

How Much Can I Make? Insights on Belief Updating in the Labor Market

Sebastian Brown*, Kenneth Chan†

March 7, 2024

Abstract

We use a nationally representative survey (the labor supplement of the Survey of Consumer Expectations) to study individuals' behavioral biases in learning about their wage distributions. Using a recently developed test for Bayesianism ([Augenblick and Rabin, 2021](#)), we find strong evidence of non-Bayesian learning. Among respondents who responded at least twice to the survey, we find an average movement in beliefs that is roughly 518% of the reduction in their beliefs' uncertainty, 418% more than the Bayesian benchmark, a result consistent with base-rate neglect and/or overreaction to signals. This estimate remains large and significant when examining subsets of the data separated by different labor market and demographic variables. We further examine the heterogeneity in agents' learning by identifying specific non-Bayesian patterns that can be explained by multiple priors models and updating rules that accommodate asymmetric updating.

JEL: D83, D84, D90, J64

*Department of Economics, UC Santa Barbara. sebastian_brown@ucsb.edu.

†Department of Economics, UC Santa Barbara. kenneth_chan@ucsb.edu.

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 strategy 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 (Rothschild, 1978; Rosenfield and Shapiro, 1981; Talmain, 1992; Li and Yu, 2018), where the agent has to learn about the underlying distribution over time by observing the offers they receive.¹ While these models are more realistic, the empirical application of these models has been limited due to insufficient high-quality data on job searchers' beliefs. Consequently, there has been a lack of research on how job seekers learn and update their beliefs in light of new information. Examining the process of learning and belief updating can provide valuable insights into the effectiveness of information provision policy in a job search setting. Our paper hopes to fill this gap and provide some insights into how job searchers incorporate new information into existing beliefs.

Moreover, providing information has been shown to be a potent policy tool for changing people's beliefs and behavior. In a series of recent field experiments, providing information has been shown to help students make more informed decisions about their college majors (Wiswall and Zafar, 2015), improve people's knowledge about COVID-19 (Sadish et al., 2021), eliminate statistical racial discrimination in a patient's choice of medical professional (Chan, 2022), and reduce disagreement about the extent of racial discrimination (Haaland and Roth, 2023). In a job search context, Arni (2016) found that a coaching program was successful in increasing job-finding rates among the treated job searchers by 9 percent.²

¹These models are 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 extended easily to a job search context.

²We interpret this coaching program as a form of providing information to help the job searcher.

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 about the number of workers applying for a specific job. She showed that this additional labor market information increased the probability of a worker completing a job application by 2.5 percent. Overall, these studies highlighted have shown the effectiveness of information treatments in affecting people’s behavior in various settings.

Unlike most of the above papers which use a treatment intervention to manipulate people’s beliefs, in this paper, we utilize a government dataset (the Survey of Consumer Expectations) to study an individual’s capacity to learn about the wage distribution without any experimental manipulation. As our data is representative and free of the cognitive influences of an experimental setting, it is also a more externally valid measure of how individuals process information in a labor market context.

Our main analysis employs a recently developed test for Bayesian updating in beliefs from [Augenblick and Rabin \(2021\)](#). There are two statistics required for this test. Belief movement which is the squared difference of changes in belief, and uncertainty reduction which is the amount of reduction in the belief’s variance from updating. If the agent is Bayesian, the average movement in beliefs will be equal to the average uncertainty reduction. Using this test, we find an average movement in beliefs that is over five times the average reduction in uncertainty. We can reject that people are updating their beliefs in a Bayesian manner, and our results suggest that on average people overupdate relative to the Bayesian benchmark. This is consistent with biases such as overreaction to signals and base rate neglect.

Our remaining analysis attempts to examine specific behavioral patterns that are inconsistent with the Bayesian prediction. With these results, we are able to distinguish between various non-Bayesian updating rules without the martingale property ([Grether, 1980](#); [Ortoleva, 2012](#); [Hagmann and Loewenstein, 2017](#); [Ba, 2022](#)) to determine which model can best describe how job searchers update their beliefs.

