unrelated to macroeconomic conditions, the signal value, i.e., 20%, is identical to that of the Professional_{20} and AI_{20} recession forecasts. Moreover, the information is presented in the same visual format as the recession signals. If Placebo group participants update their recession beliefs, in response to the health statistic, at a rate comparable to those in the primary treatment conditions, it would suggest that treatment participants' updating is driven by numerical anchoring on the stated probability rather than by the economic content of their recession forecast. Conversely, significantly attenuated updating among Placebo participants provides evidence that is consistent with the belief revision among our key treatment participants reflecting substantive economic information processing.

The coefficients in Columns (1) and (2) of Table IA1 show that participants who received the tobacco use statistic place significantly less weight on the information when forming their posterior recession beliefs. Specifically, the estimates from the fully specified model indicate that the average Placebo participant's learning rate is about 0.093, which is approximately 40% of the genuine updating rate of 0.235 among key treatment participants. The presence of a small but statistically significant placebo effect is consistent with a long-established literature documenting numerical anchoring in judgment under uncertainty (Tversky and Kahneman, 1974). Importantly, this anchoring accounts for less than half of the updating observed in the genuine forecast conditions. Accordingly, this attenuation suggests that the belief revision identified among participants who receive AI or professional aggregate risk forecasts reflects substantive economic information processing rather than numerical anchoring alone.[10]

Second, we examine the extent to which our participants are responsive to experimenter demand effects by deliberately trying to influence their belief responses. Specifically, we incorporate a Demand condition in our experimental design. Individuals in this condition receive an information provision, adopted from Roth and Wohlfart (2020), to identify the extent to which social desirability bias influences participants' responses to our survey questions. The provision states: "In this experiment, people are randomly assigned to receive different instructions. We hypothesize that participants who are shown the same instructions as you report more optimistic expectations about the US economy." The narrative does not contain any economic information or forecast. Instead, it suggests a direction for belief revisions. If participants in the key treatment arms primarily update their beliefs to align with perceptions of our expectations, rather than the economic information content of their recession forecast, then this demand provision should produce belief updating among participants in this condition.

The estimates in Columns (3) and (4) of Table IA1 show that participants in the Demand condition behave similarly to individuals in the Control condition. That is, the results suggest that our key treatment participants' economic beliefs are unlikely to be sensitive to overt signals about our key research hypotheses.

[10]Additionally, numerical anchoring on the shared 20% value cannot explain the equivalence in learning rates we document across the AI and professional conditions. If that equivalence merely reflected a common anchor, the AI learning rate should diverge from the rate observed in the Professional_{20} condition once the AI signal's value changes. As we show in Section 3.2, it does not. The learning rate in the AI_{32} condition remains statistically indistinguishable from that in the Professional_{20} condition. Because the learning rates remain stable and track the professional benchmark even under a different potential numerical anchor, the evidence indicates that individuals are responding to the economic content of their recession forecast rather than merely anchoring on its specific integer.

3.2. Belief Updating Across AI and Professional Signals

Next, we examine whether participants update differently depending on the source of their aggregate risk signal. How households weight AI-attributed economic information is theoretically ambiguous. Algorithm aversion predicts that individuals discount information from algorithmic sources and would therefore update less from an AI forecast than from an identical professional one (Dietvorst et al., 2015). Algorithm appreciation predicts the opposite, i.e., that individuals weight algorithmic input more heavily (Logg et al., 2019). A third possibility is that households treat the two sources as informationally equivalent and assign similar weight to each. Ultimately, whether households update from AI-attributed forecasts, and whether they do so to a degree comparable to professional forecasts, determines whether AI can serve as a viable channel for the transmission of macroeconomic news.

We test whether the source of an economic forecast differentially influences households' belief updating by expanding Equation (1) to include interaction terms between *Shock* and indicators for each of our three primary treatment conditions. This specification provides condition-specific learning rates relative to benchmark participants in the Control group. Table 3 reports the results.

Our primary test of the messenger effect compares the learning rates of the AI_{20} and Professional_{20} condition participants, as both receive 20% recession probability signals which differ only in their source. Specifically, we conduct a linear combination test of the two learning rate coefficients using the full variance-covariance matrix of the regression. The estimated differences in learning rates across the regression specifications have an economically negligible range of -0.018 to -0.017 (standard errors $\approx$ 0.039–0.040) and are not statistically different from zero. To assess whether the null reflects an absence of a meaningful source effect rather than insufficient statistical power, we apply a two one-sided tests (TOST) procedure for equivalence testing (Lakens, Scheel, and Isager, 2018). With an equivalence margin set at 35% of the baseline learning rate, corresponding to an absolute difference of 0.08, the procedure rejects source-driven differences larger than this bound in either direction (p-value

= 0.049). We therefore find no detectable difference in learning rates between the AI and professional conditions within this testing framework.

Comparisons among the remaining experimental conditions reinforce this pattern. The learning rates associated with the AI_{32} and $\mathrm{Professional}_{20}$ conditions are not statistically different, nor are the rates between the AI_{32} and AI_{20} conditions. Because the latter comparison holds the source constant while varying the magnitude of the signal, it indicates that learning rates for AI-attributed economic news do not, on average, differ significantly with signal magnitude. This result suggests that households apply a similar proportional weight to AI forecasts whether they are more or less optimistic. Collectively, these findings indicate that the average individual updates their recession beliefs similarly regardless of whether the forecast is from an AI model or a professional, and regardless of the magnitude of the information signal.

3.2.1. Alternative Specification

For robustness, we re-estimate our learning rate tests using the approach of D'Acunto, Gao, Liu, Lu, Wang, and Yang (2024). Specifically, we regress posterior recession beliefs on interactions terms between experimental condition indicators and participants' priors. Under Bayesian learning, treated participants should place less weight on their priors when forming posteriors in response to receiving their aggregate risk signal, corresponding to interaction coefficients falling in the interval $[-1, 0]$.

Table IA2 reports the results. Across all treated conditions, the estimated prior-interaction coefficients are negative and highly significant, confirming that treated participants reduce reliance on their prior beliefs when forming their posterior recession expectations in response to receiving their aggregate risk signal. Notably, both the univariate and multivariate specifications yield learning rate differences between AI_{20} and $\mathrm{Professional}_{20}$ condition participants that are statistically indistinguishable from zero. Collectively, the evidence suggests individuals update their recession beliefs in response to their aggregate risk signal and reduce their reliance on their prior beliefs by statistically equivalent amounts regardless of whether the recession signal is from an AI model or a professional.

3.2.2. Forecast Source Encoding

A potential concern is that participants may have failed to notice the source framing of their recession forecast signal. If source label information was not encoded, our baseline evidence would reflect ignorance of the source rather than genuine equivalence in its treatment effects. To examine this concern, we include a question at the end of the survey that states: Earlier in this survey, you may have received some information. What was the source of that information? Answer options are designed to be relevant for both treatment and control participants, and include: A professional economist, An artificial intelligence model, A public health agency, A general statement about the survey, I did not receive any specific information, and I do not remember. Conditional on the participant's assigned experimental condition, we create an indicator that equals one if the individual correctly recalls the source of their aggregate risk signal, and zero otherwise.

Panel A of Figure 2 shows, across experimental conditions, the average rate of correct encoding of the information source. Active treatment participants display high rates of source encoding. About 77% of participants in the AI_{32} condition and 76% in the AI_{20} condition correctly recall the source of their recession forecast signal. Among individuals in the Professional_{20} condition, the encoding rate is 70%. Participants in the Control, Placebo, and Demand display encoding rates that are lower than active treatment participants, but substantially above random guessing.[11]

In Table IA3, we re-estimate our empirical tests but restrict the sample to respondents who correctly identified their information source. Under this restriction, the learning rate dif-

[11]The lower encoding rate is partially attributable to our strict definition of the correct answer choice associated with each experimental condition. Some answer options may be perceived as ambiguous by participants, particularly in the Control condition where no forecast was received. For instance, Panel B of Figure 2 shows that expanding the correct response set for Control participants from only "I did not receive any specific information" to also include "A general statement about the survey," which some participants are likely to interpret to be an acceptable answer, results in the encoding rate increasing from 35% to 60%.

ference between AI_{20} and $Professional_{20}$ participants remains -0.017 to -0.020 across both the univariate and multivariate specifications, which is economically negligible and statistically indistinguishable from zero. The difference in learning rates between AI_{32} and $Professional_{20}$ condition participants is similarly negligible. We find a small difference in the learning rates across participants in the AI conditions, suggesting signal magnitude can influence belief responsiveness. Overall, among participants who register and retain the source framing of their recession information signal, AI-attributed and professional forecasts produce statistically equivalent belief updating.[12]

3.3. Forecast Credibility and Belief Updating

Participants' perceptions of their recession signal's credibility should influence their belief updating. If participants systematically perceive AI-attributed forecasts as less credible than those from professionals, e.g., they exhibit algorithm aversion instead of algorithm appreciation, then their diminished credibility assessments should translate into lower learning rates. We therefore examine whether participants perceive credibility differences in recession signals across sources, and if they can account for our baseline finding of statistically equivalent belief updating.

To quantify a recession signal's credibility, we focus on perceptions of its accuracy and the trustworthiness of its source. We ask participants in the treatment conditions to respond, on five-point Likert scales, to the questions: "How accurate do you think the forecast you just received was?" and "How much do you trust the source of the forecast you just received?" Based on their answers, we construct two measures, *Forecast Accuracy* and *Forecast Trust*, that are increasing in perceived accuracy and trust, respectively. We regress these measures on indicators for participants' experimental condition and the full vector of control variables, using $Professional_{20}$ condition participants as the base group. To assess whether any

[12]The matched arms also share common support on the belief shock. The distribution of *Shock* is nearly identical across the AI_{20} and $Professional_{20}$ conditions (Figure IA3), confirming that the equivalence in learning rates is not driven by differences in the distribution of the belief shock across the two arms.

credibility gap reflects a pre-existing disposition that participants bring to the survey, we additionally control for a pre-treatment measure of the participant's general attitude toward the accuracy and trustworthiness of AI models.

Columns (1) and (3) of Table 4 show that participants in the AI conditions rate their forecasts as less credible than individuals in the Professional_{20} condition. AI_{32} participants rate their forecast about 0.162 points lower on accuracy and 0.180 points lower on trust than Professional_{20} participants. The credibility deficits are larger among AI_{20} participants, about 0.355 points on accuracy and 0.385 points on trust, or roughly one-fifth of the mean of each measure. Overall, the evidence shows a meaningful credibility gap between AI- and professionally-attributed macroeconomic information.

The estimates in Columns (2) and (4) show that this credibility gap is predominantly a response to the AI label rather than a pre-existing attitude that participants bring into the experiment. Specifically, at the start of the experiment, all participants report the extent to which they think AI models provide accurate and trustworthy information. We use their responses to create *AI Accuracy* and *AI Trust*, which are increasing in their respective dimension. We find that dispositional AI attitudes strongly predict participants' stated ratings of the forecast they receive. A one-unit increase in *AI Accuracy* (*AI Trust*) is associated with a 0.255 (0.262) increase in the rated accuracy (trust) of the forecast. While including these measures as controls raises the adjusted R^2 substantially, the AI credibility discount is almost unchanged. That is, the AI_{20} accuracy discount moves from 0.355 to 0.339, and the trust discount from 0.385 to 0.371. In other words, pre-existing attitudes explain a large share of why some participants rate forecasts as more credible than others, but they explain little of the gap between AI- and professionally-attributed forecasts.

