How Much Can I Make? Insights on Belief Updating in the Labor Market^{*}

Sebastian Brown[†], Kenneth Chan[‡]

April 21, 2025

Abstract

We use a nationally representative survey to study how people update their wage expectations. Using the recently developed excess belief movement test (Augenblick and Rabin, 2021), we find strong evidence of non-Bayesian learning at the aggregate level. We find that the average belief movement is 518% of the average uncertainty reduction in their beliefs, 418% more than the Bayesian benchmark. Our simulations show that this result is unlikely to be explained solely by measurement error. We also find evidence of asymmetric updating, where individuals update their beliefs more when they receive higher wage offers than expected.

Keywords: Job Search, Wage Expectations, Belief Updating, Non-Bayesian, Martingale Property **IEL Codes:** D83, D84, D90, 164

JEL Codes: D83, D84, D90, J64

^{*}We are grateful to Gary Charness, Erik Eyster, Lorenz Goette, Peter Kuhn, Ryan Oprea, Jessica Pan, Peter Rupert, Dick Startz, Sevgi Yuksel, and Basit Zafar for their helpful comments and suggestions. All remaining errors are our own.

[†]University of California, Santa Barbara. <u>sebastian_brown@ucsb.edu</u>.

[‡]National University of Singapore. ken_chan@nus.edu.sg.

1 Introduction

Canonical job search models typically assume that the agent knows the underlying wage distribution (Mortensen, 1970; McCall, 1970). This class of model has an optimal decision rule of setting a reservation value and accepting the first offer that is better than the reservation value (Weitzman, 1979). While these models give us valuable insights into job search behavior, it is often unrealistic to assume that people know the underlying wage distribution.

Another class of search models assumes an unknown underlying distribution, where the agent has to learn about the underlying distribution over time by observing the offers they receive (Rothschild, 1978; Rosenfield and Shapiro, 1981; Talmain, 1992; Li and Yu, 2018; Potter, 2021).¹ In this class of search models, the agent sets a reservation value that is a function of the agent's beliefs, which varies over time as the agent learns about the underlying distribution.

Although the latter class of models is more realistic, the empirical application of these models has been limited due to insufficient high-quality data on job searcher beliefs.² Most studies that accommodate learning assume Bayesian updating in the calibrated model (e.g. Potter (2021)). Consequently, there has been a lack of research on how people learn and update their beliefs in light of new information. Our paper provides further insight into how people update their wage expectations over time. These beliefs can influence people's search behavior, which will affect their earnings and the duration of unemployment.

In the last two decades, nudges (Thaler and Sunstein, 2008) have become a popular policy tool for policy makers to influence people's behavior. One of the more popular forms of nudges is information provision, which is a potent policy tool for changing people's beliefs and behavior. Examining the process of belief updating can provide valuable insights into understanding the effectiveness of information provision policy. In a similar spirit, Coffman, Featherstone, and Kessler (2024) presents a model of information nudges where the prior belief of the marginal agent determines the effectiveness of the information provision policy.³

¹Most of these models are consumer search models where an agent is looking for a good with the lowest price or a product with the highest quality. These models can be easily extended to a job search context.

²The beliefs refer to their wage expectations or arrival rate of job offers.

³Their model assumes Bayesian updating and heterogeneity in the agents' prior belief. How the marginal agent responds to the new information will determine the sign of the treatment effect.

In a series of field experiments, information provision has been shown to help college students form more accurate wage expectations (Wiswall and Zafar, 2015; Jiang and Zen, 2023). Arni (2016) found that a coaching program was successful in increasing job-finding rates among the treated job searchers by 9 percent. The author argues that coaching helped workers to have more realistic expectations and search more effectively. In another experiment, Gee (2018) randomly displayed information to LinkedIn users on the number of workers applying for a specific job and found that this additional information increased the probability that a worker completes a job application by 3.5 percent. Although several studies have found that information provision is effective in changing people's beliefs and influencing their actions, Jones and Santos (2022) showed that over-optimism can persist because people are unresponsive to negative news.⁴

Our paper differentiates itself from information provision experiments in two key ways. First, unlike many studies mentioned above that use treatment interventions to alter people's beliefs, we rely on nationally representative survey data—the Survey of Consumer Expectations—to examine how individuals learn about their wage distribution without any experimental manipulation. Since our data is free from the influences of an experimental setup, it offers new insights into how people update their beliefs in a non-experimental labor market context. Second, our approach allows us to analyze how individuals update their beliefs compared to the Bayesian benchmark, revealing the systematic biases to which they are prone.

Testing for Bayesian updating in the field is challenging because it typically requires constructing a Bayesian benchmark, which demands strong assumptions or a highly detailed dataset on individuals' beliefs.⁵ We deviate from this approach by using an excess belief movement test (Augenblick and Rabin, 2021) that tests the martingale property in belief updating. It is important to note that the martingale property is necessary but not sufficient for Bayesian updating.⁶ The martingale prop-

⁴Some other context in which information provision policy is used includes improving people's knowledge about COVID-19 (Sadish et al., 2021), eliminating statistical racial discrimination in a patient's choice of medical professional (Chan, 2022) and reducing disagreement about the extent of racial discrimination (Haaland and Roth, 2023).

⁵It is easier to study belief updating in the lab as the lab offers the experimenter control over the environment, allowing the experimenter to easily compute the Bayesian benchmark. Refer to Benjamin (2019) for a detailed literature of belief-updating experiments in the lab.

⁶Cripps (2018) showed that Bayes' rule can be primarily characterized by the martingale and the divisibility property, and Chan (2025) showed that Bayes' rule can be primarily characterized by the martingale and the preservation of the monotone likelihood ratio property.

erty requires that before seeing a signal, the expected updated beliefs over the signal realizations should be equal to the prior. The intuition of this property is the agent should not expect his beliefs to change before seeing a signal. There are two statistics required for the excess belief movement test. Belief movement is defined as the squared difference of changes in belief, and uncertainty reduction is defined as the amount of reduction in the belief's variance from updating. If the agent's updating rule satisfies the martingale property, the average belief movement will be equal to the average uncertainty reduction.

Using this test, we find an average belief movement that is more than five times the average uncertainty reduction. We can reject that people are updating their beliefs in a manner consistent with the martingale property, allowing us to reject that people are Bayesian. Our result also suggests that on average, people over-update relative to the Bayesian benchmark. This updating pattern is consistent with biases such as overreaction to signals and base rate neglect. This could potentially explain why information provision policy has been mostly effective in a job search context.

Given that people are non-Bayesian in updating their beliefs, the rest of our analysis focuses on determining whether a policy intervention is necessary. The key criterion we use to assess the need for such a policy intervention is whether individuals can eventually learn the true wage distribution. Some non-Bayesian updating behaviors, such as overreacting to signals, can still allow individuals' beliefs to converge to the true wage distribution. However, other updating biases like base rate neglect and asymmetric updating can hinder individuals from learning their actual wage distribution, even in an information-rich environment. In extreme cases of asymmetric updating, individuals' beliefs may even converge to an incorrect distribution. In the survey data, we found patterns of asymmetric updating and suggestive evidence of base rate neglect. This suggests a need for policy intervention to help people learn about their wage distribution.

The paper that is closest to ours is Conlon, Pilossoph, Wiswall, and Zafar (2018). Despite having a similar research question and using the same dataset, our papers differ in three ways. Firstly, at the general level, our paper is primarily interested in using the survey data to test various updating rules and identifying models that best describe people's updating behavior. In contrast, their paper focuses on how information friction affects job search behavior. In our paper, we further examine and identify different belief updating patterns instead of just comparing them against

the Bayesian benchmark.

Secondly, we adopt a different methodology for studying non-Bayesian updating patterns. Their analysis assumes that the agent is updating in a Gaussian framework where the priors and signals (with a correctly perceived variance) are normally distributed. The wage offers the respondents received between the surveys are assumed to be the only signals. This allows them to compute the Bayesian benchmark and compare the survey respondent updating behavior from the survey against this Bayesian benchmark, where they found over-updating relative to the Bayesian benchmark. Our approach uses the excess belief movement test which tests for the martingale property instead of Bayesian updating. This circumvents the need to compute the Bayesian benchmark to compare if people are Bayesian, which allows us to make fewer and more conservative assumptions. Our test allows us to reject all updating rules without the martingale property. Rejecting the martingale property will allow us to reject a larger class of updating rules, as well as identify mistakes like having incorrect priors.⁷

Lastly, the excess belief movement test requires only the initial and updated beliefs. This approach does not restrict our sample to respondents who received a wage offer, which is advantageous since job offers are not the only signals conveying information about one's wage distribution in the labor market. For instance, the absence of a wage offer may itself provide insight into a respondent's wage offer distribution (Milgrom, 1981; Jin, Luca, and Martin, 2021), or individuals may update their beliefs through social learning (DeGroot, 1974; Golub and Jackson, 2010) by observing their peers. This allows us to use a larger sample from the SCE data, which makes our result more representative of the population.

We contribute to the existing literature in two main ways. Firstly, while field experiments and surveys on information provision have demonstrated that people update their beliefs in light of new information, they are often unable to assess how people update their beliefs relative to the Bayesian benchmark. Using the excess belief movement test, we provide insights into how people are updating relative to the Bayesian benchmark and uncover systematic biases in their updating behavior. These biases give us insights into why information provision policies are effective and how we can better design information policies.

⁷We are grateful to Sevgi Yuksel for pointing out that incorrect priors can cause the martingale property to fail even when agents are Bayesian.

Secondly, most theoretically driven belief-updating experiments are conducted in laboratory settings using abstract belief updating tasks. We offer an empirical test of different belief-updating rules using survey data about people's wage expectations. Our study provides supporting evidence for belief updating biases similar to those documented in lab experiments, thereby reinforcing the findings from lab experiments.

The remainder of the paper proceeds as follows. Section 2 reviews relevant literature. Section 3 describes the data. Section 4 outlines a descriptive theoretical model of labor market updating that relates to the SCE survey questionnaire. Section 5 explains our empirical strategy and the excess belief movement test. Section 6 then presents our results. We conclude in section 7 with a broader discussion of our approach, suggesting future directions for research.