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 a few ways. Firstly, our paper is primarily interested in using the data to test various updating rules and distinguish between them, while their paper primarily focuses on how information friction affects job search behavior. Secondly, our approach uses the [Augenblick and Rabin \(2021\)](#) test which circumvents the need to compute the Bayesian benchmark to compare if people are Bayesian. This test also allows us to test for Bayesianism with fewer and more conservative assumptions. 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, and they only focus on respondents who received at least one job offer. Our analysis does not rely on this parametric assumption, we only assumed log-normality in the wage distribution to compare our data across time as described in Section 5. Thirdly, our approach allows for updating without a job offer, as people may change their beliefs with news about the labor market. This allows us to use a larger sample from the SCE data, allowing us to extrapolate our results from respondents without a job offer. This makes our result Lastly, we build on this paper by conducting additional *individual-level* analyses that allow us to uncover additional details about labor market learning. Specifically, we measure and discuss the prevalence of various learning patterns in the data.

The remainder of the paper proceeds as follows. Section 2 reviews relevant literature. Section 3 describes the data and our process of fitting distributions to the data. Section 3 outlines a simple theoretical model of labor market updating and briefly describes the general theory of the Augenblick-Rabin test for Bayesianism. Section 5 explains our empirical strategy, detailing the steps that we took to apply the Augenblick-Rabin test to this dataset. Section 6 then lists our main results. Section 7, shows that our result is robust to measurement errors. In section 8, we look at other belief updating patterns at the individual level. Finally, section 9 concludes and provides suggestions for future work.

2 Literature Review

Our work is related to two strains of literature: search theoretic models with unknown distributions and non-Bayesian updating. In these models, the underlying distribution from which offers are drawn, while fixed, is also unknown to the agent.³ A key characteristic of this class of model is that the agent has to update his beliefs about the underlying distribution as he searches.

[Rothschild \(1978\)](#) wrote the first search model with unknown distribution, and he showed that the static reservation value no longer applies in most scenarios as the agent will continuously update his expectations about the future. He also pointed out that having the ability to recall a previously rejected offer will affect the optimal search strategy when the distribution is unknown. When recall is not possible, the optimal strategy can be characterized by reservation value that is a function of the agent’s beliefs ([Rothschild, 1978](#); [Rosenfield and Shapiro, 1981](#)). When recall is possible, the strategy is to set a decreasing sequence of reservation values ([Talmain, 1992](#)).

For tractability of the model, the agent is typically assumed to have a Dirichlet prior on the wage distribution ([Rothschild, 1978](#); [Rosenfield and Shapiro, 1981](#); [Talmain, 1992](#)) or a Gamma distribution on the arrival rate of offers [Potter \(2021\)](#). This is because the Dirichlet and Gamma distribution has a conjugacy property⁴ which provides a closed-form expression for the posterior. Recently, [Li and Yu \(2018\)](#) deviated from the conjugate prior assumption and they assume the agent believes that there is a set of possible distributions and all of the distributions in this set satisfy the monotone likelihood ratio property (MLRP). The MLRP allowed them to establish first-order stochastic dominance between distributions and posterior beliefs, which allows them to perform monotone comparative statics ([Milgrom and Shannon, 1994](#)) on the reservation value.

³An alternative way to model this is to let the arrival rate of offers be unknown, see [Potter \(2021\)](#). However, assuming the underlying distribution is unknown is the main way of introducing uncertainty into the model.

⁴Under Bayesian updating, the posterior belief and the prior beliefs are from the same parametric family of distribution.

While these search models typically assume people are updating their beliefs in a Bayesian manner, behavioral economists have shown that people are not Bayesian. Behavioral theorists have come up with belief updating models to accommodate these non-Bayesian behavior (Grether, 1980; Epstein, Noor and Sandroni, 2010; Ortoleva, 2012; Hagmann and Loewenstein, 2017; Ba, 2022). Some of the more well-known biases include base rate neglect (Kahneman and Tversky, 1972), 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 is from lab experiments (see Benjamin (2019) for a survey of the experimental literature on belief updating), 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 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.

Contrary to Mueller, Spinnewijn and Topa (2021) findings, Conlon, Pilossoph, Wiswall and Zafar (2018) use the SCE data and find that people overupdate their 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 focusing on wage expectation, Potter (2021) examines how the expectation of offer arrival rate changes over time using data from the Great Recession. Using a cali-

brated model, he showed that learning can explain the job search dynamics during the Great Recession.

Aside from these papers, our work also relates more generally to a broader, yet still relatively new, literature on behavioral job search (e.g. [DellaVigna et al. \(2017\)](#); see [Cooper and Kuhn \(2020\)](#) for a recent review), as well as more general work studying employed job search (e.g. [Faberman et al. \(2022\)](#), [Ahn and Shao \(2021\)](#)).