In Columns (1) and (2) of Table 5, we examine whether participants' perceptions of their aggregate risk signal's credibility moderate belief updating. We standardize the *Forecast Accuracy* and *Forecast Trust* measures to have mean zero and standard deviation one, and expand Equation (1) to include triple interactions between *Shock*, the experimental condition indicators, and the credibility measures. Consistent with credibility-based updating, the positive coefficients on the key interaction terms show that participants who rate their forecast as more accurate or trustworthy update more strongly in response to their recession signal. The credibility–updating slopes are positive and significant in every condition, and the linear combination tests show that the AI slopes are at least as steep as the corresponding Professional_{20} slope. The one slope that differs significantly is the accuracy-updating sensitivity in the AI_{20} condition, which exceeds that in the Professional_{20} condition ($\Delta = 0.068$, $p = 0.017$). Thus, among participants who receive a 20% recession likelihood forecast, recipients of the AI signal are no less sensitive and, on the accuracy dimension, more sensitive to their credibility assessment than recipients of the Professional_{20} forecast.

Collectively, these results document a robust dissociation: individuals rate AI-attributed forecasts as less credible than professional forecasts yet weight them at least as heavily when revising their beliefs. While our experimental design cannot fully isolate the cognitive determinants of this behavior, it allows us to eliminate some plausible explanations. Specifically, the behavior is unlikely to be driven by numerical anchoring, demand effects, or source inattention. We can also rule out that the credibility gap and the equivalence in updating reflect systematic differences in participants' disposition towards AI, as Table 1 shows that the *AI Accuracy* and *AI Trust* measures do not systematically vary across conditions. Random assignment thus ensures that the dissociation is not an artifact of how AI-favorable participants are distributed across experiment arms. One reading consistent with the evidence is that stated skepticism toward AI is an expressed attitude that does not bind on belief revision. However, there could potentially be other mechanisms that support this pattern, and we therefore view the determinants of the dissociation as an open question for future research.

3.4. *Individual Heterogeneity in Source-Specific Belief Updating*

Having established the population-level equivalence in updating and ruled out several confounds, we now examine which individual characteristics shape how strongly participants respond to AI- and professionally-attributed forecasts. We consider two intuitively motivated moderators, the frequency with which participants use AI systems and their trust in the Federal Reserve, alongside a broader set of demographic characteristics that we use to assess the generality of our results. For each moderator, we standardize continuous measures to mean zero and unit standard deviation and expand our baseline learning-rate regression (Equation 1) to include triple interactions between *Shock*, the experimental condition indicators, and the moderator. The coefficients of interest are the triple-interaction terms, which measure how each source's learning rate varies with the moderator. The associated pairwise linear combinations test whether these slopes differ between the AI and professional conditions. The results are reported in Table IA4.

We first consider the behavioral and institutional moderators. We find that participants who use AI systems more frequently update their beliefs more strongly in response to the AI_{32} forecast. That is, the *AI Use* slope is significantly steeper for the AI_{32} forecast than for the professional forecast. The AI_{20} and professional slopes are statistically indistinguishable. The effect is therefore specific to the AI forecast at the higher recession signal rather than to AI attribution in general. Turning to institutional trust, participants who express greater trust in the Federal Reserve update more strongly in response to the professional forecast, while the corresponding AI slopes are statistically insignificant. Together, these results indicate that belief responsiveness to AI- and professionally-attributed forecasts varies with individuals' behavioral engagement with AI and with their trust in a traditional economic institution, respectively.

We next explore whether the source-specific dissociation is consistent across demographic subgroups. To do so, we interact the source-specific learning rates with participants' age, income, gender, education, marital status, race, political affiliation, and financial literacy.

The evidence indicates that the dissociation is broadly stable. Across the demographic characteristics, the difference between AI and professional updating is statistically insignificant for the majority of the source-specific comparisons. Two exceptions reach conventional significance. Older participants update less strongly than younger participants in response to the AI_{32} forecast relative to the professional forecast, a pattern consistent with prior evidence that receptiveness to new technologies tends to decline with age (e.g., Morris and Venkatesh, 2000; Venkatesh, Morris, Davis, and Davis, 2003). Higher-income participants update somewhat more strongly in response to the professional forecast relative to the AI_{20} forecast. Though the demographic results are generally consistent with existing literature, we interpret them with caution, as the effects arise on different experimental arms and they do not form a consistent pattern. Overall, the evidence points toward a dissociation that generalizes across demographic subgroups rather than definitive evidence of demographic heterogeneity in belief updating.

3.5. The Persistence of Recession Beliefs

Following prior literature which studies the role of beliefs in economic decision-making (e.g., Roth and Wohlfart (2020)), we conduct a follow-up survey two weeks after our initial experiment. While no information treatment is administered in our follow-up survey, we re-elicit individuals' beliefs expectations by asking them to forecast the likelihood of a recession in the U.S. in the first quarter of 2027. We use these data to distinguish two distinct notions of persistence: the persistence of the randomized treatment effects, created by the randomly assigned aggregate risk signals, and the persistence of individual belief levels over time.

We first examine whether the belief changes induced by the aggregate risk treatments persist. In Table IA5, the condition indicators identify each experimental condition's treatment effect relative to the Control group. For context, Column (1) estimates these effects using the full sample of participants in the initial experiment, while Column (2) re-estimates them on the subsample of participants who completed the follow-up questionnaire. Across the two specifications, the coefficients on the experimental condition indicators are similar in sign, magnitude, and ordering, indicating that follow-up participation is not strongly selective with respect to participants' responsiveness to their assigned recession signal. Column (3) then estimates the same effects in the follow-up wave. The treatment effects attenuate substantially over the two-week interval: relative to the baseline estimates in Column (2), roughly one-quarter to one-third of each treatment effect remains.[13] The surviving effects retain the sign and ordering of their baseline magnitudes, i.e., larger for the conditions that initially moved beliefs more, but are no longer statistically distinguishable from zero at conventional levels, consistent with the reduced precision of the follow-up sample.

The persistence of individual belief levels is stronger, but reflects a different construct. In Column (4), follow-up recession beliefs load on both the treatment-induced posterior and the pre-treatment prior, with coefficients of similar magnitude. Individuals who held higher recession beliefs during the baseline experiment tend to remain relatively more pessimistic two weeks later.

Finally, the decay dynamics do not differ by the source of the forecast. That is, the pairwise differences in the belief persistence estimates show that the surviving belief gaps are statistically indistinguishable across the AI and professional conditions. The source equivalence we document at baseline thus extends to the medium term, i.e., AI- and professionally-attributed recession signals move beliefs to a similar degree and their effects decay at a similar rate. Overall, the information treatments produce belief revisions that partially persist over two weeks, with directional consistency and symmetry across sources, though we estimate them with reduced precision in the follow-up sample.

[13]For instance, the coefficient on Professional_{20} of -2.949 (p-value = 0.102) suggests that, two weeks after the baseline experiment, recipients of the professional's forecast hold recession expectations roughly three percentage points lower than those of Control participants, which is directionally consistent with the baseline effect, though not statistically significant at conventional levels.

4. Belief Spillovers Across Economic Domains

Our prior evidence shows that individuals update their recession beliefs in response to aggregate risk information attributed to artificial intelligence models and professionals, and that they do so to a statistically equivalent degree. We now examine whether these AI-influenced beliefs propagate to other dimensions of households' economic outlook, as beliefs formed from conventional sources are known to do (Roth and Wohlfart, 2020). Whether AI-influenced beliefs propagate beyond the initial recession forecast determines whether the source equivalence is economically meaningful. That is, if they do propagate, it suggests that beliefs shaped by an AI signal can carry the same implications as beliefs shaped by a professional forecast.

To provide causal insights, we follow Roth and Wohlfart (2020) and rely on participants' random assignment to the experimental conditions because the information provisions should exogenously induce variation in participants' outlooks. Specifically, we conduct two-stage least squares regressions using experiment condition indicators as instruments for participants' posterior recession expectations such that:

$$y_i = \beta_0 + \beta_1 \widehat{Posterior\ Recession}_i + \theta' \mathbf{X}_i + \varepsilon_i, \tag{2}$$

where $\widehat{Posterior\ Recession}_i$ is obtained from the first-stage regression:

$$Posterior\ Recession_i = \gamma_0 + \gamma_1' Z_i + \gamma_2' \mathbf{X}_i + \nu_i, \tag{3}$$

The dependent variable, y_i, is each of our measures of a participant's economic beliefs. In particular, we examine macroeconomic expectations related to national unemployment, unemployment in the participant's county of residence, inflation rate and uncertainty, and firms' profit outcomes. We also examine participants' microeconomic beliefs related to their perceived financial prospects, income growth and uncertainty, as well as the stability of their employment status. Z_i comprises indicators for the Professional$_{20}$, AI$_{20}$, and AI$_{32}$ conditions, with Control participants as the reference group.

4.1. Effects on Other Macroeconomic Beliefs

We begin by examining whether participants adjust their macroeconomic beliefs in response to recession signals from AI and professional sources. Panel A of Table 6 shows the results from the instrumental variable (IV) regressions using our full sample. We have high first-stage F-statistics, suggesting concerns related to a weak instrument are limited. The second-stage estimates show that participants extrapolate their posterior recession expectations to beliefs about other future macroeconomic outcomes, including the national unemployment rate, unemployment rate in their county of residence, and inflation rate. In particular, as individuals perceive a higher likelihood of a future recession, they expect higher national and county unemployment rates and higher inflation, consistent with a stagflationary view of the economy documented in households' expectations (Andre, Pizzinelli, Roth, and Wohlfart, 2022; Coibion, Georgarakos, Gorodnichenko, and Van Rooij, 2023; D'Acunto, Charalambakis, Georgarakos, Kenny, Meyer, and Weber, 2024; Kamdar and Ray, 2025). In terms of magnitudes, the results suggest that a one percentage point increase in the average individual's aggregate risk expectation causes a 0.665 (p-value < 0.001) percentage point increase in their perceived likelihood that the national unemployment rate will rise and a 0.030 (p-value $= 0.039$) percentage point increase in the inflation rate they expect for the U.S. economy. Taken together, the evidence shows that households are likely to update multiple dimensions of their macroeconomic beliefs in response to receiving aggregate risk signals from AI models and professionals.

In Panel B, we examine whether AI-attributed recession signals have similar downstream effects on macroeconomic beliefs as signals of equivalent magnitude from professional economists. We estimate a restricted specification of Equation 3 in which Z_i contains only the AI$_{20}$ and Professional$_{20}$ condition indicators, exploiting the fact that both information treatments deliver an identical 20% recession signal with their source label as the sole point of variation.

The estimates convey two notable findings. First, the coefficients are similar in magnitude and statistical significance to those found when using the full set of instruments, again indicating that participants extend their recession beliefs to expectations about unemployment and inflation. Second, the Hansen J p-values exceed conventional significance levels, so we fail to reject the overidentifying restrictions. Because the only point of variation across our two instruments is their source label, this failure to reject indicates no detectable difference in the downstream responses induced by AI- and professionally-attributed recession signals.

4.2. Effects on Microeconomic Beliefs

In Table IA6, we report the results from examining whether recession beliefs affect participants' personal economic beliefs. We find no statistical evidence that individuals extrapolate their posterior recession expectations to their beliefs about their personal financial prospects. Among employed participants, posterior recession beliefs do not influence their perceived income trajectory or the likelihood of losing their job. However, consistent with higher recession beliefs prompting higher expected unemployment rates, both AI and professional recession signals influence participants' perceived likelihood of finding a new job if they become unemployed. In terms of magnitude, when using the matched 20% recession signals, the estimates in Column (5) of Panel B indicate that a one percentage point increase in aggregate risk expectation causes a 0.279 (p-value = 0.081) percentage point decline in the average participant's expected likelihood of finding a new job should they become unemployed. Additionally, since the Hansen J p-value exceeds conventional significance levels, we fail to reject the overidentifying restrictions. This indicates no detectable difference across sources and is consistent with a common local average treatment effect.

The pattern of microeconomic spillovers is indicative about the mechanism underlying belief propagation. Recession pessimism does not uniformly depress all forward-looking beliefs, as undirected priming would imply. Rather, the one personal belief that does respond, the perceived difficulty of finding a new job, is the one tightly linked to the aggregate unemployment expectations that recession beliefs influence. This selectivity suggests that participants integrate AI-attributed recession signals into a mental model of the economy, propagating them to connected variables rather than revising all beliefs in a common direction.