2 Literature Review

Our work is related to two strains of literature: non-Bayesian updating, and learning in job search. Over the last two decades, behavioral economists have shown robust evidence that people are not Bayesian and some systematic biases that people are susceptible to. Behavioral theorists have come up with belief updating models to accommodate these non-Bayesian behavior (Grether, 1980; Epstein, Noor, and Sandroni, 2010; Hagmann and Loewenstein, 2017). Some of the more well-known biases include base rate neglect (Kahneman and Tversky, 1972; Esponda, Vespa, and Yuksel, 2024), a phenomenon where people underweight their priors, and conservatism bias (Phillips and Edwards, 1966), a situation where people are insensitive to new information. While most of the empirical evidence comes from laboratory experiments,⁸ some recent papers have used field data to show that people update their beliefs in a non-Bayesian manner (Conlon, Pilossoph, Wiswall, and Zafar, 2018; Bordalo, Gennaioli, Porta, and Shleifer, 2019; Bordalo, Gennaioli, Ma, and Shleifer, 2020; Augenblick and Rabin, 2021).

There is a burgeoning literature that empirically studies learning in job search and how people's behavior deviates from the Bayesian benchmark. Firstly, Kudlyak, Lkhagvasuren, and Sysuyev (2014) finds that job seekers first apply to jobs that match their education levels. But with prolonged unemployment, they apply to jobs that require a lower education level. They argue that this is evidence that searching workers

⁸See Benjamin (2019) for a survey of the experimental literature on belief updating.

learn to adjust their expectations downward over the unemployment spell. Similarly, Mueller, Spinnewijn, and Topa (2021) found that unemployed workers adjust their beliefs downwards but not sufficiently. This results in the long-term unemployed displaying an optimistic bias in job finding. Jones and Santos (2022) found evidence of asymmetric updating where individuals underreact to negative news using a field experiment with Mozambican undergraduates. This asymmetric updating pattern is attributed to the limitation of information provision policy and the persistence of over-optimism bias.

Contrary to the findings of the above mentioned papers, Conlon, Pilossoph, Wiswall, and Zafar (2018) use the SCE data and find that people over-update their wage beliefs relative to the Bayesian benchmark. They estimated that on average, wage expectations increase by \$0.47 for every one-dollar increase in observed wage offer, while the Bayesian benchmark is estimated to be \$0.16.

Instead of studying wage expectations, Potter (2021) examines how the expectation of offer arrival rate changes over time using data from the Great Recession. Using a calibrated model, he showed that learning can explain the job search dynamics observed during the Great Recession.

Our work also relates more generally to a broader, yet still relatively new, literature on behavioral job search (DellaVigna et al., 2017; Cooper and Kuhn, 2020), as well as more general work studying employed job search (Faberman, Mueller, Şahin, and Topa, 2022; Ahn and Shao, 2021).

3 Data

3.1 Overall Description

We are using a public dataset from the New York Federal Reserve's Survey of Consumer Expectations (SCE), which is a nationally representative survey. The survey is divided into two parts, a core set of questions that remain the same every month and a supplementary set of questions that rotates between several different economic topics. Our analysis focuses on the labor market supplement, which is administered every March, July, and November. Throughout this paper, any reference to "the survey" pertains specifically to this labor supplement. Subjects can be surveyed for up to three times, and are replaced on a rolling basis. This feature allows us to observe the same individual for up to a maximum of three times, allowing us to observe the dynamics of wage expectations. In return for completing the survey, respondents are paid \$15 for each survey.⁹

The main advantage of this dataset is that it explicitly elicits wage expectations from survey respondents, meaning there is no need to indirectly infer beliefs as in many previous studies (e.g. Potter (2021)). This survey contains a question that elicits a belief distribution over wages. This is the main question that allows us to implement the excess belief movement test.¹⁰ The last benefit is that the survey is representative of the American population and it features mostly employed individuals, allowing us to study the wage expectations of employed individuals, a previously understudied group (Faberman, Mueller, Şahin, and Topa, 2022). However, this comes at the cost of including fewer unemployed individuals.¹¹

3.1.1 Timeline

To study belief updating, we require at least two responses from the same individual to observe the changes in belief. Hence, we restrict our sample to only respondents who completed at least two surveys. Respondents can complete the survey up to three times, which allows us to observe at most two updates from each individual. The timeline with respect to the survey is shown in figure 1.



Figure 1: Timeline of Survey

Between two surveys, the respondent will receive information about their underlying wage distribution which enables the updating of beliefs. The survey also asks

 $^{^9{\}rm For}$ a more in-depth description of the structure and administration of the SCE, see Armantier et al. (2017).

¹⁰Just eliciting the wage expectation is insufficient to study Non-Bayesian depending. Any changes in the wage expectations can be rationalized by how disperse the beliefs are.

¹¹While the SCE has been included in analyses of unemployed workers (e.g. Mueller, Spinnewijn, and Topa (2021); Conlon, Pilossoph, Wiswall, and Zafar (2018); Faberman, Mueller, Şahin, and Topa (2022)), the fraction of the sample that is unemployed is very small, chiefly because the SCE is a representative sampling of the entire population, not of the unemployed who are actively searching for a job. As noted in table 1, only 3.5% of our final sample is unemployed.

for the wages of the job offer that the respondent received in the last four months which is the duration between the two surveys.¹²

3.1.2 Descriptive Statistics

The composition of our final dataset, categorized by employment status, job search status, and the receipt of a job offer, is presented in the table 1. We define survey 1 as the first survey where a response to the question eliciting the belief distribution was provided.¹³ As this survey is representative of the American population, most of our sample are employed individuals, individuals who are not searching for a job and individuals who did not receive any job offers. For individuals who are not searching for a job, we interpret their responses about their wage expectations as hypothetical or wages from unsolicited offers.

Count	Survey 1	Survey 2	Survey 3	Total	
Number Unemployed	83	93	36	212	
Number Employed	2,063	2,047	904	$5,\!014$	
Number Not in Labor Force	286	299	105	690	
Missing Employment Status	27	20	11	58	
Received Offer(s) in last 4 Months	521	440	165	1,126	
No Offers in last 4 months	1,938	2,019	891	4,848	
Searched for Jobs	651	589	250	1,490	
Did Not Search for Jobs	$1,\!606$	$1,\!673$	720	$3,\!999$	
Missing Search Status	202	197	86	485	
Number of Survey Responses	2,459	$2,\!459$	1,056	5,974	
Data Collection Period	3/2015-11/2019				

Table 1: Dataset composition by survey responses. Notes: Survey counts reflect the number of individuals responding to the first, second, or third consecutive survey. Individuals with at least one survey in 2020 were excluded to avoid measuring the effects of the pandemic on expectations. Individuals reporting expected annual wages (using either question OO2a or OO2a2 for either period) below \$10,000 were dropped, as it is suggestive that these individuals are not reporting their annual salary or they are looking for part-time jobs. 618 observations were dropped this way, 257 of which had a reported wage or one of the expectation questions less than \$100. 24 observations were also dropped for individuals who moved to a different state between surveys, who would be unlikely to have a stable wage offer distribution across surveys.

¹²The individual can update their beliefs with information other than wage offers. Our main analysis does not restrict to individuals who have received a wage offer.

¹³Some individuals only provide the belief distribution in the second survey that they completed. For expositional purposes, we label these responses as survey 1 as we are using this response as the initial belief for the excess belief movements test.

Most job search studies primarily look at unemployed individuals while our sample comprises mostly of employed individuals who are not searching for a job. Despite the majority our sample not actively searching for a job, it is still important to study the wage expectations of employed individuals as this belief would influence their decision to quit their job and their perceived outside option (Jäger, Roth, Roussille, and Schoefer, 2024).

3.2 Survey Questions

Our main analysis centers around the questions about wage expectations which merit more detailed discussions. Figure 2 shows the survey questionnaire used to elicit the respondent's wage expectations and belief distribution. It is important to note that as most of the survey respondents are not looking for a job, we interpret their responses as a hypothetical response assuming they are searching for a job or expectations from unsolicited offers.

OO2a2 - OO2a2 (Added March 2015) Think about the job offers that you may receive within the coming four months. Roughly speaking, what do you think the annual salary for the best offer will be for the first year?

Note the best offer is the offer you would be most likely to accept.

dollars

OO2b - OO2b (shown if OO2a2 > 0 each response is % of OO2a2 ranging from .8 to 1.2) (Added November 2014)

Think again about the job offers that you may receive within the coming four months. What do you think is the percent chance that the job with the best offer will have an annual salary for the first year of...

The best offer is the offer you would be most likely to accept.

Less than [0.8* OO2a2] dollars (1)	%(1)
Between [0.8* OO2a2] dollars and [0.9* OO2a2] dollars (2)	% (2)
Between [0.9* OO2a2] dollars and [1.0* OO2a2] dollars (3)	% (3)
Between [1.0* OO2a2] dollars and [1.1* OO2a2] dollars (4)	% (4)
Between [1.1* OO2a2] dollars and [1.2* OO2a2] dollars (5)	% (5)
More than [1.2* OO2a2] dollars (6)	% (6)

OO2a - OO2a (Added November 2014)

Think about the job offers that you may receive within the coming four months. Roughly speaking, what do you think the average annual salary for these offers will be for the first year?

____ dollars

Figure 2: Survey Questions OO2a2, OO2b and OO2a from SCE

The survey elicits the expectation of the best offer that the survey respondent

may receive in the next four months. In the follow-up question, the survey elicits the belief distribution of receiving the best offer that belongs to a range that is determined by the subject's response in "OO2a2". The excess belief movement test uses the distribution of the best wage offer elicited in the follow-up question.¹⁴ Hereafter we will refer to wage "bins" as the ranges of wages from the above survey question.

A challenge working with this question is that the wage bins are different depending on the respondent's initial response to "OO2a2". To make the responses comparable between an individual's surveys, we fit various distributions to each individual's responses. The details of the fitting process is discussed in section 5.1.