3 Data

3.1 Data Source Description

Our primary data source is the Survey of Consumer Expectations (SCE), which is a nationally representative survey administered monthly by the New York Federal Reserve. 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. Subjects can remain in the survey for up to a year, and are replaced on a rolling basis. Therefore, we are able to observe the same subject for up to a maximum of three times for the labor force questions and twelve times for the core questions. For a more in-depth description of the structure and administration of the SCE, see [Armantier et al. \(2017\)](#).

The main advantage of this dataset is that it explicitly elicits expectations from survey respondents, meaning there is no need to indirectly infer beliefs as in many previous papers (for example [Potter \(2021\)](#) and [Spinnewijn \(2015\)](#)). Another benefit is its panel structure, which, while short, still does allow for within-subject observations of the evolution in beliefs over time. A third benefit is that it features mostly employed individuals, allowing study of a previously-understudied class of workers (as highlighted in [Faberman et al. \(2022\)](#)).

This last benefit, however, comes at the cost of including fewer unemployed individuals.

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)

Figure 1: Survey Question oo2b from SCE

While 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. As noted in table 1 in section 3.3, only 3.6% of our final sample is unemployed.⁵

3.2 Survey Question Description

Our analysis uses a few survey questions which merit more detailed explanations. First, we use the distributional question labeled “oo2b” as shown in Figure 1 to measure the mean and variance of an individual’s prior.

Here the label “oo2a2” refers to the following question: “*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?*” Therefore, the categories of this question may be taken to mean, “What is the probability your maximum offer will be 80% or less of your guess of the average maximum offer?”, “What is the probability your maximum offer will be 80% to 90% of your guess of the average maximum offer?”, and so on.

Hereafter we will refer to wage “bins,” by which we mean the ranges of wages implied by the survey questions. These range from bin 1, the lowest bin, consisting of all wages 80%

⁵It is therefore not surprising that [Mueller et al. \(2021\)](#) augments the SCE by including a second dataset, the Survey of Unemployed Workers in New Jersey, which as its name suggests, focuses entirely on unemployed individuals.

times the expected best offer and below, to bin 6, the highest bin, consisting of all wages 120% times the expected best offer and above. A challenge working with this question is that the bins of wages are redefined with each new survey taken, including within subjects. The size of the wage bins is determined by their response to “oo2a2,” which differs across every response to the survey. To make the responses comparable between an individual’s consecutive surveys, we fit a log-normal distribution to each individual’s responses. This fitting also gives us the mean and variance used in our principal analysis for each individual distribution. For an explanation of this fitting process, see section 5.1.

While our main analysis uses the mean from the fitted log-normal distribution, we also use the question “oo2a” to measure an individual’s expectation of the average wage for a given period to graph the sample distribution of expected offers in figure 2 in the next section. This question reads “*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?*” We do not use this question for our main analysis since unlike oo2a2, it does not have a corresponding question (oo2b) to elicit a probability distribution. Currently, we only use it to perform an alternate version of our asymmetry analysis from section 8.2. This alternate analysis can be found in appendix C.

Finally, we also have questions about the wages and number of offers received, as well as a question on the number of offers expected. The survey asks individuals to report their three best 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. While these questions are not used directly in our main analyses, we still present them as part of our data description in the following section. The expected number of offers is also used as part of an alternate analysis we present in appendix B.1.

3.3 Data Structure and Summary Statistics

Our final dataset is organized such that each observation contains information from two consecutive labor market surveys for a single individual. While most individuals in the dataset appear only twice, for those that appear three times, there will therefore be two separate observations: one for the update from period 1 to period 2, and one for the update from period 2 to period 3.⁶

The composition of our final dataset by employment status (defined as the employment status of the individual in each before-after pair’s first period) and whether they received at least one wage offer is given Table 1:

Count	All Individuals	Got Offer?		Searched?		
		Yes	No	Yes	No	Unknown
Total Updates	4,301	731	3,570	1,040	2,924	337
Number unemployed	156	45	111	82	66	8
Number employed	3,587	625	2,962	894	2,375	318
Number not in labor force	518	55	463	46	465	7
Missing employment status	40	6	34	18	18	4
Unique individuals	3,053	674	2,379	867	1,967	219
Data Range	3/2015-3/2020					

Table 1: Dataset composition. Observations consist of information from two consecutive surveys for a given individual.