5. Behavioral Consequences of Recession Beliefs

Canonical models predict that innovations in individuals' economic outlooks should influence their consumption and investment decisions (Carroll, 1997; McKay et al., 2016; Challe and Ragot, 2016; Kuchler and Zafar, 2019). Whether AI-induced belief changes carry through to behavior determines whether they are economically meaningful. We, therefore, examine whether, and on which margins, individuals adjust their behavior in response to the recession signal. We focus on four behavioral categories: nondurable consumption intentions, realized nondurable spending, durable goods purchase climate and buying behavior, and stock market participation and trading.

For nondurables, we construct three composite spending intention indices by standardizing each component to mean zero and unit standard deviation before taking the equal-weighted average: (i) Total spending, which aggregates food, leisure, and clothing purchase intentions; (ii) Necessities, which aggregates food-at-home and restaurant food spending intentions; and (iii) Discretionaries, which averages leisure and clothing purchasing intentions. In addition, we measure follow-up survey participants' realized nondurable spending growth as the difference in the natural log of their spending on nondurables the week before the follow-up survey and the natural log of their spending the week before the primary experiment.

For consumption with respect to durable goods, we elicit whether participants believe it is currently a good or bad time to purchase durable goods. Responses are coded on a five-point Likert scale that is increasing in the favorability of the purchase climate. Using the follow-up questionnaire, we elicit whether an individual purchased a durable good during the two-week period between surveys. To identify participants' stock investment decisions, we ask, in the follow-up survey, whether they increased, decreased, or did not change their holdings of publicly traded stocks, mutual funds, or investment trusts over the two-week time span.

We use the consumption and investment decision measures as dependent variables when estimating IV regressions. In the IV regressions, we instrument for the difference between individuals' posterior and prior recession beliefs, i.e., *Belief Revision*, using *Shock* and experiment condition indicators as instruments. Panel A of Table 7 reports the full sample results.

We find that recession belief revisions significantly increase planned spending on necessities, with the estimates indicating that a one percentage point increase in recession belief updating is associated with a 0.007 standard deviation increase in planned food spending. The key coefficient when examining realized nondurable spending is statistically insignificant, suggesting no shift in net spending behavior.[14] For consumption of durable goods, we find no significant effects. With respect to investment decisions, we find that becoming more pessimistic about the economy reduces participants' propensity to trade. Moreover, non-investors are significantly more likely to remain out of the market (Columns (10) and (11)).

Panel B shows results when restricting the sample to participants in the AI_{20}, Professional_{20}, and Control conditions. Overall, the evidence is similar to that based on the full sample. Becoming more pessimistic about the economy influences individuals' nondurable consumption through increasing their intentions to spend on necessities. As in the main sample, non-investors are deterred from entering into the stock market as recession expectations are revised upwards. Additionally, the Hansen J p-values exceed traditional significance lev-

[14] Our data cannot cleanly identify a compositional shift in realized spending from discretionaries to necessary nondurables.

els. We therefore fail to reject the overidentifying restrictions, providing no evidence of a difference in downstream behavioral outcomes across sources.

Overall, the direction of these behavioral responses is consistent with patterns documented for beliefs formed from conventional information sources. Prior work shows that pessimistic macroeconomic conditions shift household consumption toward necessities and away from discretionary categories (Kamakura and Yuxing Du, 2012; Dutt and Padmanabhan, 2011), and that more pessimistic stock market and economic expectations are associated with lower equity exposure among households (Amromin and Sharpe, 2014; Giglio, Maggiori, Stroebel, and Utkus, 2021; Hoffmann, Post, and Pennings, 2013). Further, our finding that recession beliefs deter non-investors from entering the market aligns with life-cycle models in which a higher perceived probability of an economic downturn depresses stock market participation (Kolusheva, 2011; Alan, 2012). We extend these literatures by showing that the behavioral shifts can arise from AI-influenced belief revisions.

6. Conclusion

Households' beliefs about how the economy will evolve shape how they consume and save. We provide evidence that macroeconomic information signals from artificial intelligence models influence individuals' beliefs and behavior. By randomly assigning recession forecasts from AI models and a professional, we exogenously shift individuals' beliefs about the likelihood of a future recession. Our analysis of the consequences yields several findings.

First, individuals meaningfully revise their recession beliefs in response to AI-attributed forecasts. We document an average learning rate near one-quarter of the signal, which aligns with the partial weighting of advice identified in both the information-provision literature (Coibion et al., 2022) and in studies of how individuals incorporate algorithmic judgment (Logg et al., 2019). The observed incomplete learning is also consistent with imperfect-information models in which households update sparingly from new signals (Mankiw and

Reis, 2002; Carroll, 2003).

Second, our estimates indicate that individuals place similar weights on AI and professional forecasts when updating their recession beliefs. However, this source-equivalence at the population level masks a credibility deficit: individuals rate the AI-attributed forecasts as less accurate and less trustworthy. Taken together, the evidence indicates that algorithm aversion, in our setting focused on macroeconomic news, manifests in individuals' stated assessments of AI-attributed information more than in the weight they ultimately place on the signal.

Third, we find that individuals extrapolate their AI-influenced recession beliefs to other macroeconomic expectations, particularly to beliefs about unemployment and inflation, in a manner consistent with a stagflationary view of the economy. These spillover effects are statistically equivalent in magnitude across AI and professional forecast sources. Finally, these belief revisions carry behavioral consequences, increasing planned spending on nondurable necessities and deterring stock market entry among non-investors.

These findings have broader implications that suggest avenues for future research. As AI tools become a more pervasive conduit for economic information, they may alter how macroeconomic news diffuses through the economy, potentially reaching households that are inattentive to traditional media and accelerating the speed at which beliefs respond to aggregate signals (Beaudry and Portier, 2006; Barsky and Sims, 2012).[15] Whether such broader and faster diffusion compresses or widens the dispersion of household expectations is an open question (Angeletos and La'o, 2013; Dovern, Fritsche, and Slacalek, 2012). Additionally, investigating how heterogeneity in AI access and trust translates into differences in attention allocation and consumption responses is likely to be important, given the central role of such heterogeneity in heterogeneous-agent models (Kaplan, Moll, and Violante, 2018). For policymakers operating in an increasingly AI-mediated information environment,

[15] Consistent with this possibility, we find that more frequent users of AI hold baseline recession beliefs closer to the professional consensus and follow economic news more closely (Table IA7), though whether AI causally alters information acquisition is beyond our design.

the secondary dissemination of communications through AI tools could extend their reach and amplify their effects, but it also raises the risk that misinterpretation of the intended message propagates more widely. Ultimately, understanding how households come to rely on AI as a source of economic information, and how that reliance reshapes the transmission of macroeconomic news, is likely to be a promising direction for future research.

References

Alan, S. 2012. Do disaster expectations explain household portfolios? *Quantitative Economics* 3:1–28.

Amromin, G., and S. A. Sharpe. 2014. From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science* 60:845–866.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart. 2022. Subjective models of the macroeconomy: Evidence from experts and representative samples. *The Review of Economic Studies* 89:2958–2991.

Angeletos, G.-M., and J. La'o. 2013. Sentiments. *Econometrica* 81:739–779.

Armantier, O., W. Bruine de Bruin, S. Potter, G. Topa, W. Van Der Klaauw, and B. Zafar. 2013. Measuring inflation expectations. *Annual Review of Economics* 5:273–301.

Armantier, O., S. Nelson, G. Topa, W. Van der Klaauw, and B. Zafar. 2016. The price is right: Updating inflation expectations in a randomized price information experiment. *Review of Economics and Statistics* 98:503–523.

Barsky, R. B., and E. R. Sims. 2012. Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review* 102:1343–1377.

Beaudry, P., and F. Portier. 2006. Stock prices, news, and economic fluctuations. *American Economic Review* 96:1293–1307.

Bhandari, A., J. Borovička, and P. Ho. 2025. Survey data and subjective beliefs in business cycle models. *Review of Economic Studies* 92:1375–1437.

Blankespoor, E., J. Croom, and S. M. Grant. 2024. Generative AI and investor processing of financial information. *Working paper* .

Bloomberg. 2024. Bloomberg Launches AI-Powered Earnings Call Summaries. URL https://www.bloomberg.com/company/press/bloomberg-launches-ai-powered-earnings-call-summaries/. Accessed: 2026-05-27.

Bloomberg. 2025. Bloomberg Launches Gen AI Summarization for News Content. URL https://www.bloomberg.com/company/press/bloomberg-launches-gen-ai-summarization-for-news-content/. Accessed: 2026-05-27.

Burton, J. W., M.-K. Stein, and T. B. Jensen. 2020. A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making* 33:220–239.

Campbell, J. Y. 2006. Household finance. *The Journal of Finance* 61:1553–1604.

Capponi, A., S. Olafsson, and T. Zariphopoulou. 2022. Personalized robo-advising: Enhancing investment through client interaction. *Management Science* 68:2485–2512.

Carroll, C. D. 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics* 112:1–55.

Carroll, C. D. 2003. Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics* 118:269–298.

Cavallo, A., G. Cruces, and R. Perez-Truglia. 2017. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics* 9:1–35.

CDC. 2024. Current Cigarette Smoking Among Adults in the United States.

Challe, E., and X. Ragot. 2016. Precautionary saving over the business cycle. *The Economic Journal* 126:135–164.

Choukhmane, T., T. de Silva, W. Lin, and M. Akuzawa. 2026. AI Financial Advice: Supply, Demand, and Life Cycle Implications. *Working paper* .

Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber. 2024. The effect of macroeconomic uncertainty on household spending. *American Economic Review* 114:645–677.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. Van Rooij. 2023. How does consumption respond to news about inflation? Field evidence from a randomized control trial. *American Economic Journal: Macroeconomics* 15:109–152.

Coibion, O., and Y. Gorodnichenko. 2015. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105:2644–2678.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2022. Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy* 130:1537–1584.

Cottrill, J., and I. P. Affairs. 2024. Nearly two in five Americans turn to AI for financial management advice. https://www.ipsos.com/en-us/nearly-two-five-americans-turn-ai-financial-management-advice.

Da, Z., J. Engelberg, and P. Gao. 2015. The sum of all FEARS investor sentiment and asset prices. *The Review of Financial Studies* 28:1–32.

D'Acunto, F., J. Gao, L. Liu, K. Lu, Z. Wang, and J. Yang. 2024. Subjective expectations and financial intermediation. *Working paper.* .

D'Acunto, F., U. Malmendier, and M. Weber. 2023. What do the data tell us about inflation expectations? In *Handbook of Economic Expectations*, pp. 133–161. Elsevier.

D'Acunto, F., and A. G. Rossi. 2023a. Robo-advice: Transforming households into rational economic agents. *Annual Review of Financial Economics* 15:543–563.

D'Acunto, F., and A. G. Rossi. 2023b. Taming behavioral biases in consumer decision-making: The role of robo-advisors. *Available at SSRN 4666240* .

Dietvorst, B. J., J. P. Simmons, and C. Massey. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144:114.

Douglas, B. D., P. J. Ewell, and M. Brauer. 2023. Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *Plos one* 18:e0279720.

Dovern, J., U. Fritsche, and J. Slacalek. 2012. Disagreement among forecasters in G7 countries. *Review of Economics and Statistics* 94:1081–1096.

Dutt, P., and V. Padmanabhan. 2011. Crisis and consumption smoothing. *Marketing Science* 30:491–512.

D'Acunto, F., E. Charalambakis, D. Georgarakos, G. Kenny, J. Meyer, and M. Weber. 2024. Household inflation expectations: An overview of recent insights for monetary policy .

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber. 2021. Exposure to grocery prices and inflation expectations. *Journal of Political Economy* 129:1615–1639.

D'Acunto, F., N. Prabhala, and A. G. Rossi. 2019. The promises and pitfalls of robo-advising. *The Review of Financial Studies* 32:1983–2020.

D'Acunto, F., and M. Weber. 2023. Information and macroeconomic expectations: Global evidence. *Working paper* .