The survey also included the question "OO2a" to measure an individual's expectation of the average wage offer. We do not use this question for the excess belief movement test since the follow-up question that elicits a probability distribution is necessary for the excess belief movement test, and the belief distribution is not elicited for the average wage offer. We do, however, use this response to test for asymmetric updating to determine how the expected wage offer changes depending on the wage offer they receive.¹⁵ Finally, the survey asks about the offers they received,¹⁶ as well as the number of offers expected.

One potential concern is that the beliefs from the survey are not incentivized and that this may affect the quality of the responses. Figure 3 below shows the scatter plot showing the relationship between the expected best wage offer and the highest wages offer that the survey respondents received.

 $^{^{14}}$ We are not using the response in "OO2a2" for the excess belief movement test. The role of the initial survey survey question "OO2a2" is to partition the wage distribution into different wage bins for the follow-up question.

¹⁵We use this response to examine if the number of offers affects how much people are updating their beliefs.

¹⁶The survey asks individuals to report wages of their best three offers, but it also collects information on the number of offers expected and received. Hence, we can tell if an individual received more offers than they have space to report.



Figure 3: Relationship between expectation and actual maximum wage offer. Notes: The vertical axis shows each individual's (logged) response to the initial survey question OO2a2, while the horizontal axis shows their (logged) largest received offer on the subsequent survey.

We can see that there is strong positive correlation between expected best wage offer and the highest wages offer received. This suggests that survey respondents are not providing random or inaccurate responses. On average, the survey respondents are slightly optimistic about their wage offers, and this is consistent with the findings in existing literature (Spinnewijn, 2015; Krueger and Mueller, 2016).

4 Theoretical Framework

We will now provide a descriptive model of belief dynamics in a job search context. Suppose there is an agent who is searching for a job and believes there is a set of possible wage distributions, \mathcal{F} , that he is drawing his wages from. We will index the density functions in \mathcal{F} , by a parameter θ taking values in an ordered set Θ . The agent has a non-degenerate belief g_t over θ at time t. We will also assume that the agent places a non-zero probability weight on the true wage distribution, F, he is drawing from. The agent's belief about the wage he will receive is a mixture distribution of the possible wage distributions in \mathcal{F} .¹⁷

In each period, the agent can observe a signal that reveals some information about the wage distribution he is drawing from. We let the set of possible signal realizations be $X \subseteq \mathbb{R}^n$; we can think of the signal realization $x \in X$ as a vector of wages of the job offer the agent received or any news that reveals information about his wage distribution. The conditional density function is denoted as $p(x_t|\theta)$, which reflects the likelihood of observing signal x at time t conditioned on the wage distribution being f_{θ} .

In the survey question "OO2b", the respondents provided a probability distribution about the best wage offer they will receive in the form of assigning probabilities to "binned" ranges of potential wages. We denote the probability of the wages being in bin $i \in \{1, 2, ..., n\}$ at time t as π_t^i . We will partition the wages into n wage bins $\{[a_0, a_1), [a_1, a_2), ..., [a_n, a_{n+1})\}$.¹⁸ The agent's reported belief in the survey can be represented by the expression below

$$\pi_t^i = \overbrace{\int_{\theta' \in \Theta} g_t(\theta') \underbrace{\int_{a_{i-1}}^{a_i} f(w|\theta') dw \ d\theta'}_{\text{a wage within the bin from distribution } \theta'}}$$
(1)

where w is the wage drawn from the distribution. The inner integral is the probability of receiving a wage offer within the wage bin from the distribution indexed by θ' . The outer integral integrates the agent's beliefs over which distribution the agent is drawing from.

Unlike Augenblick and Rabin (2021), where the state is an outcome (e.g. Democrats winning the US election), in our setting, the state is a probability distribution. We will need to first verify that for Bayesian agents, the reported beliefs from the mixture distribution have the martingale property to apply their test.

The martingale property requires that the expected posterior beliefs have to be

¹⁷If we are interested performing monotone comparative statics (Milgrom and Shannon, 1994; Athey, 2002), we can further assume the set of wage distributions obey the monotone likelihood ratio property in the set of belief distribution to get useful comparative statics result as shown by Li and Yu (2018).

¹⁸In the SCE survey, there are six wage bins, hence n = 6, and $a_0 = 0$ and $a_6 = \infty$. This partition covers the entire support of the wage distribution. The SCE elicits the belief distribution of the best wage offer, we can also replace the distribution of wages to the distribution of maximum wage to directly apply our model.

equal to the prior belief, $\mathbb{E}(g_{t+1}(\theta|x)|g_t(\theta)) = g_t(\theta)$. Intuitively, this property requires that the agent should not expect himself to change his belief before seeing the signal. If the agent expects himself to change his belief, he should have already done so which leads to inconsistency with the current prior.

We will now show that the beliefs π^i satisfy the martingale property if the agent's updating rule on g has the martingale property. Due to π^i_{t+1} being probability measure, $\mathbb{E}(\pi^i_{t+1}|g_t)$ converges absolutely since it has to be less than 1 and the integrands are non-negative. We can then apply Fubini's theorem and interchange the order of expectation and the integral.

$$\mathbb{E}(\pi_{t+1}^{i}|g_{t}) = \int_{\theta' \in \Theta} \mathbb{E}(g_{t+1}(\theta'|x_{t+1})|g_{t}) \int_{a_{i-1}}^{a_{i}} f(w|\theta') dw \int_{\theta' \in \Theta} g_{t}(\theta') \int_{a_{i-1}}^{a_{i}} f(w|\theta') dw = \pi_{t}^{i}$$

We see that the reported beliefs satisfy the martingale property. This allows us to perform the excess belief movement test on the statistics reported to determine if people are updating their beliefs in a Bayesian manner.

4.1 Updating Rules

We will be testing the data against some common updating rules used in the literature to determine which rules can best describe how people update their beliefs.

4.1.1 Updating Rules with Martingale Property

1. Bayesian updating

$$g_{t+1}^{bayes}(\theta|x_{t+1}) = \frac{g_t(\theta)p(x_{t+1}|\theta)}{\int_{\theta'\in\Theta} g_t(\theta')p(x_{t+1}|\theta')}$$
(2)

Bayesian updating is the standard updating rule in economics. It is the objectively correct way to update one's belief, and there are some microfoundations for Bayes' rule (Ortoleva, 2022). This updating rule has many desired properties such as the martingale property.¹⁹

¹⁹Cripps (2018), Jakobsen (2021) and Chan (2025) characterized Bayes' rule with various beliefupdating axioms or properties. These axioms and properties are normatively desirable ways of how an people should update their beliefs.

2. Affine Transformation of Bayesian Belief and prior

$$g_{t+1}^{bias}(\theta|x_{t+1}) = (1-\lambda)g_t(\theta) + \lambda g_{t+1}^{bayes}(\theta|x_{t+1})$$
(3)

The Epstein, Noor, and Sandroni (2010) updating rule is attractive for theoretical models as it preserves the martingale property, and it accommodates over- and underupdating relative to the Bayesian benchmark. $\lambda \geq 0$ is the parameter that determines the degree of over and underreaction. When $\lambda < 1$ we have underreaction and $\lambda > 1$ we have overreaction.²⁰ When $\lambda = 1$ we have the standard Bayesian updating. This model nests cursed belief (Eyster and Rabin, 2005), where the updated belief is just a convex combination between the Bayesian posterior and the prior.

4.1.2 Updating Rules without Martingale Property

3. Exponential distortion to prior and conditional probabilities

$$g_{t+1}^{bias}(\theta|x_{t+1}) = \frac{g_t(\theta)^a p(x_{t+1}|\theta)^b}{\int_{\theta' \in \Theta} g_t(\theta')^a p(x_{t+1}|\theta')^b}$$
(4)

The Grether (1980) updating rule can accommodate several updating biases but it violates the martingale property. $a \ge 0$ is the weight the agent places on the prior; when a < 1, we have base rate neglect and when a > 1, we have confirmation bias. $b \ge 0$ is the weight the agent places on the signal; when b < 1 we have underreaction to signals, and when b > 1 we have overreaction to signals. When a = 1 and b = 1we have the standard Bayes' rule.

This model is widely used in the analysis of experimental data because the odds ratio can be linearized by taking logs and estimated with a linear regression. Chan (2025) provides an axiomatic characterization of the Grether (1980) updating rule and provides experimental evidence validating this updating rule.

4. Convex combination of Bayesian belief and reference belief

$$g_{t+1}^{bias}(\theta|x_{t+1}) = (1-\lambda)\mu(\theta) + \lambda g_{t+1}^{bayes}(\theta|x_{t+1})$$
(5)

In the (Hagmann and Loewenstein, 2017) updating rule, μ is a reference belief, ²⁰There is an upper bound on λ to ensure that it is a valid probability measure. which is a belief that the agent wants to have, and $0 \le \lambda \le 1$ is a parameter that draws the updated belief towards the reference belief. This model has been use to explain motivated beliefs (Bénabou and Tirole, 2002; Eil and Rao, 2011). A key prediction from this model is that the agent will update asymmetrically: the agent over-updates when the signal moves the prior towards the reference belief and underupdates when the signal moves the prior away from the reference belief. For instance, if an individual is optimistic about the wage offers he can receive, we set μ to be an optimistic belief. When $\lambda \neq 1$, the agent will update towards this reference belief.