For the unique individual counts, an individual was counted as having received an offer or having searched if they did so for at least one of their updates. Some individuals left the question blank for the search question. If the individual had two updates and answered the question for only one, this answer was used to proxy for whether the individual had searched or not at any point in the “unique individuals” line (although for individual updates in the “total updates” line, a blank response was classified as “unknown”). Observations reporting annual wages or expected annual wages (using either question oo2a or oo2a2 for either period) below \$10,000 were dropped, as it was assumed these responses were mistaking the question

⁶This format allows us to analyze updates separately in section 8. For the analysis in section 6 which uses an individuals’ entire sequence of updates, We add the period 3 information to the observation covering the update from period 1 to 2 where relevant and instead use only one observation per individual.

for asking for hourly or monthly wages.⁷

In the following graphs, we report the expected and realized wage offers, both in the wage amount and number of offers. One concern is that the beliefs from the survey are not incentivized and that this may affect the quality of the responses. In figure 2 below, we see that in aggregate, the responses are close to the average offers the survey respondents reported receiving (although respondents are slightly optimistic about their wage offers). This suggests that at the aggregate level, people are not randomly responding to the survey questions. We can also see that the aggregate log distributions for expected offers and reported offers seem to follow a bell curve, consistent with our assumptions of log normality.

⁷558 updates were dropped this way, 192 of which had a reported wage or one of the expectation questions less than \$100. Individuals who reported having received a wage of 0 (5 updates) were kept in the dataset and classified as having not received an offer, since there was no evidence that this misunderstanding of the instructions would impact their ability to answer the expectations questions. There were no individuals who reported an expectation of a zero wage for either expectations question.