Enders, Z., S. T. Trautmann, and D. V. Vinogradov. 2026. Macroeconomic expectations: Formation, measurement, and consequences.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. 2021. The joint dynamics of investor beliefs and trading during the COVID-19 crash. *Proceedings of the National Academy of Sciences* 118:e2010316118.

Gomes, F., M. Haliassos, and T. Ramadorai. 2021. Household finance. *Journal of Economic Literature* 59:919–1000.

Greig, F., T. Ramadorai, A. G. Rossi, S. Utkus, and A. Walther. 2023. Algorithm aversion: Theory and evidence from robo-advice. *SSRN: https://ssrn. com/abstract* 4301514.

Haaland, I., C. Roth, and J. Wohlfart. 2023. Designing information provision experiments. *Journal of Economic Literature* 61:3–40.

Hansen, A. L., J. J. Horton, S. Kazinnik, D. Puzzello, and A. Zarifhonarvar. 2024. Simulating the Survey of Professional Forecasters. *Available at SSRN* .

Hoffmann, A. O., T. Post, and J. M. Pennings. 2013. Individual investor perceptions and behavior during the financial crisis. *Journal of Banking & Finance* 37:60–74.

Imbens, G. W., and D. B. Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Imbens, G. W., and J. M. Wooldridge. 2009. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature* 47:5–86.

Jussupow, E., I. Benbasat, and A. Heinzl. 2024. An integrative perspective on algorithm aversion and appreciation in decision-making. *MIS Quarterly* 48:1575–1590.

Kamakura, W. A., and R. Yuxing Du. 2012. How economic contractions and expansions affect expenditure patterns. *Journal of Consumer Research* 39:229–247.

Kamdar, R., and W. Ray. 2025. *Attention-driven sentiment and the business cycle*. Working paper.

Kaplan, G., B. Moll, and G. L. Violante. 2018. Monetary policy according to HANK. *American Economic Review* 108:697–743.

Kennedy, B., A. Tyson, and E. Saks. 2023. Public awareness of artificial intelligence in everyday activities. *Pew Research Center* .

Kolusheva, D. 2011. Life-cycle Portfolio Allocation When Disasters are Possible. *Working paper* .

Kuchler, T., and B. Zafar. 2019. Personal experiences and expectations about aggregate outcomes. *The Journal of Finance* 74:2491–2542.

Lakens, D., A. M. Scheel, and P. M. Isager. 2018. Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science* 1:259–269.

Logg, J. M., J. A. Minson, and D. A. Moore. 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes* 151:90–103.

Longoni, C., A. Bonezzi, and C. K. Morewedge. 2019. Resistance to medical artificial intelligence. *Journal of Consumer Research* 46:629–650.

Lusardi, A., and O. S. Mitchell. 2008. Planning and financial literacy: How do women fare? *American Economic Review* 98:413–17.

Lusardi, A., and O. S. Mitchell. 2011. Financial literacy and planning: Implications for retirement wellbeing .

Lusardi, A., and O. S. Mitchell. 2014. The economic importance of financial literacy: Theory and evidence. *American Economic Journal: Journal of Economic Literature* 52:5–44.

Mankiw, N. G., and R. Reis. 2002. Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics* 117:1295–1328.

Mankiw, N. G., R. Reis, and J. Wolfers. 2003. Disagreement about inflation expectations. *NBER Macroeconomics Annual* 18:209–248.

McKay, A., E. Nakamura, and J. Steinsson. 2016. The power of forward guidance revisited. *American Economic Review* 106:3133–3158.

Morris, M. G., and V. Venkatesh. 2000. Age differences in technology adoption decisions: Implications for a changing work force. *Personnel Psychology* 53:375–403.

OpenAI. 2025. Introducing ChatGPT Pulse. URL https://openai.com/index/introducing-chatgpt-pulse/. Accessed: 2026-05-27.

Roth, C., and J. Wohlfart. 2020. How do expectations about the macroeconomy affect personal expectations and behavior? *Review of Economics and Statistics* 102:731–748.

Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50:665–690.

Stradi, F., and G. Verdickt. 2025. Algorithm aversion, appreciation, and investor return beliefs. *The University of Auckland Business School Research Paper (forthcoming)* .

Tversky, A., and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185:1124–1131.

Venkatesh, V., M. G. Morris, G. B. Davis, and F. D. Davis. 2003. User acceptance of information technology: Toward a unified view1. *MIS Quarterly* 27:425–478.

White, H. 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48:817–838.

Figure 1: Participants' Prior and Posterior Recession Expectations

The figure shows histograms of participants' prior and posterior expectations with respect to the likelihood of a recession in the U.S. in the first quarter of 2027 across our primary experimental conditions. In the plots, the x-axis reports the likelihood of a recession. The y-axis reports the density of participants' forecasts. The black vertical line denotes the probability assigned to a recession in the experiment condition's information treatment, i.e., 20% for the Professional$_{20}$ and AI$_{20}$ conditions and 32% for the AI$_{32}$ condition.

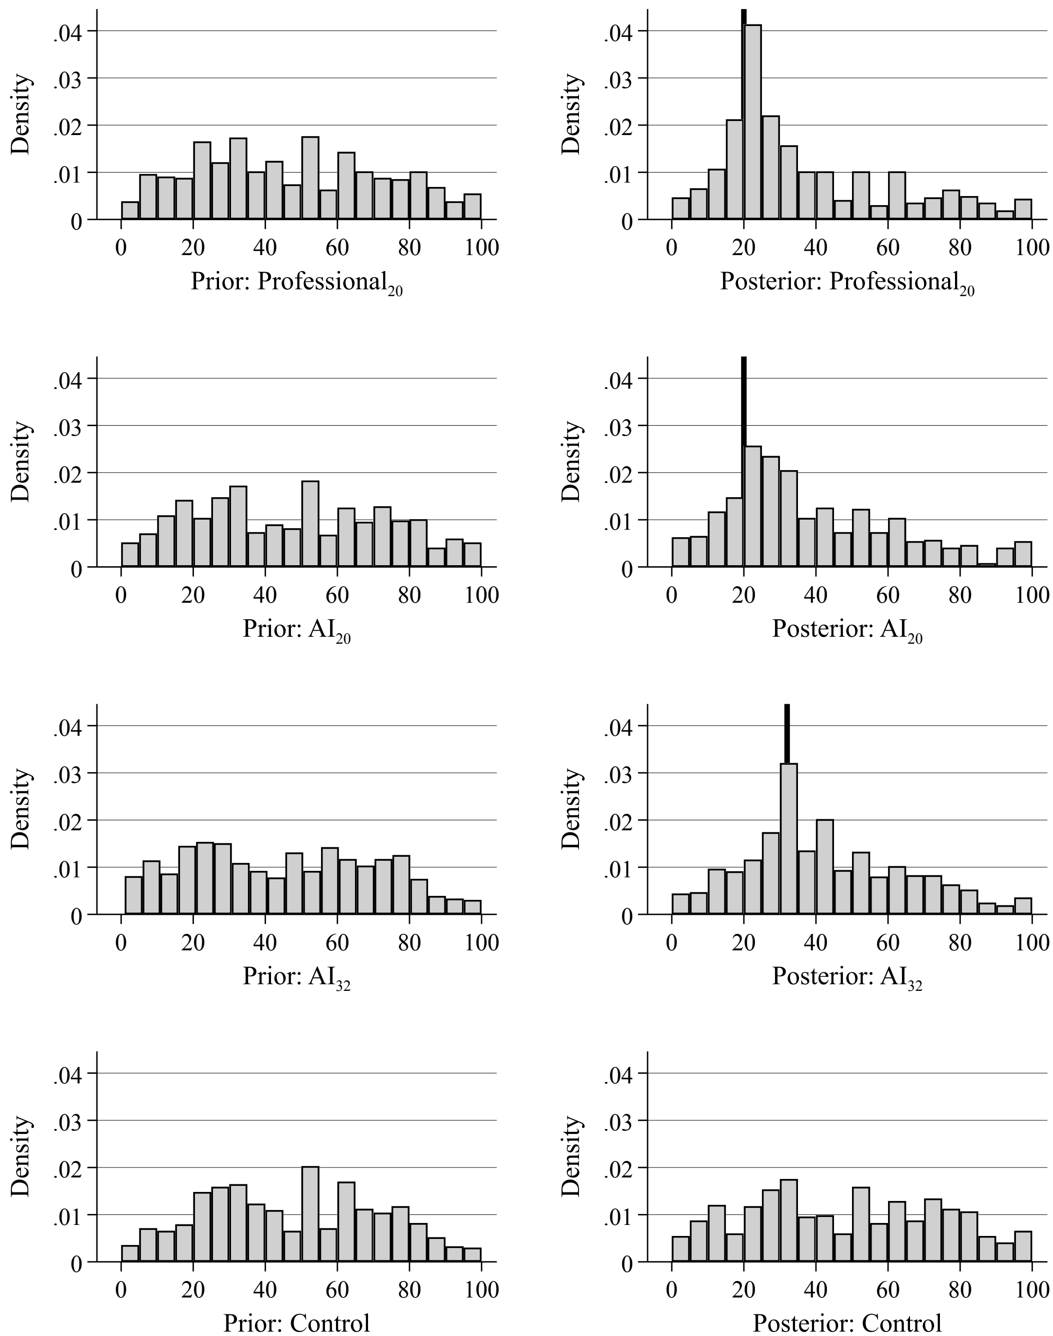

Figure 2: Information Signal Source Encoding

The figure presents evidence on participants' encoding of the source of their information treatment. Panel A shows the proportion of participants in each experimental condition who correctly identify the source of their information signal. The dashed horizontal line denotes the proportion expected under random guessing given six available answer options. Panel B shows the distribution, within each condition, of selected answers options for each information source.