5 Empirical Strategy

5.1 Distribution Fitting

	w < 52,000	$52,000 \le w < 58,500$	$58,500 \le w < 65,000$	$65,000 \le w < 71,500$	$71,500 \le w < 78,000$	$w \ge \\78,000$
$p(\cdot)$	0	0.2	0.3	0.4	0.1	0

First Survey's Belief

Second Survey's Belief

	w <	$56,800 \le$	$63,900 \leq$	$71,000 \le$	$78,100 \leq$	$w \ge$
	56,800	w <	w <	w <	w <	85,200
		63,900	71,000	78,100	85,200	
$p(\cdot)$	0	0	0.7	0.2	0.1	0

Table 2: Example of a response from the SCE survey (ID: 70048914)

The main challenge in working with this data is that the wage bins vary across each individuals' survey responses. These wage bins are constructed based on percentages of the expected best wage offers reported in each survey. Since individuals may report different wage expectations in different surveys due to updating of beliefs, this leads to variations in the wage bin intervals for each response. Table 2 provides an example individual's response from the SCE survey, where the wage bins are not directly comparable. To make the bins comparable across survey responses, we fit various distributions to the second survey's belief. This allows us to estimate what the updated beliefs are in the initial survey's wage bins. The excess belief movement test can be biased by measurement errors in the prior beliefs, but not by measurement errors in the updated beliefs.²¹ Fitting the distribution to the updated beliefs in subsequent surveys minimize the bias in our test statistics.

We first convert the survey responses to a CDF and we use the Simulated Method of Moments to fit various distributions to the five points on individuals' updated beliefs from the survey questionnaire.²² We try to minimize the "distance" between the estimated probability from the fitted distribution and the actual probability. Since this process gives us an entire distribution for the updated beliefs, we can directly estimate the probability weight placed over the ranges of the updated beliefs that align with the prior wage bins defined by the survey questions by directly estimating the weight over these bins from the simulated distribution.

We present the results of fitting a log-normal distribution to the survey responses in the main text, as it yielded the lowest mean squared error compared to all other distributions we used to fit subjects' beliefs about the best wage offer.²³ The results from these alternative fitting methods can be found in appendix C of the Supplemental Appendix. Overall, our main results are robust to the different distribution fitting processes.

5.2 Martingale Test: Excess Belief Movement

5.2.1 Test Description

We will be using the excess belief movement test to test for the martingale property. The main benefit of the excess belief movement test over other martingale tests is that the test statistics are closely related to some of the more prominent belief updating biases, such as base rate neglect. This will give us insight into the type of belief updating bias to which the survey respondents are prone. Moreover, Augenblick

²¹This measurement errors are unbiased and introduced in an additive manner.

 $^{^{22}}$ The sixth bin is unbounded, and the CDF on the sixth bin should always be equal to 1.

²³Fitting the wage distribution to the recovered single wage distribution provides a slightly better fit than fitting the log-normal distribution directly to beliefs about the best wage offer. However, we choose to use the direct fit to respondents' beliefs, as recovering the single wage distribution requires additional assumptions and offers no significant improvement in fit.

and Rabin (2021) showed that their test has a higher power in detecting non-Bayesian updating compared to other existing martingale tests in their simulations.

The test involves computing two statistics: (1) belief movement, m_{t_1,t_2} , and (2) uncertainty reduction, r_{t_1,t_2} , as shown in equations 6 and 7 respectively. We will also assume $t_2 > t_1$ which denotes the time periods. Let π_t^i denote the probability assigned to bin *i* at time *t*. With 6 bins, the two statistics are defined as

$$m_{t_1,t_2} \equiv \sum_{i=1}^{6} \sum_{\tau=t_1}^{t_2-1} (\pi_{\tau+1}^i - \pi_{\tau}^i)^2$$
(6)

$$r_{t_1,t_2} \equiv \sum_{i=1}^{6} \sum_{\tau=t_1}^{t_2-1} \pi^i_{\tau} (1 - \pi^i_{\tau}) - \pi^i_{\tau+1} (1 - \pi^i_{\tau+1})$$
(7)

Both statistics have an intuitive interpretation. The belief movement is the total squared difference between beliefs in consecutive periods. This captures how much beliefs are changing regardless of the direction of the change. For uncertainty reduction, the statistics can be interpreted as a measurement of the "variance" of the belief. If we treat each bin as a Bernoulli distribution, the expression in the summation is the variance of the Bernoulli distribution in the first period minus the variance of the Bernoulli distribution in the first period minus the variance of the Bernoulli distribution in the first period minus the bins. This gives us a proxy of the amount of uncertainty in the belief distribution. If the belief updating rule satisfies the martingale property, the expected belief movement will be equal to the expected uncertainty reduction.²⁴ This means that if the agent's belief is expected to move greatly, we will expect the agent to become more certain.

The ideal test would require us to elicit the respondents' beliefs at every possible signal realization to compute the expected belief movement and uncertainty reduction for each individual. This would allow us to determine if an individual's updating rule satisfies the martingale property. Since we only observe a single updated belief based on the realized signals observed by the respondent in the last four months, we can only test if the population's updating rule satisfies the martingale property at the aggregate level. Given that there are n observations, we compute the average belief movement as $\overline{m}_{t_1,t_2} \equiv \frac{1}{n} \sum_{j=1}^{n} m_{t_1,t_2}^{j}$ and the average uncertainty reduction as $\overline{r}_{t_1,t_2} \equiv \frac{1}{n} \sum_{j=1}^{n} r_{t_1,t_2}^{j}$. We then compute the average excess belief movement statistic:

$$X = \overline{m}_{t_1, t_2} - \overline{r}_{t_1, t_2} \tag{8}$$

²⁴To show this we only need to apply the law of iterated expectations.

where "excess" refers to the amount of movement exceeding uncertainty reduction. An excess belief movement has a different interpretation depending on the amount of uncertainty reduction. For instance, with an excess belief movement of 0.01, the agent is closer to Bayesian if the uncertainty reduction is large compared to the uncertainty reduction when it is small. To allow for comparability across different studies, we also compute the normalized excess movement, which can be interpreted as the percentage of excess belief movement relative to the amount of uncertainty reduction. Under the null hypothesis that people are Bayesian, we expect $X_{norm} = 1$.

$$X_{norm} = \frac{\overline{m}_{t_1, t_2}}{\overline{r}_{t_1, t_2}} = \frac{X}{\overline{r}_{t_1, t_2}} + 1$$
(9)

If we reject the null hypothesis, we reject that the martingale property is satisfied. There are two possible interpretations of this result. Firstly, the updating rule that people use does not satisfy the martingale property. Secondly, the updating rule obeys the martingale property but people have incorrect priors, which causes the test to reject the martingale property.

To see why having a correct prior matters, consider a simple two-state model with a fully revealing signal. Suppose the correct prior that state 1 is drawn is 0.5 but the Bayesian agent holds an incorrect prior of 0.7. If this trial is repeated many times the agent will expect that 70% of the time, the posterior belief is 1 and 0 in the remaining 30% of the time. However, when the data is collected, half the time the posterior belief will be 1 and 0 otherwise, and the beliefs average to 0.5 instead of the agent's prior of 0.7. We may incorrectly conclude that the agent is non-Bayesian even when the agent is Bayesian but has a wrong prior.²⁵

5.2.2 Test Assumption: Stable Wage Distribution During Survey Period

It is expected that the worker's wage distribution will change depending on labor market conditions and the phase of the worker's career. For the test to be valid, we require the worker's wage distribution to be stable for the survey duration (up to a year). If the wage distribution changes during the survey period, the equality of the expected belief movement and uncertainty reduction will not hold. While we cannot test how each individual's offer distribution is changing, we can test whether

²⁵The average belief movement collected from the data in this example will be $\frac{1}{2}(0.7-1)^2 + \frac{1}{2}(0.7-0)^2 = 0.29$ while the average uncertainty reduction is $\frac{1}{2}(0.7)(1-0.7) + \frac{1}{2}(0.7)(1-0.7) = 0.21$. We can see that the excess belief movement is 0.08.

the aggregate distribution of all reported wage offers changes from period to period for individuals who reported offers on at least two surveys.²⁶



Figure 4: Wage offer distributions reported in the survey. Notes: The first distribution plots the distribution of offers reported in the first survey with a reported wage offer and the second distribution plots the distribution of offers reported in subsequent surveys with reported offers.

Figure 4 shows the density plot of the wage offers that were received in the first survey where an individual reported a wage offer and subsequent surveys.²⁷ We performed a non-parametric Kolmogorov–Smirnov (KS) test to determine if the distribution of the pooled wage offers received by all the survey respondents is the same across time.²⁸ We only have 124 wage offers reported on third surveys by

²⁶To test if the wage distribution is stable for each individual workers, we will need to observe multiple offers from each worker to estimate the distribution of wage offers for each worker.

²⁷Note that this differs slightly from the definition used in in table 1, where the survey 1 is the first survey where individuals answered the wage expectations questions. The analyses in this section are based on the same set of individuals, but focus on the change in reported offers rather than the change in reported expectations.

²⁸The ideal test would involve observing multiple offers from each individual to compare their offer distributions at the individual level. However, since such data is unrealistic, we can only conduct this test at the aggregate level.

respondents who received at least one offer in two previous surveys.²⁹ There were not many observations to estimate the wage distribution separately for the third survey, hence we pooled the offers reported in the second and third survey together to test this assumption. The KS test fails to reject that the distribution of wage offers received in the first survey is different from the wage offers received in the subsequent survey (p-value = 0.255).

The KS test pooled wages received by individuals in different professions together. To control for individual fixed effects, we also perform an individual fixed effect regression to test if the average wage offer of survey respondents received changed between surveys. This analysis also required an individual to report receiving job offers in at least two surveys. It is also important to note that this test is necessary but not sufficient to show that the individual wage distribution did not change in the survey period. This fixed effects regression is a complement to the KS test that we have conducted earlier.

	Wage Offer
Post Initial Survey (=1)	1 329
1 ost initial survey (-1)	(1.622)
Constant	74,374***
	(870.1)
	1.049
Observations	1,042
No. of Individuals	282
Individual Fixed Effects	Υ
R-squared	0.002

Table 3: Fixed effect regression of wage offers over survey period. Notes: The independent variable an indicator variable that takes on a value 1 for follow-up surveys (survey 2 or 3) and 0 for the first survey. Robust standard errors clustered by state in parentheses.

As shown in table 3, we see that wage offers are not statistically different across the survey period. These tests do not find evidence of wage distribution instability during the survey period for each individual.

 $^{^{29}}$ We have 483 offers in survey 1 and 435 wage offers in survey 2.