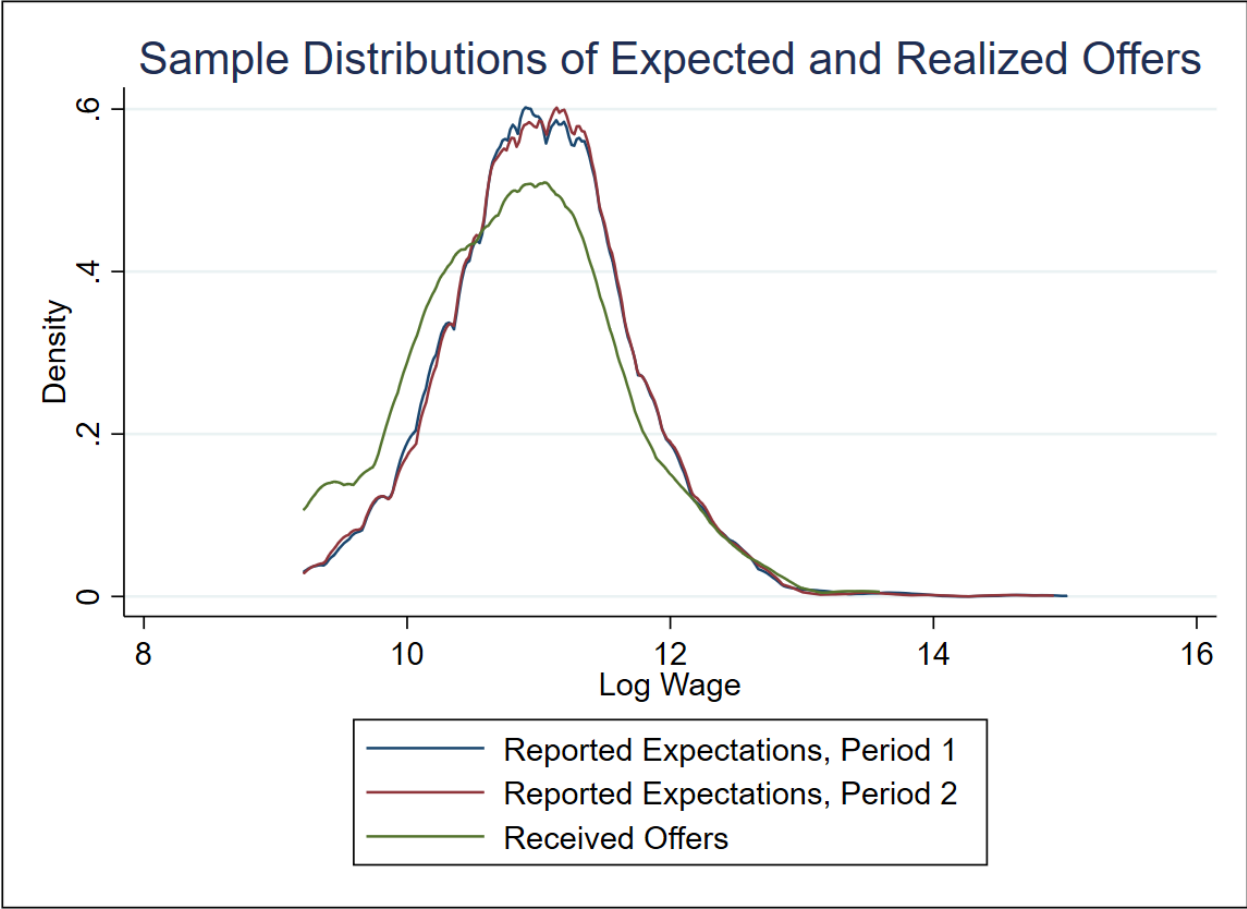


Figure 2: Distribution of reported expected wages, fitted expected wages, and received wages. Reported expected wages are based on question oo2a2 and include all individuals, including those who did not receive an offer.

Figure 3 (below) shows that individuals tend to overestimate how many offers they will receive. This graph groups observations with more than ten offers into one category, of which there are very few. In addition to the information on the number of offers presented in the graph, we also note that the number of individuals who are unable to report all received offers is relatively small: Of the 731 observations in our final dataset with realized offers, 457 reported one offer, 163 reported two offers, and 111 reported three offers. Of those 111, 54 reported receiving more than three offers (the maximum number of wages that can be reported).