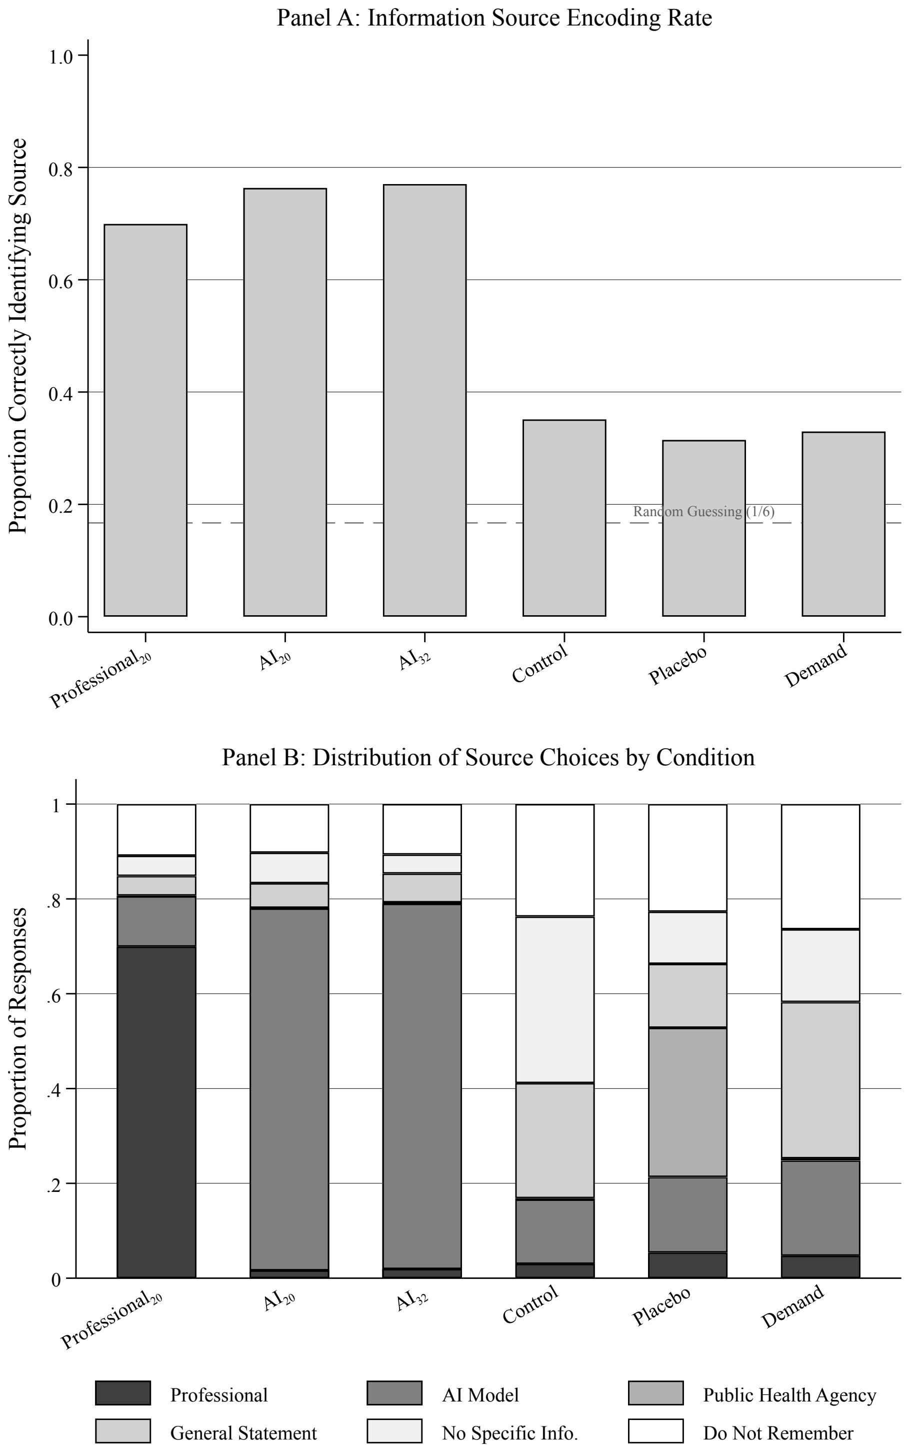

Table 1: Participants' Descriptive Statistics

The table reports participants' descriptive statistics. For each variable, condition-specific means appear in the first row, standard deviations in the second row, and normalized differences relative to the Control condition in the third row. The sample consists of participants in the Professional$_{20}$, AI$_{20}$, AI$_{32}$, and Control conditions. All variables are defined in Appendix I.

	Professional$_{20}$	AI$_{20}$	AI$_{32}$	Control
Age	2.218	2.445	2.343	2.144
	(1.471)	(1.554)	(1.567)	(1.422)
	[0.051]	[0.202]	[0.133]	–
Male	0.492	0.479	0.481	0.465
	(0.500)	(0.500)	(0.500)	(0.499)
	[0.055]	[0.028]	[0.033]	–
White	0.749	0.758	0.766	0.760
	(0.434)	(0.429)	(0.424)	(0.427)
	[-0.025]	[-0.005]	[0.015]	–
Married	0.446	0.38	0.402	0.429
	(0.497)	(0.486)	(0.491)	(0.495)
	[0.033]	[-0.100]	[-0.054]	–
College Degree	0.571	0.583	0.585	0.578
	(0.495)	(0.493)	(0.493)	(0.494)
	[-0.013]	[0.011]	[0.015]	–
Democrat	0.440	0.494	0.461	0.486
	(0.497)	(0.500)	(0.499)	(0.500)
	[-0.091]	[0.016]	[-0.050]	–
Income	3.641	3.356	3.473	3.587
	(1.756)	(1.837)	(1.805)	(1.758)
	[0.031]	[-0.129]	[-0.064]	–
Employment Status	2.212	2.423	2.336	2.102
	(1.781)	(1.835)	(1.793)	(1.690)
	[0.064]	[0.182]	[0.135]	–
Financial Literacy	2.593	2.591	2.568	2.524
	(0.731)	(0.724)	(0.751)	(0.741)
	[0.094]	[0.091]	[0.060]	–
Attentive	0.954	0.951	0.957	0.957
	(0.209)	(0.217)	(0.203)	(0.202)
	[-0.015]	[-0.032]	[-0.002]	–
News Consumption	0.655	0.642	0.668	0.627
	(1.162)	(1.134)	(1.172)	(1.190)
	[0.024]	[0.013]	[0.035]	–
AI Accuracy	0.499	0.457	0.549	0.547
	(1.496)	(1.384)	(1.451)	(1.424)
	[-0.033]	[-0.064]	[0.001]	–
AI Trust	0.227	0.149	0.233	0.230
	(1.687)	(1.512)	(1.608)	(1.532)
	[-0.002]	[-0.053]	[0.002]	–
AI Use Frequency	2.568	2.591	2.651	2.700
	(1.145)	(1.145)	(1.16)	(1.139)
	[-0.115]	[-0.095]	[-0.042]	–
Trust Fed. Reserve	1.937	1.941	1.917	1.937
	(0.899)	(0.913)	(0.912)	(0.904)
	[-0.000]	[0.005]	[-0.022]	–
Prior Recession	45.654	45.847	45.994	46.416
	(25.701)	(26.139)	(25.575)	(23.990)
	[-0.031]	[-0.023]	[-0.017]	–
Prior Confidence	1.960	1.902	1.925	1.944
	(1.136)	(1.166)	(1.089)	(1.089)
	[0.015]	[-0.037]	[-0.017]	–
N	725	731	723	729

Table 2: Baseline Learning Rates

The table reports OLS estimates of household belief updating in response to AI and professional recession forecasts. The dependent variable is *Belief Revision*, which is the difference between the participant's posterior and prior recession probability. *Shock* is the difference between the recession forecast signal the participant receives and their prior belief. *Prior Confidence* measures the participant's confidence in their priors, standardized to mean zero and unit standard deviation. *News Consumption* measures the participant's tendency to follow economic news, standardized to mean zero and unit standard deviation. Baseline controls include age, household income, gender, college education, marital status, race, financial literacy, employment status, and political affiliation, as well as indicators for demographic measures the participant chooses not to answer. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, $Professional_{20}$, and Control groups. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Shock	0.237***	0.236***	0.235***	0.283***	0.237***
	(0.028)	(0.028)	(0.028)	(0.030)	(0.028)
Shock × Prior Confidence				-0.088***	
				(0.025)	
Shock × News Consumption					-0.048**
					(0.023)
Prior Recession	-0.148***	-0.151***	-0.151***	-0.158***	-0.152***
	(0.024)	(0.024)	(0.024)	(0.026)	(0.024)
$Professional_{20}$	-4.522***	-4.526***	-4.528***	-4.004***	-4.478***
	(1.028)	(1.031)	(1.032)	(1.040)	(1.029)
AI_{20}	-2.522**	-2.529**	-2.528**	-1.678	-2.362**
	(1.043)	(1.050)	(1.051)	(1.064)	(1.046)
AI_{32}	-1.513*	-1.423	-1.425	-1.469	-1.385
	(0.904)	(0.908)	(0.908)	(0.905)	(0.905)
Controls	No	Yes	Yes	Yes	Yes
Attentive	No	No	Yes	Yes	Yes
N	2,908	2,908	2,908	2,908	2,908
Adj. R-sq.	0.247	0.254	0.254	0.276	0.259

Table 3: Learning Rates Across AI and Professional Sources

The table reports OLS estimates of participants' learning rates across information sources. The dependent variable in the regressions is *Belief Revision*, which is the difference between the participant's posterior recession expectation and their prior belief. *Shock* is the difference between the recession forecast signal the participant receives and their prior belief. AI_{32} equals one if the individual is randomly assigned to receive the AI recession forecast of 32%, and zero otherwise. AI_{20} equals one if the individual is randomly assigned to receive the AI recession forecast of 20%, and zero otherwise. $Professional_{20}$ equals one if the individual is randomly assigned to receive the professional's recession forecast of 20%, and zero otherwise. Baseline controls include age, household income, gender, college education, marital status, race, financial literacy, employment status, and political affiliation, as well as indicators for demographic measures the participant chooses not to answer. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, Professional_{20}, and Control groups. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)
Shock × Professional_{20}	0.231***	0.230***	0.229***
	(0.038)	(0.037)	(0.037)
Shock × AI_{20}	0.213***	0.212***	0.212***
	(0.035)	(0.035)	(0.035)
Shock × AI_{32}	0.269***	0.266***	0.266***
	(0.035)	(0.035)	(0.035)
Professional_{20}	-4.686***	-4.678***	-4.686***
	(1.036)	(1.039)	(1.040)
AI_{20}	-3.148***	-3.138***	-3.132***
	(1.076)	(1.089)	(1.091)
AI_{32}	-1.065	-1.000	-1.001
	(0.907)	(0.908)	(0.908)
Prior Recession	-0.148***	-0.150***	-0.150***
	(0.024)	(0.024)	(0.024)
Controls	No	Yes	Yes
Attentive	No	No	Yes
N	2,908	2,908	2,908
Adj. R-sq.	0.247	0.254	0.254
Pairwise Learning Rate Differences:			
AI_{20} − Professional_{20}	-0.018	-0.018	-0.017
	(0.040)	(0.039)	(0.039)
AI_{32} − Professional_{20}	0.038	0.036	0.037
	(0.039)	(0.039)	(0.039)
AI_{32} − AI_{20}	0.056	0.054	0.054
	(0.037)	(0.037)	(0.037)

Table 4: Forecast Credibility: Source Discount

The table reports OLS estimates of the credibility participants assign to the recession forecast they receive. The dependent variable in the regression is denoted in the column header. *Forecast Accuracy* is the participant's perceived accuracy of their recession forecast signal. *Forecast Trust* is participants' degree of trust in the source of their recession forecast signal. AI_{32} and AI_{20} are experimental condition indicators. *AI Accuracy* and *AI Trust* are the participant's pre-treatment attitudes toward the accuracy and trustworthiness of AI information. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Estimates on control variables are suppressed for presentation purposes. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, and $Professional_{20}$ conditions. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	Forecast Accuracy		Forecast Trust	
	(1)	(2)	(3)	(4)
AI_{32}	-0.162***	-0.156***	-0.180***	-0.184***
	(0.051)	(0.048)	(0.050)	(0.046)
AI_{20}	-0.355***	-0.339***	-0.385***	-0.371***
	(0.050)	(0.047)	(0.050)	(0.046)
AI Accuracy		0.255***		
		(0.014)		
AI Trust				0.262***
				(0.012)
N	2,179	2,179	2,179	2,179
Adj. R-sq.	0.039	0.173	0.045	0.220
Pairwise Credibility Differences:				
AI_{32} - AI_{20}	0.192***	0.183***	0.205***	0.187***
	(0.050)	(0.045)	(0.050)	(0.043)

Table 5: Credibility Moderation of Belief Updating

The table reports OLS estimates of participants' learning rates and how they vary with the perceived credibility of the recession forecast. The dependent variable in the regressions is *Belief Revision*, the difference between the participant's posterior recession expectation and their prior belief. *Shock* is the difference between the recession forecast the participant receives and their prior belief. *Forecast Accuracy* is the participant's perceived accuracy of their recession forecast signal. *Forecast Trust* is participants' degree of trust in the source of their recession forecast signal. AI_{32}, AI_{20}, and $Professional_{20}$ are experimental condition indicators. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Estimates on control variables are suppressed for presentation purposes. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, and Professional_{20} conditions. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)
Shock × AI_{32}	0.462***	0.465***
	(0.026)	(0.027)
Shock × AI_{20}	0.506***	0.484***
	(0.027)	(0.029)
Shock × Professional_{20}	0.426***	0.412***
	(0.029)	(0.028)
Shock × AI_{32} × Forecast Accuracy	0.163***	
	(0.025)	
Shock × AI_{20} × Forecast Accuracy	0.228***	
	(0.019)	
Shock × Professional_{20} × Forecast Accuracy	0.160***	
	(0.023)	
Shock × AI_{32} × Forecast Trust		0.175***
		(0.025)
Shock × AI_{20} × Forecast Trust		0.193***
		(0.021)
Shock × Professional_{20} × Forecast Trust		0.154***
		(0.022)
Forecast Accuracy	1.089***	
	(0.408)	
Forecast Trust		0.710
		(0.433)
AI_{32}	1.535*	1.530*
	(0.866)	(0.868)
AI_{20}	2.382**	2.017**
	(0.938)	(0.934)
N	2,179	2,179
Adj. R-sq.	0.354	0.349
Pairwise Differences in Credibility-Updating Slopes:		
AI_{20} - Professional_{20}	0.068**	0.040
	(0.028)	(0.028)
AI_{32} - Professional_{20}	0.003	0.021
	(0.033)	(0.032)