6 Results



6.1 Excess Belief Movement Test

Figure 5: Boxplot of belief movement and uncertainty reduction.

We first present our main result from the excess belief movement test. Overall we found larger belief movement relative to the amount of uncertainty reduction, as shown in the box plot in figure 5 and column 1 of table $4.^{30}$ The normalized excess belief movement of $X_{norm} = 5.18$ means that beliefs are moving 418% more than the amount of uncertainty reduction. The standard errors clustered at the state level is reported in parenthesis and we can see that X_{norm} is statistically different from 1.³¹ Excess belief movement suggests that people are over-updating relative to the

³⁰The observation counts in table 4 are slightly less than in the summary statistics table because there were 3 individuals with missing state in both survey, and we are not able to assign them to a cluster to compute the clustered standard error.

³¹The excess belief movement statistic is X = 0.7556 and it is statistically different from 0. This allows us to reject the null hypothesis that people's updating rule follows the martingale property,

Bayesian benchmark.

Statistic	All Individuals	Got (Offer?		Searched	!?
Statistic	All multiduals	Yes	No	Yes	No	Unknown
\overline{m}	.9363	1.0557	.9023	1.0126	.9012	.9576
	(.0193)	(.0437)	(.0209)	(.0368)	(.0209)	(.0579)
\overline{r}	.1807	.1986	.1756	.2010	.1677	.2179
	(.0079)	(.0190)	(.0086)	(.0134)	(.0081)	(.0259)
$X = \overline{m} - \overline{r}$.7556	.8571	.7268	.8116	.7335	.7398
	(.0209)	(.0477)	(.0243)	(.0360)	(.0220)	(.0566)
$X_{norm} = \frac{\overline{m}}{\overline{r}}$	5.1824	5.3161	5.1393	5.0374	5.3744	4.3955
,	(.2501)	(.5542)	(.3013)	(.3402)	(.2807)	(.5178)
<i>p</i> -value of t-test:						
$X = \overline{m} - \overline{r} = 0$	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
Observations	2,456	544	1,912	682	1,591	183

Table 4: Excess movement statistics: Log normal-fitted results. Notes: Clustered standard errors (at state level) in parenthesis. The standard error for X_{norm} is computed using the Delta method. Some of the subjects in our sample did not answer the question of whether they searched. We include them here as they are included in our full sample estimate.

We analyze the heterogeneous effects between respondents who received a wage offer and those who did not, as well as between those actively searching for a job and those who are not. We hypothesize that individuals who receive a job offer will update their wage expectations in a more Bayesian manner, as job offers provide personalized and direct feedback about their wage distribution. In contrast, other sources of information, such as general labor market news, offer less direct feedback. Additionally, we conjecture that individuals actively searching for a job are more likely to engage in deliberate expectation formation, making their belief updates more Bayesian compared to those who are not actively searching.

Contrary to both of our hypotheses, we find no heterogeneous effects along these dimensions, and each has a normalized excess belief movement statistic that is similar in magnitude. Figure 6 plots our test statistics across the different categorizations of individuals. The box plots show the surprising result that movement and reduction remain fairly similar across all categories.

which includes Bayesian updating.



Figure 6: Boxplot of belief movement and uncertainty reduction by categories.

6.1.1 Updating Rules that Produce Excess Belief Movement

Given the result that people are non-Bayesian, what are some updating rules that are consistent with the data? Firstly, we can reject all updating rules with the martingale property, which includes Bayesian updating and the Epstein, Noor, and Sandroni (2010) updating rule.

The Grether (1980) model as described in equation 4 can be consistent with the results obtained from this test. If we set a < 1 and/or b > 1, which allow the agent to exhibit base rate neglect or overreaction to signals, we can get excess belief movement. The intuition is that if people overweight the signal they observe or underweight their prior beliefs, this will result in over-updating in the beliefs compared to the Bayesian benchmark. This over-updating, in turn, results in excess belief movement.

The Hagmann and Loewenstein (2017) model can produce excess belief movement, but the conditions to do so are more complicated than the Grether (1980) model in a setting with multiple states or wage bins. From Augenblick and Rabin (2021), we know that the excess belief movement statistic has the following formula:³²

$$X = \sum_{i=1}^{6} \mathbb{E}[(2\pi_t^i - 1)(\pi_t^i - \pi_{t+1}^i)].$$
(10)

In this model, $\mathbb{E}\pi_{t+1}^i = (1-\lambda)\pi_{ref}^i + \lambda \pi_t^i$, where π_{ref}^i is the reference belief and $0 \le \lambda \le 1$. If we expand the above expression, we get

$$X = (1 - \lambda) \sum_{i=1}^{6} (2\pi_t^i - 1)(\pi_t^i - \pi_{ref}^i).$$
(11)

There are various combinations of priors and reference beliefs where we can obtain excess belief movement from this model.³³ The sign of the excess belief movement statistic depends on the summation. However, the parameter λ affects the magnitude of the test statistic. If λ is small, the agent puts more weight on the reference belief and less weight on the Bayesian beliefs. This makes the agent more non-Bayesian, resulting in an excess movement statistic that is larger in magnitude.

6.2 Robustness Checks

6.2.1 Measurement Error

In surveys, it is possible to potentially misreport beliefs or round off some of the probabilities. In this section, we address the concern that our primary result might be driven by this measurement error.

Consider that the theoretical framework where the agent has a true belief of π_t but reports a distorted $\hat{\pi}_t = \pi_t + \epsilon_t$, where ϵ_t is the measurement error. In a two-state model, assuming that the measurement error term is mean zero with variance σ_{ϵ}^2 and that the measurement error in the updated belief is uncorrelated with the prior, updated beliefs and error realizations ($\mathbb{E}(\epsilon_{t+1}\pi_{t+1}) = \mathbb{E}(\epsilon_{t+1}\pi_t) = \mathbb{E}(\epsilon_{t+1}\epsilon_t) = 0$), Augenblick and Rabin (2021) showed that the excess belief movement will be equal to $2\sigma_{\epsilon_t}^2$.

 $^{^{32}\}overline{\rm While}$ we use a two-period model to provide intuition, the result generalizes to more than two periods.

³³In a two-state case, $X = (1 - \lambda)(2\pi_t - 1)(\pi_t - \pi_{ref})$. We can see that we get excess belief movement in two cases: (1) $\pi_t < 0.5$ and $\pi_{ref} > \pi_t$, (2) $\pi_t > 0.5$ and $\pi_{ref} < \pi_t$. The idea is that if the expected updated belief is biased towards 0.5, we will get excess belief movement. This is also shown in Augenblick and Rabin (2021).

Generalizing this to n states, we show that the excess belief movement will be equal to $\sum_{i=1}^{n} 2\sigma_{e_i}^2$.³⁴ With measurement error, the excess belief movement is equal to the total variance of the measurement error in the prior belief multiplied by 2. This result tells us that only the measurement error in the first period will affect the excess belief movement statistics. This motivates our method of fitting the distribution to the updated beliefs instead of the prior beliefs, which minimizes the measurement error in the prior. In the calibration exercise, we only have to introduce the measurement error to the prior beliefs.

To test how much measurement error is required to rationalize our data, we perform a Monte Carlo simulation. The Monte Carlo simulation is designed to mirror the setting in our data as closely as possible. Since there are 6 wage bins and most subjects responded to the survey only twice, we have 6 states and with only 2 periods we will have only 1 signal realization and 1 update. Since we have 2,456 individuals for which we can calculate movement and reduction statistics, we generated 2,456 pairs of priors and posteriors in each simulation.

Since our statistics of interest only depend on measurement errors in the prior, we only introduce measurement errors to the prior beliefs. We draw our simulated prior data from a symmetric Dirichlet distribution that is centered around a uniform prior distribution.³⁵ Since we are working with uniform priors, the parameters of the Dirichlet distribution are equal. We then scale the parameters to adjust the variance of the distribution to match the normalized excess belief movement that we obtained.

Assuming that only measurement error is driving the results we observe, we can compute the supposed Bayesian belief movement and uncertainty reduction statistics. Based on the result we obtained, the supposed Bayesian belief movement and uncertainty reduction are 0.9363 and 0.1807, respectively, and we try to match these statistics in our calibration exercise.

It is important to note that there are infinitely many possible combinations of prior beliefs and distributions of posterior beliefs that can produce the same desired belief movement statistics. For this calibration exercise, we assume a uniform prior. The main reason for this assumption is that we can construct a uniform distribution of the posterior beliefs with symmetric posteriors, which is easier to work with. The

 $^{^{34}\}mathrm{Refer}$ to appendix A for the poof.

³⁵The draws from the Dirichlet distributions are a valid probability distribution that has to sum to 1. This approach is better than adding an error that is normally distributed which may cause the probability measure to be negative or exceed 1.

posterior belief is then selected to match the supposed Bayesian belief movement and reduction statistics.³⁶

Finally, the variance of the error term is not an intuitive measure of the amount of measurement error. To have a more intuitive measure of measurement error, we define $\hat{\pi}_1^i$ as the prior beliefs drawn from the Dirichlet distribution. The measurement error, $\Delta \equiv \sum_{i=1}^6 |\hat{\pi}_1^i - \pi_{prior}^i|$, is the total "distance" between the beliefs drawn from the Dirichlet distribution and true prior. Δ is bounded above by 2, which is a complete misreporting of beliefs.

We simulate this process 10,000 times. In Table 5, we report the average statistics computed from the 10,000 Monte Carlo simulations and the interval that contains 95% of our simulation result.

Statistics	Simulation Results	Target Values
\overline{m}	0.9365	0.9363
	[0.9185, 0.9543]	
\overline{r}	0.1772	0.1807
	[-0.2403, 0.4543]	
$X = \overline{m} - \overline{r}$	0.7556	0.7556
	[0.7333, 0.7777]	
$X_{norm} = \frac{\overline{m}}{\overline{r}}$	5.1819	5.1824
	[4.9130, 5.4695]	
Δ	1.1633	
	[1.1533, 1.1731]	

Table 5: Measurement error calibration statistics. The measurement error is generated by a Dirichlet prior and we pick Bayes' plausible distribution of posterior beliefs in order match the belief movement and uncertainty reduction statistics.