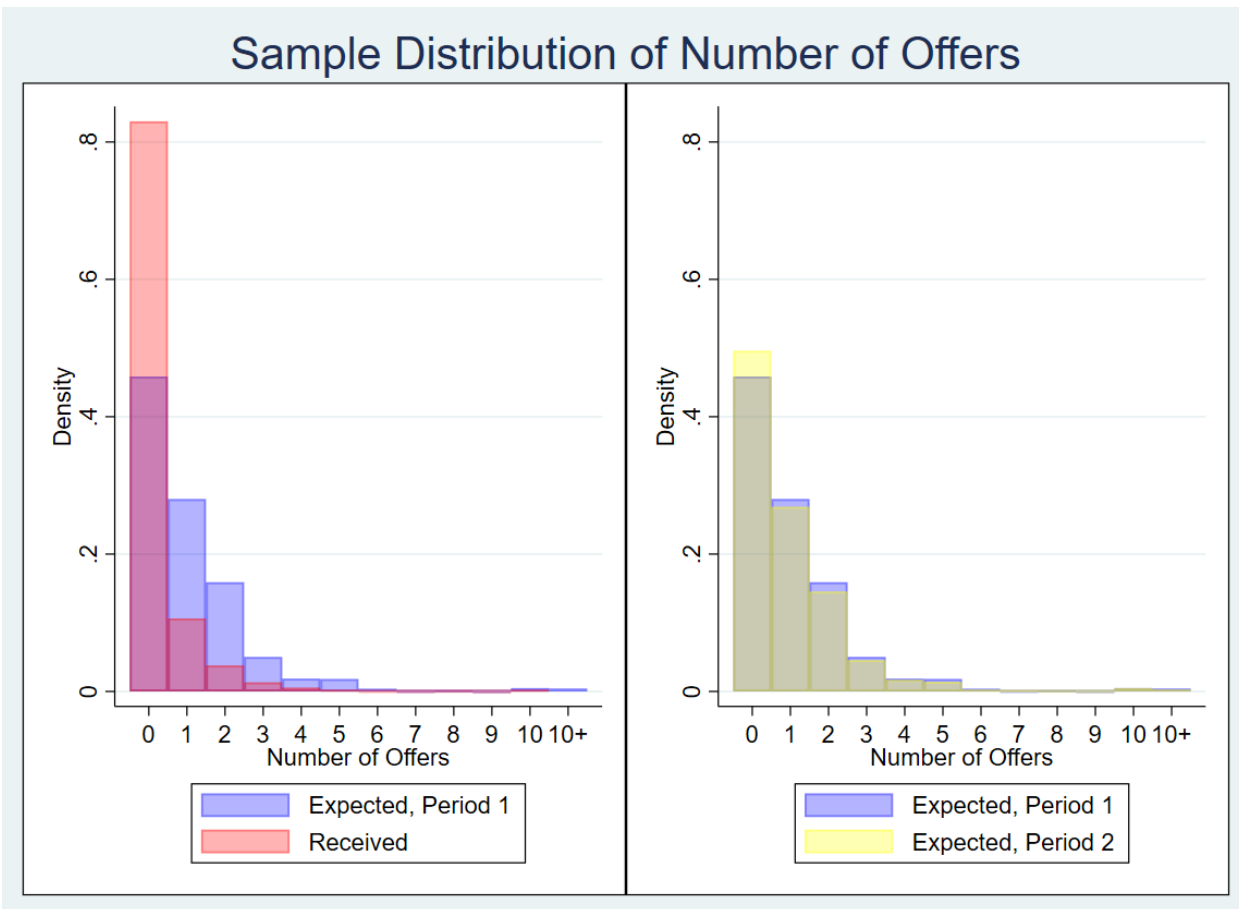


Figure 3: Distribution of expected and realized wage offers. Expected offers include the expectations of all individuals, whether or not they received an offer. There were only 18 observations in the final sample for which individuals expected more than ten offers for period 1 and 13 observations for period 2. No observations in the final sample reported receiving more than ten offers.

4 Theoretical Model

We will now provide a model of belief dynamics related to the job search setting we are studying. We will follow [Li and Yu's \(2018\)](#) method of modeling uncertainty. In their model, the agent believes there is a set of distribution he is drawing from and the distributions from this set obeys the monotone likelihood ratio (MLRP) ordering. Unlike their model which is interested in obtaining comparative statics across different signal realizations, we are primarily interested in the belief dynamics, hence we do not require the assumption

about MLRP.

This model also differs from the [Augenblick and Rabin \(2021\)](#) setting. In their setting, the state of the world is an outcome, while in our setting, the state of the world is a probability distribution. We assume that the agent’s reported belief in the SCE survey is a mixture distribution. We show that their results generalize to this setting as well. 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.

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} .

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 $i \in \{1, 2, \dots, n\}$ at time t as π_t^i .⁸

Under the model we have constructed, we will partition the wages into n wage bins $\{[a_0, a_1), [a_1, a_2), \dots, [a_n, a_{n+1})\}$. In the survey, $a_0 = 0$ and $a_{n+1} = \infty$. This partition covers the entire support of the wage distribution. We will assume that the agent’s reported belief can

⁸In the survey, there are 6 bins, but we are keeping the model general.

be represented by the expression below

$$\pi_t^i = \int_{\theta' \in \Theta} g_t(\theta') \int_{a_{i-1}}^{a_i} f(w|\theta') dw d\theta' \quad (1)$$

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 over the agent's beliefs about the distribution he is drawing from.

Our primary test tests if people's belief follows the martingale property. Formally, the martingale property states that the expected posterior beliefs have to be equal to the prior belief, $\mathbb{E}(g_{t+1}(\theta|x)|g_t(\theta)) = g_t(\theta)$. Intuitively, this property states 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.

If the agent's updating rule on g has the martingale property, the probabilities π^i will have the martingale property as well.⁹

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

This allows us to perform a martingale 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 updating rules to determine which rules can best explain people's behavior.

⁹We can easily verify that $\mathbb{E}(\pi_{t+1}^i | g_t)$ converges absolutely since it has to be less than 1 and the integrands are non-negative. We can apply Fubini's theorem and interchange the order of expectation and the integral.

4.1.1 Updating Rules with Martingale Property

1. Bayesian updating

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

Bayesian updating is the standard updating rule in microeconomic theory. This updating rule has many desired properties such as the martingale property, which makes models tractable.¹⁰

2. Affine Transformation of Bayesian Belief and prior (Epstein, Noor and Sandroni, 2010)¹¹

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

The Epstein, Noor and Sandroni (2010) updating rule is attractive as it preserves the martingale property, and yet it accommodates over- and under-updating 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.¹² When $\lambda = 1$ we have the standard Bayesian updating.