Table 6: Extrapolating to Other Macroeconomic Beliefs

The table reports estimates from instrumental variable regressions of the effects participants' posterior recession expectations on additional dimensions of their macroeconomic beliefs. Panel A reports estimates based on the full sample while Panel B reports regression estimates using participants in the AI_{20}, $Professional_{20}$, and Control conditions. The dependent variable in each regression is denoted in the column header. *Posterior Recession* is the participant's posterior recession forecast. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Their coefficients are suppressed for presentation purposes. All variables are defined in Appendix I. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	National Unemployment (Percent) (1)	National Unemployment (Categorical) (2)	County Unemployment (Percent) (3)	Inflation Rate (Average) (4)	Inflation Rate (Std. Dev.) (5)	Firms' Profits (6)
Panel A: Full Sample						
Posterior Recession	0.665***	0.011***	0.503***	0.030**	0.004	-0.005
	(0.103)	(0.004)	(0.101)	(0.014)	(0.010)	(0.004)
Prior Recession	0.074	0.003	0.123*	0.016	-0.005	-0.005
	(0.071)	(0.003)	(0.070)	(0.010)	(0.007)	(0.003)
N	2,908	2,908	2,908	2,908	2,908	2,908
First stage F-statistic	54.76	54.76	54.76	54.76	54.76	54.76
Hansen J statistic	0.513	3.807	0.647	2.151	5.596	1.124
Hansen J p-value	0.774	0.149	0.724	0.341	0.061	0.570
Panel B: Restricted Sample						
Posterior Recession	0.677***	0.010***	0.484***	0.025*	-0.000	-0.005
	(0.105)	(0.004)	(0.103)	(0.015)	(0.010)	(0.005)
Prior Recession	0.066	0.004	0.128*	0.022**	-0.003	-0.005
	(0.076)	(0.003)	(0.074)	(0.011)	(0.007)	(0.003)
N	2,185	2,185	2,185	2,185	2,185	2,185
First stage F-statistic	78.97	78.97	78.97	78.97	78.97	78.97
Hansen J statistic	0.044	0.377	0.477	0.055	0.959	0.868
Hansen J p-value	0.834	0.539	0.490	0.814	0.327	0.351

Table 7: Behavioral Consequences

The table reports estimates from instrumental variable regressions of the effects participants' recession belief revisions on behavioral outcomes. Panel A reports estimates based on the full sample while Panel B reports regression estimates using participants in the AI_{20}, $Professional_{20}$, and Control conditions. The dependent variable in each regression is denoted in the column header. *Belief Revision* is the difference between the participant's posterior recession expectation and their prior belief. *Posterior Recession* is the participant's posterior recession forecast. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Their coefficients are suppressed for presentation purposes. All variables are defined in Appendix I. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	Consumption Growth: Total	Consumption Growth: Necessities	Consumption Growth: Discretionaries	Consumption Growth: Realized	Durables Purchase Climate	Purchased Durables	Purchased Stocks	Sold Stocks	No Trades	Non-investors: Purchased Stocks	Non-investors: No Trades
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A: Full Sample											
Belief Revision	0.005*	0.007**	0.004	0.003	-0.003	0.003	-0.002	-0.002	0.004*	-0.004*	0.004*
	(0.003)	(0.003)	(0.004)	(0.005)	(0.004)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)
Prior Recession	-0.003**	0.000	-0.006***	0.002	-0.008***	0.001	-0.001	-0.000	0.001	-0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
N	2,908	2,908	2,908	1,289	2,908	1,291	1,291	1,291	1,291	462	462
First stage F-statistic	50.61	50.61	50.61	20.00	50.61	20.00	20.00	20.00	20.00	9.27	9.27
Hansen J statistic	3.653	7.127	0.931	2.914	0.548	1.254	2.495	1.758	1.071	5.006	3.140
Hansen J p-value	0.301	0.068	0.818	0.405	0.908	0.740	0.476	0.624	0.784	0.171	0.371
Panel B: Restricted Sample											
Belief Revision	0.007**	0.008**	0.005	0.002	-0.003	0.003	-0.002	-0.002*	0.004	-0.004*	0.005**
	(0.003)	(0.003)	(0.004)	(0.005)	(0.004)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.003)
Prior Recession	-0.003**	0.001	-0.006***	0.001	-0.008***	0.001	-0.001	-0.000	0.001*	-0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
N	2,185	2,185	2,185	963	2,185	964	964	964	964	340	340
First stage F-statistic	59.99	59.99	59.99	23.49	59.99	23.45	23.45	23.45	23.45	10.70	10.70
Hansen J statistic	0.588	0.837	0.773	0.853	0.332	0.774	2.636	0.490	1.154	3.219	0.639
Hansen J p-value	0.745	0.658	0.680	0.653	0.847	0.679	0.268	0.783	0.561	0.200	0.727

Appendix I. Variable Definitions

Name	Definition
Demographics Variables	
Age	Age group of the participant: 0. 18 - 24; 1. 25 - 34; 2. 35 - 44; 3. 45 - 54; 4. 55 - 64; 5. 65 - 74 6. 75 or older.
AI Accuracy	Participant's pre-treatment, general belief that information produced by AI systems is accurate. Measured on a seven-point Likert scale ranging from -3 (Strongly disagree) to 3 (Strongly agree).
AI Trust	Participant's pre-treatment, general belief that AI systems can be trusted. Measured on a seven-point Likert scale ranging from -3 (Strongly disagree) to 3 (Strongly agree).
AI Use	Participant's reported frequency of using AI systems during a typical day, on a five-point Likert scale ranging from 1 (Not at all) to 5 (Seven or more times per day).
Attentive	Equals one if the participant correctly answers the attention check, zero otherwise.
College Degree	Equal to one if the participant has a bachelor's degree or higher, zero otherwise.
Democrat	Equals one if the participant identifies with the Democratic party, zero otherwise.
Employment Status	Participant's employment status: 1. Full-time employed; 2. Part-time employed; 3. Self-employed; 4. Unemployed and seeking work; 5. Unemployed and not seeking work; 6. Retired; 7. Student.
Financial Literacy	Index from 0 (Low literacy) to 3 (High literacy) based on three literacy questions developed by Lusardi and Mitchell (2008, 2011).
Income	Participant's household income in 2025: 0. Less than \$15,000; 1. Between \$15,000 and \$25,000; 2. Between \$25,000 and \$50,000; 3. Between \$50,000 and \$75,000; 4. Between \$75,000 and \$100,000; 5. Between \$100,000 and \$150,000; 6. Between \$150,000 and \$200,000; 7. More than \$200,000.
Male	Equal to one if the participant identifies as male, zero otherwise.
Married	Equal to one if the participant is married, zero otherwise.
News Consumption	Participant's self-reported frequency of following news about the national economy, on a five-point Likert scale ranging from -2 (Strongly disagree) to 2 (Strongly agree).
Trust Fed. Reserve	Participant's degree of trust in the Federal Reserve, on a five-point Likert scale ranging from 0 (Not at all) to 4 (A great deal).
White	Equal to one if the participant identifies as White, zero otherwise.
Experimental Conditions	
AI_{20}	Equal to one if the participant is randomly assigned to receive the AI recession forecast of 20%, zero otherwise.
AI_{32}	Equal to one if the participant is randomly assigned to receive the AI recession forecast of 32%, zero otherwise.
Demand	Equal to one if the participant is randomly assigned to receive the experimenter demand information treatment, zero otherwise.
Control	Equal to one if the participant is randomly assigned to receive no recession forecast, zero otherwise.
Placebo	Equal to one if the participant is randomly assigned to receive the health information treatment, zero otherwise.
$Professional_{20}$	Equal to one if the participant is randomly assigned to receive the professional economist's recession forecast of 20%, zero otherwise.

Name	Definition
Recession Beliefs and Forecast Credibility	
Belief Revision	Difference between the participant's posterior recession expectation and their prior belief.
Follow-up Recession	Participant's recession probability forecast in the follow-up survey.
Forecast Accuracy	Participant's perceived accuracy of the recession forecast signal they received, on a five-point Likert scale ranging from 1 (Not at all accurate) to 4 (A great deal).
Forecast Trust	Participant's degree of trust in the source of the recession forecast signal they received, on a five-point Likert scale ranging from 0 (Not at all) to 4 (A great deal).
Posterior Recession	Participant's recession probability forecast after receiving their information treatment.
Prior Confidence	Participant's self-reported confidence in their initial recession belief, on a five-point Likert scale ranging from 0 (Not at all confident) to 4 (Extremely confident).
Prior Recession	Participant's initial recession probability forecast.
Shock	Difference between the recession probability signal in the participant's information treatment and their prior recession belief. Equals zero for Control participants.
Economic Beliefs	
County Unemployment	Participant's probabilistic belief that unemployment in their county of residence will be higher over the next twelve months.
Financial Prospects	Participant's response on a seven-point Likert scale ranging from -3 (Much worse off) to 3 (Much better off) to: "Looking ahead, do you think you (and any family living with you) will be financially better or worse off twelve months from now than you are these days?"
Firms' Profits	Participant's expectation about U.S. firms' profitability over the next twelve months on a five-point Likert scale ranging from -2 (Strongly decrease) to 2 (Strongly increase).
Income Growth	Participant's expected earnings growth rate over the next twelve months (mean of the participant's density forecast), conditional on remaining in their main job.
Income Uncertainty	Standard deviation of the participant's density forecast for earnings growth over the next twelve months.
Inflation Rate	Mean of the participant's density forecast for the U.S. inflation rate over the next twelve months.
Inflation Uncertainty	Standard deviation of the participant's density forecast for the U.S. inflation rate over the next twelve months.
Personal Reemployment	Participant's perceived probability that they would find an acceptable replacement job if they lost their current position.
Personal Unemployment	Participant's perceived probability of losing their main job within the next twelve months.
National Unemployment (Categorical)	Participant's response on a five-point Likert scale ranging from -2 (Unemployment will strongly decrease) to 2 (Unemployment will strongly increase) to the question: "What do you think will happen to unemployment in the U.S. over the next twelve months?"
National Unemployment (Percent)	Participant's perceived probability that the national unemployment rate will increase over the next twelve months.

Name	Definition
Consumption and Investment Behavior	
Consumption Growth: Discretionaries	Equal-weighted index of spending intentions for leisure activities and clothing. Each component is standardized to mean zero and unit standard deviation before averaging.
Consumption Growth: Necessities	Equal-weighted index of spending intentions for food at home and restaurant food. Each component is standardized to mean zero and unit standard deviation before averaging.
Consumption Growth: Realized	Difference in the natural log of nondurable consumption spending during the week before the follow-up survey and the week before the main experiment.
Consumption Growth: Total	Equal-weighted index of nondurable goods spending intentions.
Durable Purchase Climate	Participant's perceived climate for purchasing durable goods, on a five-point Likert scale ranging from -2 (A very bad time) to 2 (A very good time).
No Trades	Indicator equal to one if the participant reports having not traded stocks, mutual funds, or investment trusts during the two-week period between surveys, zero otherwise.
Purchased Durables	Indicator equal to one if the participant reports having purchased a durable good during the two-week period between surveys, zero otherwise.
Purchased Stocks	Indicator equal to one if the participant reports having increased their holdings of publicly traded stocks, mutual funds, or investment trusts during the two-week period between surveys, zero otherwise.
Sold Stocks	Indicator equal to one if the participant reports having decreased their holdings of publicly traded stocks, mutual funds, or investment trusts during the two-week period between surveys, zero otherwise.