We see that to obtain an excess belief movement from our dataset assuming the respondents are Bayesian, we need the survey respondents to misreport their prior beliefs by a total of 116 percentage points across the six bins. For example, this means that survey respondents who have a belief of 70 in one of the wage bins would have to misreport the belief as 12, with the misassigned weight redistributed across the remaining wage bins. The measurement error required to rationalize the behavior is

³⁶The posterior beliefs that is selected in this calibration exercise is (0.8356, 0.1644, 0, 0, 0, 0). To obtain the full distribution of posterior beliefs, we can permute the order of the probabilities and each of these permutations will be realized with equal probability due to the uniform prior.

very large and is unlikely to be the sole explanation for the result we obtained.

The analysis here assumes all survey respondents have the same measurement error in reporting their beliefs. This assumption can be relaxed and we require the average total variance in the error to be equal to half of the excess belief movement statistics. However, Δ will be lower in the case of heterogeneous measurement error due to the concave relationship between Δ and excess belief movement. Figure 11 plots the relationship between Δ and the excess belief movement statistics from the simulation exercise that is described above.

We can see that if we allow for heterogeneous measurement error and fix the average excess movement statistics, the average Δ required to rationalize the statistics will be lower. Assuming homogeneous measurement error gives us an upper bound for Δ .

6.2.2 Two Wage Bins

It is possible that having six different wage bins could inflate the excess belief movement statistics. Most of the other studies that use the excess movement test have binary states (Augenblick and Rabin, 2021; Augenblick, Lazarus, and Thaler, 2023).³⁷ For comparability with other studies, we also perform a robustness test by collapsing the six wage bins into just two bins. We combine the first 3 wage bins, beliefs that the maximum wage offer is lower than the response in OO2a2, and the last 3 wage bins, beliefs that the maximum wage offer is greater than the response in OO2a2, into a single wage bin.

The excess belief movement statistics in our study are still significantly larger than other studies.³⁸ As shown in table 6, we see that the statistics that we got are still significantly larger than the other studies. In the next section, we provide some explanation for the large excess movement statistics that are observed in our dataset. To understand why these statistics are larger, we further examine the survey respondents' updating patterns.

³⁷Some of the other settings include expert forecasters predicting the likelihood of geopolitical events, as well as beliefs in a British prediction market

³⁸In Augenblick and Rabin (2021) the largest normalized excess belief movement was 1.2, while for Augenblick, Lazarus, and Thaler (2023) it is about 1.24 as computed from their regression estimates in Table 2 for the basketball game; the statistic from their financial data is 1.46.

Statistic	All Individuals	Got (Offer?	Searched?			
Statistic	All multiduals	Yes	No	Yes	No	Unknown	
\overline{m}	.6068	.6654	.5902	.6392	.5882	.6481	
	(.0151)	(.0412)	(.0147)	(.0325)	(.0172)	(.0633)	
\overline{r}	.1072	.1220	.1030	.1166	.1039	.1011	
	(.0063)	(.0112)	(.0060)	(.0112)	(.0065)	(.0181)	
$X = \overline{m} - \overline{r}$.4996	.5434	.4871	.5226	.4843	.5469	
	(.0151)	(.0421)	(.0154)	(.0303)	(.0173)	(.0634)	
$X_{norm} = \frac{m}{\overline{r}}$	5.6593	5.4560	5.7278	5.0374	5.3744	4.3955	
,	(.3343)	(.5906)	(.3509)	(.3402)	(.2807)	(.5178)	
p-value of t-test:							
$X = \overline{m} - \overline{r} = 0$	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	
Observations	2,456	544	1,912	682	1,591	183	

Table 6: Excess movement statistics: Two-bin results. Notes: Clustered standard errors (at state level) in parenthesis. The standard error for X_{norm} is computed using the Delta method.

6.3 Updating of Unlikely Events

The main reason for the large magnitude documented in our study is that wage offers that were thought to be unlikely in the initial survey are updated to likely offers in subsequent surveys. For a Bayesian agent, this kind of update is caused by receiving a "surprising" offer that is highly unlikely according to the prior, which then causes a large change in beliefs. Such updates should be rare if the martingale property is satisfied.³⁹

Table 7 shows an example of a survey response where an individual updates an event from unlikely to likely. In the first survey, the individual assigns 0 probability to wages above \$33,000. We interpret a zero probability response as the survey respondent thinking that the offer is unlikely rather than impossible. In the second survey, the respondent thinks that there is a 90% chance that the wage offer will exceed \$36,000, which was thought to be extremely unlikely in the first survey.

³⁹If such updating patterns occur so frequently, the initial prior is unlikely to be correct, which violates the martingale property.

First Survey's Belief

	<i>w</i> <	$24,000 \le$	$27,000 \leq$	$30,000 \le$	$33,000 \le$	$w \ge$
	24,000	w <	w <	w <	w <	36,000
		27,000	30,000	33,000	36,000	
$p(\cdot)$	0	0	0.5	0.5	0	0

Second Survey's Belief

	w <	$32,000 \leq$	$36,000 \leq$	$40,000 \le$	$44,000 \le$	$w \ge$
	32,000	w <	w <	w <	w <	48,000
		36,000	40,000	44,000	48,000	
$p(\cdot)$	0	0.1	0.4	0.4	0.1	0

Table 7: Respondent 70051617's reported beliefs

We estimate the number of instances where an unlikely wage offer is updated to a likely outcome. First, we identified wage bins in the initial survey where respondents indicated zero probability of receiving a wage offer from that wage bin. We then fit a log-normal distribution to the updated beliefs as discussed in the earlier section, and set various thresholds for probability weights that are considered likely. Figure 7 shows the relationship between the proportion of observations and the different thresholds used. For thresholds between 10% and 50%, we estimate that about 22-37% of observations update an unlikely wage offer in the initial survey into a likely wage offer in the follow-up survey.

For updating rules that multiply the conditional probabilities with the prior, like Bayes' rule and Grether (1980), we require a very informative signal or extreme biases in the latter model to update unlikely events in the prior into a likely event. These models may not be useful if we want to accommodate such updating behaviors. Alternatively, some non-Bayesian updating rules accommodate updating zero probability events or unlikely events in the prior into likely events and highly possible events in the updated beliefs (Ortoleva, 2012; Hagmann and Loewenstein, 2017; Ba, 2022). These updating rules may be more appropriate to explain these behaviors.⁴⁰

⁴⁰It is important to acknowledge that these models have several free parameters and it is difficult to falsify these models.



Figure 7: Percentage of observations where unlikely wage offers (0 probability in initial survey) is updated to likely offers. The y-axis shows the percentage of observation and the x-axis shows the different cutoffs that is used to determine if the wage offer has become "likely".

6.4 Policy Implications

Given that individuals do not update their beliefs in a Bayesian manner, is a policy intervention necessary? To answer this question, we explore whether individuals can eventually learn about their wage distribution in an information-rich environment. While certain non-Bayesian updating patterns allow individuals to eventually learn the underlying wage distribution, others do not.⁴¹ In this section, we restrict our attention to updating rules that can produce excessive belief movement that is consistent with the data.

A Bayesian agent with an incorrect prior and an agent that overreacts to signals can eventually learn about the true state with sufficient information. However, the

⁴¹Based on the definition of the belief movement and uncertainty reduction statistics, any updating rule where beliefs fail to converge will produce excess belief movement over multiple periods of updating. The excess belief movement in our data, suggests that beliefs are not converging and there may be a need for policy intervention. However, the converse is not true, an example is the Grether (1980) model with overreaction.

beliefs of an agent with base rate neglect and asymmetric updating will not converge. The intuition is that a base rate neglect agent "forgets" the prior belief and old signals which prevents belief from converging. In the case of asymmetric updating, if there is a asymmetric updating against the realized state, the agent will over-update whenever the a signal that is realized. In some extreme cases of asymmetric updating, it is also possible for the beliefs to converge to the wrong state.⁴² In appendix B of the Supplementary Appendix, we ran a simulation to demonstrate the convergence of beliefs of various updating rules using a stylized binary state updating problem.

6.4.1 Test for Convergence of Beliefs

Unfortunately, the SCE dataset is unable to provide a convincing test for the convergence of beliefs as we only observe the same person's belief up to three times. The ideal test would be to observe the same individual over a longer time horizon with more frequent surveys and test if the beliefs converge over time. The evidence should be interpreted as suggestive evidence for the belief updating biases.

	Movement	Reduction
Second Update $(=1)$	0.0147 (0.0185)	-0.0165 (0.0155)
Constant	$\begin{array}{c} (0.0220) \\ 0.627^{***} \\ (0.00923) \end{array}$	(0.00776)
Updates R-squared Individuals	$2,108 \\ 0.000 \\ 1,054$	$2,108 \\ 0.001 \\ 1,054$

Table 8: Individual fixed effect regressions of movement and reduction on sequential update number for individuals with two updates. *Notes: Robust standard errors clustered by state in parentheses.*

To test for convergence of beliefs, we focus on individuals who completed three surveys. If beliefs are converging, we would expect the belief movement and uncer-

⁴²A very simple updating rule is to disregard any signal that is suggestive of the realized state and only update from signals that are suggestive of the other state or the state the agent wants to realize. Another theoretical model that can cause convergence of belief to the wrong state is Rabin and Schrag (1999), where the agent can misinterpret the signal.

tainty reduction to decrease over time.⁴³ To test this we run an individual fixed effect regression and find that belief movement and uncertainty reduction are not significantly different across periods of survey. This suggests that base rate neglect and asymmetric updating is a likely cause of the excess belief movement. This finding supports the robust finding from lab experiments where people neglect their priors (Kahneman and Tversky, 1972; Benjamin, 2019; Esponda, Vespa, and Yuksel, 2024).

Moreover, laboratory experiments have found that the main biases in lab are base rate neglect and conservatism bias, where people under-react to new information (Benjamin, 2019). If the results from the lab can be generalized beyond the lab, then it is unlikely that overreaction to signals is causing the excess belief movement since under-reaction to signals is the main bias documented from lab experiments.