4.1.2 Updating Rules without Martingale Property

3. Exponential distortion to prior and conditional probabilities (Grether, 1980)

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

This updating rule can accommodate several updating biases but it violates the martingale property. $a \geq 0$ is the weight the agent places on the prior; when $a < 1$, we have

¹⁰Another useful property is divisibility (Cripps, 2018), where updating the signals sequentially or simultaneously does not affect the posterior belief. In the ambiguity literature, dynamic consistency (Epstein and Le Breton, 1993), receiving more information about the state of the world does not change the agent's optimal contingent plan, leads to Bayesian updating under expected utility theory.

¹¹This model nests cursed belief (Eyster and Rabin, 2005), where the updated belief is just a convex combination between the Bayesian posterior and the prior.

¹²There is an upper bound on λ to ensure that it is a valid probability measure.

base rate neglect and when $a > 1$, we have confirmation bias. $b \geq 0$ 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 $b = 1$ we have the standard Bayes' rule.

This model is widely used in the analysis of experimental data because the odds ratio can be log-linearized and estimated with a linear regression.

4. Affine Transformation of Bayesian Belief and reference belief ([Hagmann and Loewenstein, 2017](#))

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

In this model, μ is a reference belief which is a belief that the agent wants to have (for utility reason), and $\lambda \geq 0$ is 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 the agent will update asymmetrically, the agent over-updates when the signal moves the prior towards the reference belief and under-updates 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 (more likely for high wages to be drawn), and when $\lambda \neq 1$, the agent's belief will update towards this reference belief.

5. Multiple Prior models (e.g. ([Ortoleva, 2012](#); [Ba, 2022](#)))

In this class of model, the agent typically entertains two or more possible prior beliefs and there is a switching rule the agent adopts to switch between prior. For example, in [Ortoleva \(2012\)](#), the agent has a belief over which prior is most likely to be correct and he updates this distribution and all the priors with Bayes' rule. When he sees a surprising signal, the probability of seeing this signal with the current prior is below a threshold, he reevaluates which prior is most likely to be correct. This class of models violates the martingale property when a switch in prior occurs.

These models are typically used in misspecified learning models, where an agent incorrectly believe that the true states are impossible. Unlike the earlier models, when the agent assigns zero probability to an event in the prior, the updated belief for the event will always remain zero. This helps us to distinguish this class of updating rules from the other updating rules that were mentioned earlier.

While we do not have a test to falsify this class model, most of the behavioral patterns appear to be consistent with this model's prediction. This model is also intuitively appealing as it allows agents to be non-dogmatic about zero-probability events.

5 Empirical Strategy

There are two main challenges in working with this data. Firstly, the survey measured the probability distribution of the best wage offer, not the individual wage offer. The distribution of the best wage offer is a function of the arrival rate of job offers which may change over time as well. As such, when the respondent's belief changes, we can't attribute it completely to the respondent updating his belief about the wage distribution. To address this, we recover the single wage offer distribution from the reported maximum.

Secondly, the bins of the wage offer differ across every individual and time period. The bins of the wage offer are constructed using percentages of the expected best wage offers. If subjects report different best wage expectations in different time periods, the wage bins will have different intervals across each response. For example, when an individual changes his wage expectation, the data is as follows. The wage bins are not comparable across time.

Period 1's Belief

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

Period 2's Belief

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

Table 2: Example of a response from the SCE survey

To make the bins comparable across time, we fit a log-normal distribution to period 2's belief. This allows us to estimate what the updated beliefs are in the period 1's wage bins.

5.1 Distribution Fitting

We use the Simulated Method of Moments (SMM) to fit log-normal distributions to the five points on individuals' posterior cumulative density functions recorded in question oo2b. The main reason we fit the posterior distribution is that unbiased measurement error in the posterior distribution does not affect the [Augenblick and Rabin \(2021\)](#) test statistics.¹³ These five points are the ordered pairs consisting of the wage bin thresholds (e.g. 80% of oo2a2, 90% of oo2a2, and so on) as the horizontal values, and the corresponding survey responses of oo2b as the vertical values. We find the mean and standard deviation which minimize the vertical distance in CDF values (using the L2 norm) between points on the simulated CDFs with the same inverse CDF values as the five points and the five points themselves. An example is given in figure 3, where the blue points are points on the CDF implied by survey responses, and the orange points are points that minimize the vertical distance between the orange curve (a log normal CDF) and the blue points.

¹³We discuss this in greater detail in Section 7.




Figure 4: Example fitted distribution.