Internet Appendix

IA1. Recession Forecast Prompt

To elicit a forecast of the probability of a recession in the United States during the first quarter of 2027 from the artificial intelligence model, we provide it with the following prompt:

> A recession entails a fall of real Gross Domestic Product (GDP) which means that in a given quarter the economy produces less final goods and services than in the preceding quarter. A fall of real GDP is normally reflected in a decline of real income, employment, industrial production, and wholesale-retail sales. Consider relevant macroeconomic theories and economic data available to you as of today, March 1, 2026. Provide your single best numeric forecast of what you think the probability is that there will be a fall of real U.S. GDP in the first quarter of 2027 compared to the fourth quarter of 2026.

IA2. Bayesian Updating of Beliefs

To quantify individuals' updating in response to receiving a recession information provision, we rely on the model of Roth and Wohlfart (2020). Like these authors, we assume that individuals' subjective probability distributions with respect to a future recession follow beta distributions: $\theta_i \sim \text{Beta}(a_i, b_i)$. The distribution is defined on the interval [0,1] with $a_i > 0$ and $b_i > 0$. Next, assuming quadratic loss, individuals' forecasts should be their best guess, i.e., the average of their personal probability distributions, as they seek to minimize the expected squared error between their forecast and the actual outcome. Accordingly, individuals' initial recession beliefs can be described as:

$$\textit{prior recession}_i = \mathbb{E}[\theta_i] = \frac{a_i}{a_i + b_i}. \tag{4}$$

We also expect that individuals interpret the information provision from both the human expert and the AI model as outcomes from a Bernoulli trial wherein n_0 denotes a failure, i.e., no recession, and n_1 marks a success, i.e., a recession. Consequently, individuals should interpret that the source of their information provision has a sample of $n = n_0 + n_1$ observations and the signal describes the fraction of recessions in its sample such that:

$$signal = \frac{n_1}{n_0 + n_1}. \tag{5}$$

If these assumptions hold, then individuals' posterior recession expectations would also follow a beta distribution described by the parameters: $c_i = a_i + n_1$ and $d_i = b_i + n_0$. As before, an individual would report the average of their probability distribution such that:

$$posterior\ recession_i = \mathbb{E}[\theta_i] = \frac{a_i + n_1}{a_i + b_i + n}. \tag{6}$$

Accordingly, the revision process can then be constructed as:

$$posterior\ recession_i - prior\ recession_i = \frac{a_i + n_1}{a_i + b_i + n}(signal - prior\ recession_i) \tag{7}$$

The term $signal - prior\ recession_i$ is the gap between the recession probability conveyed by the information signal and the individual's prior belief, which we denote as *Shock* in the main text. The coefficient on *Shock* in our empirical specifications therefore estimates the weight that individuals place on their signal, i.e., their learning rate. As noted by Roth and Wohlfart (2020), this framework generates several predictions. First, an individual's degree of learning, i.e., the weight placed on the signal relative to the prior recession forecast, will lie within the interval [0,1]. Second, households should place a greater weight on the signal if they perceive it to be more credible. For example, greater weight should be placed on information from a source that is perceived to have a larger sample size (n). Third, individuals who have greater confidence in their initial beliefs should have a lower degree of belief revision in response to the information provision. Fourth, individuals' confidence in their posterior expectations should be at least as high as the confidence in their prior belief.

Fifth, the degree of learning in response to the information signal should be positively related with the change in confidence.

Figure IA1: Information Treatment: Professional$_{20}$ Condition

The figure shows the information treatment provided to participants randomly assigned to the Professional$_{20}$ condition. In this example, the prior belief about the probability of a recession is 40%.

You said that you think that the probability of a recession in the U.S. in the first quarter of 2027 is **40 percent**.

Now, we would like to provide you with information about the expectations of a **professional economist** on the likelihood of a recession. The forecaster is knowledgeable about current macroeconomic data and the relevant economic theories.

According to a professional that regularly takes part in a survey of professional forecasters by the Federal Reserve Bank of Philadelphia, the probability of a recession in the first quarter of 2027 is **20%.**

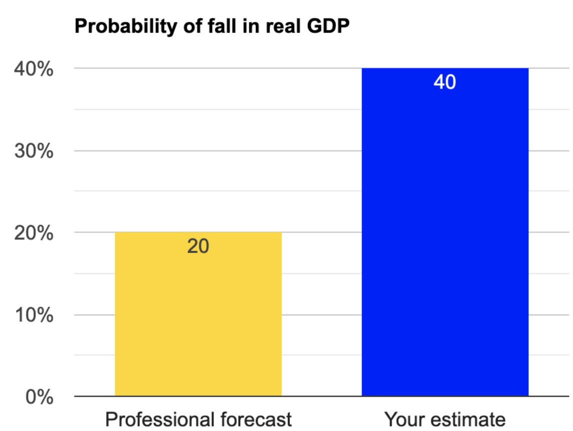

Figure IA2: Information Treatment: AI_{20} Condition

The figure shows the information treatment provided to participants randomly assigned to the AI_{20} condition. In this example, the prior belief about the probability of a recession is 40%.

You said that you think that the probability of a recession in the U.S. in the first quarter of 2027 is **40 percent**.

Now, we would like to provide you with information about the expectations of an **artificial intelligence model** on the likelihood of a recession. The AI model has access to current macroeconomic data and the relevant economic theories.

According to the AI model, the probability of a recession in the first quarter of 2027 is **20%**.

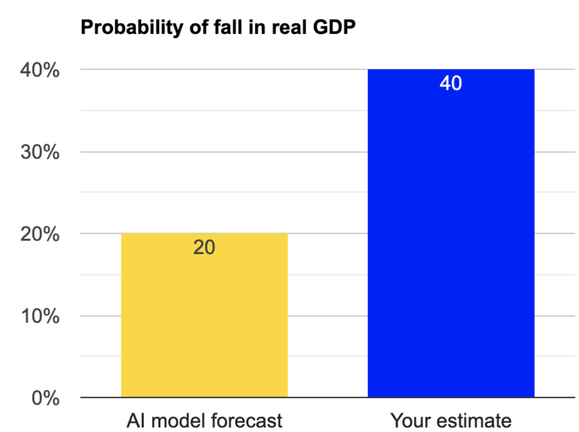

Figure IA3: Distribution of Belief Shock Across Matched Forecast Sources

The figure shows kernel density estimates of the belief shock for participants in the AI_{20} and $Professional_{20}$ conditions, the two arms that receive an economically identical 20% recession forecast and differ only in the attributed source. In the plot, the x-axis reports *Shock*, the difference between the recession forecast signal the participant receives and their prior belief. The y-axis reports the density of participants' belief shocks. The solid line denotes the AI_{20} condition and the dashed line denotes the $Professional_{20}$ condition.

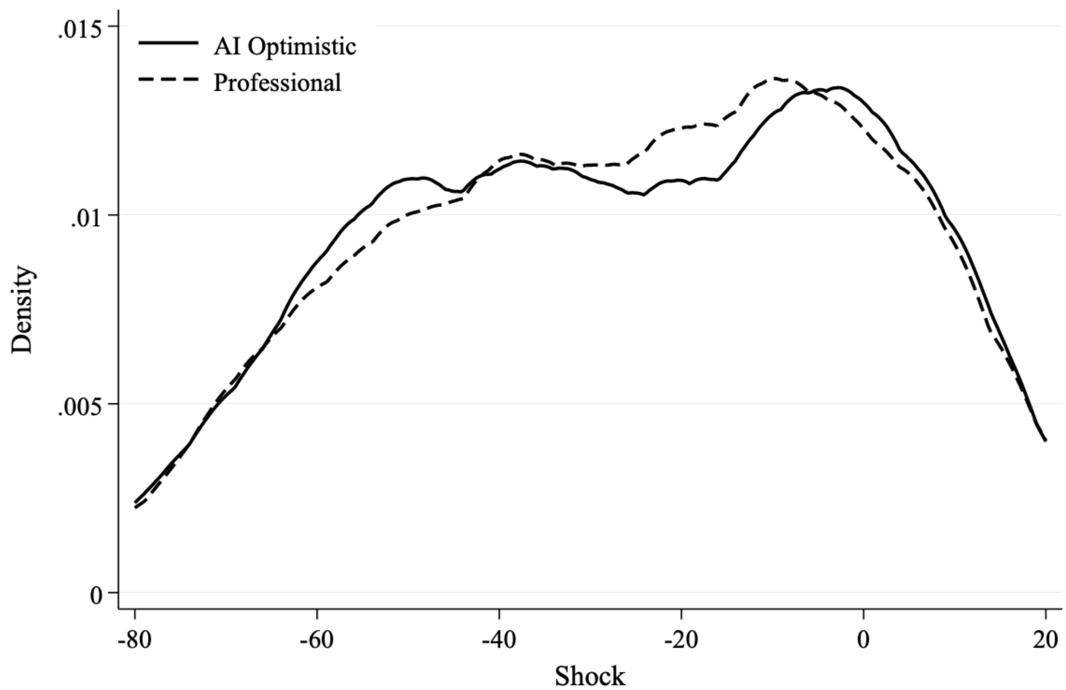

Table IA1: Alternative Explanations

The table reports OLS estimates of Placebo and Demand condition participants' belief updating. The dependent variable is *Belief Revision*, which is the difference between the participant's posterior and prior recession probability. *Shock* is the difference between the recession forecast signal the participant receives and their prior belief. *Placebo* equals one if the participant is assigned to the Place condition, zero otherwise. *Demand* equals one if the participant is assigned to the Demand condition, zero otherwise. Baseline controls include age, household income, gender, college education, marital status, race, financial literacy, employment status, and political affiliation, as well as indicators for demographic measures the participant chooses not to answer. All variables are defined in Appendix I. In Columns (1) and (2), the sample consists of participants in the Placebo and Control groups. In Columns (3) and (4), the sample consists of participants in the Demand and Control conditions. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Shock	0.222***	0.221***		
	(0.031)	(0.031)		
Shock × Placebo	-0.126***	-0.128***		
	(0.041)	(0.041)		
Placebo	-0.711	-0.520		
	(1.300)	(1.308)		
Professional_{20}	-4.920***	-4.873***		
	(1.030)	(1.038)		
AI_{20}	-2.923***	-2.856***		
	(1.060)	(1.074)		
Demand			-0.418	-0.170
			(0.905)	(0.935)
Prior Recession	-0.148***	-0.153***	-0.130***	-0.143***
	(0.024)	(0.024)	(0.017)	(0.018)
Controls	No	Yes	No	Yes
Attentive	No	Yes	No	Yes
N	2,630	2,630	1,172	1,172
Adj. R-sq.	0.213	0.224	0.037	0.050
Placebo Learning Rate:				
Total Learning Rate	0.096**	0.093**		
	(0.043)	(0.043)		

Table IA2: Alternative Specification

The table reports OLS estimates of Placebo and Demand condition participants' belief updating. The dependent variable is *Posterior Recession*, which is the participant's recession probability forecast after receiving their information provision. Baseline controls include age, household income, gender, college education, marital status, race, financial literacy, employment status, and political affiliation, as well as indicators for demographic measures the participant chooses not to answer. Estimates on control variables are suppressed for presentation purpose. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, $Professional_{20}$, and Control groups. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)
Prior Recession × $Professional_{20}$	-0.231***	-0.229***
	(0.038)	(0.037)
Prior Recession × AI_{20}	-0.213***	-0.212***
	(0.035)	(0.035)
Prior Recession × AI_{32}	-0.269***	-0.266***
	(0.035)	(0.035)
$Professional_{20}$	-0.068	-0.099
	(1.606)	(1.606)
AI_{20}	1.113	1.112
	(1.617)	(1.635)
AI_{32}	7.555***	7.510***
	(1.671)	(1.674)
Prior Recession	0.852***	0.850***
	(0.024)	(0.024)
Controls	No	Yes
Attentive	No	Yes
N	2,908	2,908
Adj. R-sq.	0.524	0.528
Pairwise Learning Rate Differences:		
AI_{20} – $Professional_{20}$	0.018	0.017
	(0.040)	(0.039)
AI_{32} – $Professional_{20}$	-0.038	-0.037
	(0.039)	(0.039)
AI_{32} - AI_{20}	-0.056	-0.054
	(0.037)	(0.037)