In the case of base rate neglect, providing a single piece of information is unlikely to have a long-term effect. Base rate neglect causes individuals to neglect the prior and past information. We will need to consistently provide information to correct the individuals' beliefs.

6.4.2 Asymmetric Updating Patterns

Among the 3,515 second and third survey responses, 605 of these responses, approximately 17% of the sample, reported receiving at least one job offer. For individuals who reported receiving a job offer, we have some information about the type of signals they received and how they should update their beliefs from the previous survey.⁴⁴

To identify this asymmetric updating pattern, we use the response about the average wage offer. We construct the following statistic, which we will refer to as the normalized change in expected wage.

$$\frac{y_2 - y_1}{|\overline{x}_2 - y_1|}$$

Here y_1 and y_2 are a respondent's expectation of the average wage offer they could earn over the next four months on survey 1 and 2 respectively (that is, the

⁴³Our main test calculates movement and uncertainty reduction measures using an individual's entire sequence of updates together, but for this test, we compare movement and reduction from one update to the next.

⁴⁴It is important to note that individuals could have observed other signals than wage offers that could have changed their beliefs.

response to survey question "OO2a").⁴⁵ \overline{x}_2 is the average of the offers reported by the subject on survey 2. In the numerator, we have the change in the respondents' beliefs. In the denominator, we have the absolute difference between the average offers the survey respondent received and the initial beliefs to control for the magnitude of the update.⁴⁶



Figure 8: Normalized Expected Wage Changes by Signal Direction. Belief variable here is the average annual salary the survey respondent expects. Error bars display 95% confidence intervals around the mean.

As shown in figure 8, we see that the normalized expected wage change is larger in magnitude when the survey respondent receives a positive signal compared to a negative signal. This pattern is robust to the number of job offers they receive. The updating pattern is consistent with motivated beliefs (Bénabou and Tirole, 2002; Eil and Rao, 2011), where individuals want to have an optimistic outlook on their job prospects.⁴⁷ This can potentially explain the persistence of over-optimism in the labor

⁴⁵This can be for survey 2 and 3 as well. We just need data from consecutive surveys.

⁴⁶We also perform the same analysis with the reported best wage offer. Figure 12 in appendix D of the Supplemental Appendix presents the results. The general asymmetric updating behavior is the same as the normalized expected wage change, except that the confidence intervals are wider.

 $^{^{47}}$ In a theoretical model like Hagmann and Loewenstein (2017), we set the reference belief to an

market (Jones and Santos, 2022).

Asymmetric updating provides a potential explanation to reconcile the seemingly disparate findings of beliefs being unresponsive over periods of unemployment (Krueger and Mueller, 2016; Mueller, Spinnewijn, and Topa, 2021), while Conlon, Pilossoph, Wiswall, and Zafar (2018) and our paper found general over-updating pattern behavior.⁴⁸ Studies that found underreaction typically requires downward adjustments in beliefs or reservation wage, this is consistent with our result where people are not adjusting their beliefs as much when they should update their beliefs downwards.

Finally, we do observe a surprising result in figure 8, where the normalized expected wages change are not too different across the number of offers received. Receiving multiple offers is more informative than receiving a single wage offer, we should expect more updating towards the average wage offers. One possible explanation is the "law of small numbers", where people hastily draw conclusions from a small sample size (Rabin, 2002).

The results from the excess belief movement test suggest that, on average, people are updating their beliefs more than the Bayesian benchmark. This implies that information provision policies are likely to be effective at the aggregate level. However, the presence of asymmetric updating indicates that individuals may only respond to good news, while remaining unresponsive to bad news. In such cases, information provision policies alone may not be sufficient to address the over-optimism bias. Other policy interventions that can help individuals update their beliefs in a non-bias way or more persuasive information like personal coaching will be necessary to eliminate the overoptimism in the survey respondent's beliefs.

6.5 Job Search Behavior

According to search theoretic models with uncertainty (Rothschild, 1978; Rosenfield and Shapiro, 1981; Talmain, 1992; Li and Yu, 2018; Potter, 2021), the agent's reservation value is a function of the agent's beliefs. When the agent is more optimistic about the job prospects, the agent will set a higher reservation wage. The SCE

optimistic belief and we can obtain this asymmetric updating pattern where beliefs would move upwards. If we set λ to be large enough, people can update in the wrong direction, as well.

⁴⁸It is important to note that Conlon, Pilossoph, Wiswall, and Zafar (2018) and our paper use the same dataset.

directly elicits the reservation wage of the individual.⁴⁹ Consistent with theoretical prediction, we see in figure 9 that there is a strong positive correlation between the reservation wage and the expectation of the best wage offer. This suggests that the beliefs formed by the survey respondents are likely to influence the reservation wages they set.

As mentioned in the previous section, asymmetric updating causes overoptimism to persist. Consequently, survey respondents are likely to set higher reservation wages than they would if they had accurate beliefs about their true wage distribution. This, in turn, leads to longer search durations and the rejection of offers that would have been beneficial for them to accept.



Figure 9: Relationship between wage expectations and the reservation wage. *Notes:* Each scatter point represents a survey respondent.

⁴⁹The survey question that elicits the reservation wage is labeled RW2 and it reads: "Suppose someone offered you a job today in a line of work that you would consider. What is the lowest wage or salary you would accept (BEFORE taxes and other deductions) for this job?". Before March 2017, the wage can be hourly, weekly, biweekly, monthly or yearly. Since March 2017, only the annual salary is elicited.

7 Discussion

Overall, we find strong evidence that people are non-Bayesian when updating their beliefs about their wage expectations using survey data. In the excess belief movement test, we find significantly more movement than uncertainty reduction in our data. This allows us to reject updating rules that have the martingale property, such as Bayesian updating and the Epstein, Noor, and Sandroni (2010) updating rule.

Having an excess belief movement is consistent with biases such as base rate neglect and overreaction to signals in the Grether (1980) model, as well as asymmetric updating. Our findings suggest that on average people are over-updating relative to the Bayesian benchmark. To determine the need for a policy intervention, we look at the patterns of convergence in beliefs. Although the data set we use is limited in that we can only observe at most three responses from survey respondents, we found suggestive evidence that subject beliefs are not converging, which suggests a need for policy interventions.

We also found evidence of asymmetric updating. This updating pattern may hinder individuals from learning about their wage distribution and cause overoptimism in beliefs to persist despite being in an information-rich environment. This insight also provides policy implications on how good news and bad news should be conveyed to individuals and the effectiveness of such news in changing people's beliefs. When providing information that can cause individuals to adjust their wage expectations downward, we may need a stronger or more informative signal.

Our results are limited in two important ways that leave scope for future work. Firstly, the current dataset does not allow us to effectively test for the non-convergence of beliefs. Understanding if beliefs could eventually converge to the actual wage distribution will give us insights into whether a policy intervention is needed to correct those beliefs and how to best structure information policies. Future studies could collect higher-frequency data, allowing a better study of how beliefs change over time to detect patterns of convergence.

Secondly, our sample chiefly includes employed individuals, so it is difficult for us to comment on the differences between employed and unemployed learning, with our analysis mostly applying to learning among employed individuals. The beliefs of unemployed individuals are likely to have the largest effect on their welfare as their beliefs will affect their decision to accept a job offer. It will be useful to have a survey where we can primarily focus on unemployed individuals. This will help us design a more effective policy to reduce unemployment.

We believe that the updating patterns documented in this paper have important implications for both behavioral theory and policy. While, the dominant approach in most learning models is to assume that people are Bayesian due to tractability of model, our findings offer valuable insights into how theoretical models can be refined to better capture the dynamics of belief updating. Given that we observe significant over-updating relative to the Bayesian benchmark, our results suggest that information provision could serve as a powerful policy tool for influencing individuals' beliefs and potentially their decision-making processes.

Appendix A Proofs

We will prove that under some assumptions, measurement errors in the prior beliefs will produce positive excess belief movement (Augenblick and Rabin, 2021). Suppose the survey respondent report a distorted prior of $\hat{\pi}_t = \pi_t + \epsilon_t$, where ϵ_t is the measurement error. In a two-state model, assuming that the measurement error term is mean zero with variance σ_{ϵ}^2 and uncorrelated with beliefs and past errors ($\mathbb{E}(\epsilon_t \pi_t) = \mathbb{E}(\epsilon_t \pi_{t-1}) = \mathbb{E}(\epsilon_t \epsilon_{t-1})$). To simplify our notation, we will use \mathbb{E} to represent the expectation of the updated beliefs and the measurement error.

The belief movement will be larger by σ_{ϵ}^2

$$\mathbb{E}(M_{t,t+1}) = \mathbb{E}(\pi_{t+1} + \epsilon_{t+1} - \pi_t - \epsilon_t)^2$$

= $\mathbb{E}[(\pi_{t+1} - \pi_t)^2 + 2(\pi_{t+1} - \pi_t)(\epsilon_{t+1} - \epsilon_t) + (\epsilon_{t+1} - \epsilon_t)^2]$
= $\mathbb{E}[(\pi_{t+1} - \pi_t)^2 + \epsilon_{t+1}^2 - 2\epsilon_{t+1}\epsilon_t + \epsilon_{t+1}^2]$
= $\mathbb{E}[(\pi_{t+1} - \pi_t)^2] + \sigma_{\epsilon_t}^2 + \sigma_{\epsilon_{t+1}}^2$

The uncertainty reduction will be smaller by σ_{ϵ}^2

$$\mathbb{E}(R_{t,t+1}) = \mathbb{E}[(\pi_t + \epsilon_t)(1 - \pi_t - \epsilon_t) - (\pi_{t+1} + \epsilon_{t+1})(1 - \pi_{t+1} - \epsilon_{t+1})]$$

= $\mathbb{E}[(\pi_t)(1 - \pi_t - \epsilon_t) + \epsilon_t(1 - \pi_t - \epsilon_t) - (\pi_{t+1})(1 - \pi_{t+1} - \epsilon_{t+1})]$
+ $\epsilon_{t+1}(1 - \pi_{t+1} - \epsilon_{t+1})]$
= $\mathbb{E}[(\pi_t)(1 - \pi_t) - \pi_{t+1}(1 - \pi_{t+1})] - \sigma_{\epsilon_t}^2 + \sigma_{\epsilon_{t+1}}^2$

The expected excess belief movement statistic for a Bayesian agent with measurement error is

$$\mathbb{E}(M_{t,t+1}) - \mathbb{E}(R_{t,t+1}) = 2\sigma_{\epsilon_t}^2$$

If we generalize this to n states,

$$\mathbb{E}(M_{t,t+1}) = \mathbb{E}\left[\sum_{i=1}^{n} (\pi_{t+1}^{i} + \epsilon_{t+1}^{i} - \pi_{t}^{t} - \epsilon_{t}^{i})^{2}\right]$$
$$= \mathbb{E}\left[\sum_{i=1}^{n} (\pi_{t+1}^{i} - \pi_{t}^{i})^{2} + 2(\pi_{t+1}^{i} - \pi_{t}^{i})(\epsilon_{t+1}^{i} - \epsilon_{t}^{i}) + (\epsilon_{t+1}^{i} - \epsilon_{t}^{i})^{2}\right]$$
$$= \mathbb{E}\left[\sum_{i=1}^{n} (\pi_{t+1}^{i} - \pi_{t}^{i})^{2}\right] + \sum_{i=1}^{n} (\sigma_{\epsilon_{t}}^{i})^{2} + (\sigma_{\epsilon_{t+1}}^{i})^{2}$$

$$\mathbb{E}(R_{t,t+1}) = \mathbb{E}\left[\sum_{i=1}^{n} (\pi_t^i + \epsilon_t^i)(1 - \pi_t^i - \epsilon_t^i) + (\pi_{t+1}^i + \epsilon_{t+1}^i)(1 - \pi_t^i - \epsilon_{t+1}^i)\right]$$

$$= \mathbb{E}\left[\sum_{i=1}^{n} (\pi_t^i)(1 - \pi_t^i - \epsilon_t^i) + \epsilon_t^i(1 - \pi_t^i - \epsilon_t^i) - (\pi_{t+1}^i)(1 - \pi_{t+1}^i - \epsilon_{t+1}^i) + \epsilon_{t+1}^i(1 - \pi_{t+1}^i - \epsilon_{t+1}^i)\right]$$

$$= \mathbb{E}\left[\sum_{i=1}^{n} (\pi_t^i)(1 - \pi_t^i) - \pi_{t+1}^i(1 - \pi_{t+1}^i)\right] + \sum_{i=1}^{n} (\sigma_{\epsilon_{t+1}}^i)^2 - (\sigma_{\epsilon_t}^i)^2$$

The excess belief movement is

$$\mathbb{E}(M_{t,t+1}) - \mathbb{E}(R_{t,t+1}) = \sum_{i=1}^{n} 2(\sigma_{\epsilon_t}^i)^2$$

References

- Ahn, H. J. and L. Shao (2021). The cyclicality of on-the-job search effort. *The B.E. Journal of Macroeconomics* 21(1), 185–220.
- Armantier, O., G. Topa, W. Van der Klaauw, and B. Zafar (2017). An overview of the survey of consumer expectations. *Economic Policy Review*, 51–72.
- Arni, P. P. (2016). Opening the blackbox: How does labor market policy affect the job seekers' behavior? a field experiment. *IZA Discussion Paper Series*.
- Athey, S. (2002). Monotone comparative statics under uncertainty. The Quarterly Journal of Economics 117(1), 187–223.
- Augenblick, N., E. Lazarus, and M. Thaler (2023). Overinference from weak signals and underinference from strong signals. arXiv preprint arXiv:2109.09871.
- Augenblick, N. and M. Rabin (2021). Belief movement, uncertainty reduction, and rational updating. *The Quarterly Journal of Economics* 136(2), 933–985.
- Ba, C. (2022). Robust misspecified models and paradigm shift.
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. Handbook of Behavioral Economics: Applications and Foundations 1 2, 69–186.
- Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review* 110(9), 2748–82.

- Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance* 74(6), 2839–2874.
- Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. The Quarterly Journal of Economics 117(3), 871–915.
- Chan, A. (2022). Discrimination against doctors: A field experiment. Unpublished Manuscript.
- Chan, K. (2025). An axiomatic model and test of grether (1980) and bayes' rule. Unpublished Manuscript.
- Coffman, L. C., C. R. Featherstone, and J. B. Kessler (2024). A model of information nudges. Technical report, National Bureau of Economic Research.
- Conlon, J. J., L. Pilossoph, M. Wiswall, and B. Zafar (2018). Labor market search with imperfect information and learning. Technical report, National Bureau of Economic Research.
- Cooper, M. and P. Kuhn (2020). Behavioral job search. Handbook of Labor, Human Resources and Population Economics, 1–22.
- Cripps, M. W. (2018). Divisible updating. Manuscript, UCL.
- DeGroot, M. H. (1974). Reaching a consensus. Journal of the American Statistical association 69(345), 118–121.
- DellaVigna, S., A. Lindner, B. Reizer, and J. F. Schmieder (2017, 05). Reference-Dependent Job Search: Evidence from Hungary*. The Quarterly Journal of Economics 132(4), 1969–2018.
- Eil, D. and J. M. Rao (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics* 3(2), 114–138.
- Epstein, L. G., J. Noor, and A. Sandroni (2010). Non-bayesian learning. The BE Journal of Theoretical Economics 10(1).
- Esponda, I., E. Vespa, and S. Yuksel (2024). Mental models and learning: The case of base-rate neglect. *American Economic Review* 114(3), 752–782.
- Eyster, E. and M. Rabin (2005). Cursed equilibrium. *Econometrica* 73(5), 1623–1672.
- Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2022). Job search behavior among the employed and non-employed. *Econometrica* 90(4), 1743–1779.
- Gee, L. K. (2018). The more you know: Information effects on job application rates in a large field experiment. *Management Science* 65(5), 2077–2094.

- Golub, B. and M. O. Jackson (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* 2(1), 112–149.
- Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly journal of economics* 95(3), 537–557.
- Haaland, I. and C. Roth (2023). Beliefs about racial discrimination and support for pro-black policies. *Review of Economics and Statistics* 105(1), 40–53.
- Hagmann, D. and G. Loewenstein (2017). Persuasion with motivated beliefs. In Opinion Dynamics & Collective Decisions Workshop.
- Jäger, S., C. Roth, N. Roussille, and B. Schoefer (2024). Worker beliefs about outside options. *The Quarterly Journal of Economics*, 1505–1556.
- Jakobsen, A. M. (2021). Coarse bayesian updating. Work. Pap., Northwest. Univ., Evanston, IL.
- Jiang, M. and K. Zen (2023). Information asymmetry in job search.
- Jin, G. Z., M. Luca, and D. Martin (2021). Is no news (perceived as) bad news? an experimental investigation of information disclosure. *American Economic Journal: Microeconomics* 13(2), 141–173.
- Jones, S. and R. Santos (2022). Can information correct optimistic wage expectations? evidence from mozambican job-seekers. *Journal of Development Economics* 159, 102987.
- Kahneman, D. and A. Tversky (1972). On prediction and judgement. ORI Research monograph 1(4).
- Krueger, A. B. and A. I. Mueller (2016). A contribution to the empirics of reservation wages. American Economic Journal: Economic Policy 8(1), 142–79.
- Kudlyak, M., D. Lkhagvasuren, and R. Sysuyev (2014). Systematic job search: New evidence from individual job application data. Technical Report 12-03r, Federal Reserve Bank of Richmond.
- Li, S. and N. N. Yu (2018). Context-dependent choice as explained by foraging theory. Journal of Economic Theory 175, 159–177.
- McCall, J. J. (1970). Economics of information and job search. The Quarterly Journal of Economics, 113–126.
- Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica:* Journal of the Econometric Society, 157–180.

- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380–391.
- Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. The American Economic Review 60(5), 847–862.
- Mueller, A. I., J. Spinnewijn, and G. Topa (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review* 111(1), 324–363.
- Ortoleva, P. (2012). Modeling the change of paradigm: Non-bayesian reactions to unexpected news. *American Economic Review* 102(6), 2410–2436.
- Ortoleva, P. (2022). Alternatives to bayesian updating. Annual Review of Economics 16.
- Phillips, L. D. and W. Edwards (1966). Conservatism in a simple probability inference task. Journal of experimental psychology 72(3), 346.
- Potter, T. (2021). Learning and job search dynamics during the great recession. Journal of Monetary Economics 117, 706–722.
- Rabin, M. (2002). Inference by believers in the law of small numbers. The Quarterly Journal of Economics 117(3), 775–816.
- Rabin, M. and J. L. Schrag (1999). First impressions matter: A model of confirmatory bias. The quarterly journal of economics 114(1), 37–82.
- Rosenfield, D. B. and R. D. Shapiro (1981). Optimal adaptive price search. *Journal* of *Economic Theory* 25(1), 1–20.
- Rothschild, M. (1978). Searching for the lowest price when the distribution of prices is unknown. In *Uncertainty in Economics*, pp. 425–454. Elsevier.
- Sadish, D., A. Adhvaryu, and A. Nyshadham (2021). (mis) information and anxiety: Evidence from a randomized covid-19 information campaign. *Journal of Develop*ment Economics 152, 102699.
- Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. Journal of the European Economic Association 13(1), 130–167.
- Talmain, G. (1992). Search from an unkown distribution an explicit solution. Journal of Economic Theory 57(1), 141–157.
- Thaler, R. H. and C. R. Sunstein (2008). *Nudge*. New Haven, CT and London: Yale University Press.

- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica:* Journal of the Econometric Society, 641–654.
- Wiswall, M. and B. Zafar (2015). Determinants of college major choice: Identification using an information experiment. *The Review of Economic Studies* 82(2), 791–824.