Table IA3: Learning Rates Among Signal Source Encoders

The table reports OLS estimates of participants' source-specific learning rates among those who correctly recall the source of their recession signal. The dependent variable in the regressions is *Belief Revision*, which is the difference between the participant's posterior recession expectation and their prior belief. *Shock* is the difference between the recession forecast signal the participant receives and their prior belief. AI_{32} equals one if the individual is randomly assigned to receive the AI recession forecast of 32%, and zero otherwise. AI_{20} equals one if the individual is randomly assigned to receive the AI recession forecast of 20%, and zero otherwise. $Professional_{20}$ equals one if the individual is randomly assigned to receive the professional's recession forecast of 20%, and zero otherwise. Baseline controls include age, household income, gender, college education, marital status, race, financial literacy, employment status, and political affiliation, as well as indicators for demographic measures the participant chooses not to answer. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, Professional_{20}, and Control groups. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)
Shock × Professional_{20}	0.208***	0.208***
	(0.046)	(0.047)
Shock × AI_{20}	0.192***	0.187***
	(0.044)	(0.044)
Shock × AI_{32}	0.263***	0.260***
	(0.043)	(0.044)
Professional_{20}	-4.881***	-4.895***
	(1.465)	(1.499)
AI_{20}	-3.794***	-3.839**
	(1.466)	(1.497)
AI_{32}	-1.357	-1.414
	(1.211)	(1.221)
Prior Recession	-0.152***	-0.146***
	(0.032)	(0.033)
Controls	No	Yes
Attentive	No	Yes
N	1,878	1,878
Adj. R-sq.	0.263	0.267
Pairwise Learning Rate Differences:		
$\text{AI}_{20} - \text{Professional}_{20}$	-0.017	-0.021
	(0.044)	(0.045)
$\text{AI}_{32} - \text{Professional}_{20}$	0.055	0.051
	(0.044)	(0.044)
$\text{AI}_{32} - \text{AI}_{20}$	0.072*	0.073*
	(0.041)	(0.042)

Table IA4: Individual Heterogeneity in Source-Specific Belief Updating

The table reports OLS estimates of the moderating effects of individual characteristics on participants' source-specific learning rates. Each column corresponds to a separate regression in which the learning rate is interacted with the moderator named in the column header. The dependent variable in all regressions is *Belief Revision*, the difference between the participant's posterior recession expectation and their prior belief. *Shock* is the difference between the recession forecast signal the participant receives and their prior belief. AI_{32}, AI_{20}, and $Professional_{20}$ are experimental condition indicators. The moderators are the participant's frequency of using AI systems (*AI Use*), degree of trust in the Federal Reserve (*Trust Fed. Reserve*), and the demographic characteristics indicated in the column headers. Continuous moderators are standardized to mean zero and unit standard deviation. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Estimates on control variables and lower-order interaction terms are suppressed for presentation purposes. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, and Professional_{20} conditions. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	AI Use (1)	Trust Fed. Reserve (2)	Age (3)	Income (4)	Male (5)	College Degree (6)	Married (7)	White (8)	Democrat (9)	Financial Literacy (10)
Shock x AI_{32}	0.416***	0.416***	0.421***	0.412***	0.412***	0.423***	0.405***	0.506***	0.350***	0.413***
	(0.026)	(0.026)	(0.026)	(0.026)	(0.034)	(0.041)	(0.033)	(0.053)	(0.036)	(0.026)
Shock x AI_{20}	0.360***	0.360***	0.361***	0.351***	0.359***	0.357***	0.370***	0.362***	0.392***	0.357***
	(0.028)	(0.027)	(0.027)	(0.027)	(0.034)	(0.037)	(0.032)	(0.052)	(0.040)	(0.027)
Shock x Professional_{20}	0.373***	0.392***	0.375***	0.376***	0.391***	0.349***	0.382***	0.451***	0.340***	0.372***
	(0.030)	(0.029)	(0.030)	(0.029)	(0.040)	(0.040)	(0.037)	(0.055)	(0.040)	(0.029)
Shock x AI_{32} x Moderator	0.071***	0.013	-0.097***	0.030	0.001	-0.019	0.019	-0.124**	0.130***	-0.023
	(0.023)	(0.023)	(0.024)	(0.025)	(0.050)	(0.051)	(0.052)	(0.060)	(0.050)	(0.024)
Shock x AI_{20} x Moderator	0.028	0.028	-0.026	-0.022	-0.005	-0.001	-0.041	-0.008	-0.055	-0.033
	(0.024)	(0.024)	(0.024)	(0.023)	(0.047)	(0.046)	(0.050)	(0.056)	(0.048)	(0.022)
Shock x Professional_{20} x Moderator	0.010	0.062***	-0.012	0.048*	-0.038	0.047	-0.022	-0.109*	0.066	-0.039
	(0.023)	(0.024)	(0.029)	(0.025)	(0.052)	(0.052)	(0.052)	(0.061)	(0.053)	(0.029)
Moderator	-0.276	-0.664	-0.714*	0.163	1.311*	-0.078	-0.671	0.084	1.878**	0.233
	(0.410)	(0.433)	(0.412)	(0.483)	(0.788)	(0.862)	(0.852)	(1.007)	(0.841)	(0.471)
N	2,179	2,179	2,179	2,179	2,179	2,179	2,179	2,179	2,179	2,179
Adj. R-sq.	0.277	0.281	0.278	0.275	0.272	0.272	0.272	0.275	0.276	0.273
Pairwise Differences in Moderator-Updating Slopes Across Sources:										
AI_{20} - Professional_{20}	0.018	-0.035	-0.014	-0.070**	0.033	-0.048	-0.019	0.100	-0.121*	0.006
	(0.031)	(0.032)	(0.035)	(0.032)	(0.065)	(0.064)	(0.068)	(0.076)	(0.066)	(0.033)
AI_{32} - Professional_{20}	0.061**	-0.049	-0.085**	-0.018	0.040	-0.066	0.041	-0.015	0.064	0.016
	(0.031)	(0.032)	(0.035)	(0.033)	(0.069)	(0.070)	(0.071)	(0.080)	(0.069)	(0.035)

Table IA5: Persistence of Posterior Recession Beliefs

The table reports OLS estimates of the persistence of randomized treatment effects and the persistence of individual belief levels over the two-week period following the experiment. The dependent variable in each regression is denoted in the column header. Column (1) is estimated on the full sample, while Columns (2)–(4) are restricted to participants who completed the follow-up survey. *Prior Recession* is the participant's initial recession belief. *Posterior Recession* is the participant's posterior recession belief from the baseline experiment. *Follow-up Recession* is the participant's re-elicited recession belief in the two-week follow-up survey. AI_{32} equals one if the individual is randomly assigned to receive the AI-attributed recession forecast of 32%, and zero otherwise. AI_{20} equals one if the individual is randomly assigned to receive the AI-attributed recession forecast of 20%, and zero otherwise. $Professional_{20}$ equals one if the individual is randomly assigned to receive the professional's recession forecast of 20%, and zero otherwise. *Pairwise Differences in Belief Persistence* report, for the follow-up participants, the differences in the surviving treatment effects across sources. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Their coefficients are suppressed for presentation purposes. All variables are defined in Appendix I. The sample consists of participants in the AI_{20}, AI_{32}, Professional_{20}, and Control groups. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	Posterior Recession (1)	Posterior Recession (2)	Follow-up Recession (3)	Follow-up Recession (4)
Professional_{20}	-10.700***	-8.802***	-2.949	-0.469
	(0.909)	(1.284)	(1.803)	(1.808)
AI_{20}	-8.728***	-8.239***	-2.518	-0.196
	(0.889)	(1.312)	(1.787)	(1.777)
AI_{32}	-4.801***	-3.729***	-0.942	0.109
	(0.882)	(1.283)	(1.716)	(1.690)
Prior Recession	0.667***	0.649***	0.455***	0.272***
	(0.014)	(0.021)	(0.027)	(0.038)
Posterior Recession				0.282***
				(0.041)
N	2,908	1,291	1,291	1,291
Adj. R-sq.	0.518	0.522	0.236	0.268
Pairwise Differences in Belief Persistence:				
AI_{20} - Professional_{20}			0.431	0.273
			(1.808)	(1.768)
AI_{32} - Professional_{20}			2.007	0.578
			(1.748)	(1.712)
AI_{32} - AI_{20}			1.576	0.305
			(1.725)	(1.678)

Table IA6: Extrapolating to Microeconomic Beliefs

The table reports estimates from instrumental variable regressions of the effects participants' posterior recession expectations on dimensions of their microeconomic beliefs. Panel A reports estimates based on the full sample while Panel B reports regression estimates using participants in the AI_{20}, $Professional_{20}$, and Control conditions. The dependent variable in each regression is denoted in the column header. *Posterior Recession* is the participant's posterior recession belief. *Prior Recession* is the participant's initial recession belief. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Their coefficients are suppressed for presentation purposes. All variables are defined in Appendix I. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	Financial Prospects (Categorical)	Income Growth (Average)	Income Uncertainty (Std. Dev.)	Personal Unemployment (Percent)	Personal Reemployment (Percent)
	(1)	(2)	(3)	(4)	(5)
Panel A: Full Sample					
Posterior Recession	-0.006	-0.002	0.010	-0.027	-0.302*
	(0.005)	(0.018)	(0.012)	(0.104)	(0.159)
Prior Recession	-0.010***	-0.013	-0.015*	0.203***	0.092
	(0.004)	(0.013)	(0.008)	(0.072)	(0.110)
N	2,908	1,988	1,988	1,968	1,987
First stage F-statistic	54.76	41.46	41.46	40.12	41.61
Hansen J statistic	2.404	0.499	2.425	4.732	1.443
Hansen J p-value	0.301	0.779	0.297	0.094	0.486
Panel B: Restricted Sample					
Posterior Recession	-0.005	-0.002	0.008	-0.055	-0.279*
	(0.005)	(0.018)	(0.012)	(0.105)	(0.160)
Prior Recession	-0.011***	-0.013	-0.015*	0.218***	0.074
	(0.004)	(0.013)	(0.009)	(0.075)	(0.115)
N	2,185	1,514	1,514	1,500	1,513
First stage F-statistic	78.97	61.92	61.92	59.52	62.14
Hansen J statistic	0.171	0.485	0.019	1.884	0.038
Hansen J p-value	0.679	0.486	0.890	0.170	0.846

Table IA7: AI Use and Households' Macroeconomic Informedness

The table reports estimates from OLS regressions of measures of participants' initial informedness on their frequency of AI use. The dependent variable in each regression is denoted in the column header. *SPF Distance* is the absolute distance between a participant's prior recession probability and the consensus probability among professionals participating in the Survey of Professional Forecasters (SPF). Lower values indicate beliefs closer to the professional consensus. *News Consumption* measures how closely the participant follows economic news. *Prior Confidence* is the participant's self-reported confidence in their initial recession belief. *AI Use* is a standardized measure of how frequently the participant uses AI systems. Controls included in the regressions are age, household income, gender, college education, marital status, race, financial literacy, employment status, political affiliation, and attentiveness, as well as indicators for demographic measures the participant chooses not to answer. Their coefficients are suppressed for presentation purposes. All variables are defined in Appendix I. Standard errors are adjusted for heteroskedasticity (White, 1980) and are presented in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	SPF Distance (1)	News Consumption (2)	Prior Confidence (3)
AI Use	-1.345***	0.179***	0.176***
	(0.335)	(0.019)	(0.018)
N	3,796	3,796	3,796
Adj. R-sq.	0.004	0.023	0.024