Since this process gives us an entire distribution for the posterior, we can directly estimate the probability weight placed over the ranges of the posterior distribution 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 selected this method for our main analysis results, since on average it minimized the fitting error between the data and simulated CDF the most out of the different methods we tried. These alternate methods, along with their fitting errors and their effects on our main result, are described in appendix B. Overall, while the method used to fit the data sometimes affected the magnitude of our point estimates, it did not change the direction, significance, or general interpretation of our main results (the Martingale test results described in section 6.1).

5.2 Martingale Tests

5.2.1 Test Description

The martingale property can be tested directly, and we will be using a test proposed by [Augenblick and Rabin \(2021\)](#) to do so. There are alternative tests, such as testing if the updated beliefs average out to be the prior belief. The main benefit of the [Augenblick and Rabin \(2021\)](#) test is that it is closely related to some of the more prominent belief-updating biases, such as base-rate neglect and underreaction to signals. This will give us insights into the type of belief updating bias the agent is exhibiting. Moreover, circumvents the need to compute the Bayesian beliefs, by estimating the conditional probabilities. This allows us to make fewer assumptions for the analysis.

The test involves computing two statistics: (1) belief movement, $m_{1,2}$, and (2) uncertainty reduction, $r_{1,2}$, as shown in equations 6 and 7 respectively. Let π_t^i denote the probability assigned to bin i at time t . With h bins, the two statistics are defined as

$$m_{1,2} \equiv \sum_{i=1}^6 (\pi_2^i - \pi_1^i)^2 \tag{6}$$

$$r_{1,2} \equiv \sum_{i=1}^6 \pi_1^i(1 - \pi_1^i) - \pi_2^i(1 - \pi_2^i) \tag{7}$$

Both statistics have an intuitive interpretation. The belief movement is the total squared difference between period 2’s belief and period 1’s belief. This captures how much belief is changing regardless of the direction of change. For uncertainty reduction, the statistics can be interpreted as a measurement of the “variance” of the belief. If we treat each bin like a Bernoulli distribution, the expression in the summation is the variance of the Bernoulli distribution in period 1 minus the variance of the Bernoulli distribution in period 2. This is summed across all 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. This means that if the agent’s

belief is expected to move greatly, we will expect the agent to become more certain about the state.¹⁴

The ideal test will require us to elicit the respondents' beliefs at every possible signal realization to compute the expected belief movement and uncertainty reduction. This will allow us to determine if an individual is Bayesian. Since we only observe single updated belief profile based on the signals the respondent observed in the last four months, we can only determine if the population is Bayesian at the aggregate level. Given that there are n observations, we compute the average belief movement as $\bar{m}_{1,2} \equiv \frac{1}{n} \sum_{j=1}^n m_{1,2}^j$ and the average uncertainty reduction as $\bar{r}_{1,2} \equiv \frac{1}{n} \sum_{j=1}^n r_{1,2}^j$. We then compute the average excess belief movement statistic:

$$X = \bar{m}_{1,2} - \bar{r}_{1,2} \tag{8}$$

where “excess” refers to the amount of movement exceeding reduction.

The standard error for this test is

$$s_{1,2} = \sqrt{\frac{1}{n} \sum \left(m_{1,2}^j - r_{1,2}^j - (\bar{m}_{1,2} - \bar{r}_{1,2}) \right)^2} \tag{9}$$

Under the null hypothesis that the population is Bayesian, we expect $X = 0$. With the central limit theorem, we can compute the Z-statistic that converges to a standard normal distribution:

$$Z \equiv \frac{\sqrt{n}}{s_{1,2}} (\bar{m}_{1,2} - \bar{r}_{1,2}) \xrightarrow{n \rightarrow \infty} N(0, 1) \tag{10}$$

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 address this concern, we also compute the normalized excess movement,

¹⁴To show this, we have to expand the terms and apply the law of iterated expectations. With the martingale property, the terms will cancel out nicely.

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 $X_{norm} = 1$.

$$X_{norm} = \frac{\bar{m}_{1,2}}{\bar{r}_{1,2}} = \frac{X}{\bar{r}_{1,2}} + 1 \quad (11)$$

As proposed by [Augenblick and Rabin \(2021\)](#), we will report the excess belief movement, the Z statistics and the normalized belief movement in our results. Overall, this test will help us to reject the hypothesis that people’s updating rule has the martingale property, which includes Bayesian updating.

5.2.2 Test Assumptions

The validity of this excess belief movement test hinges on two critical assumptions. Firstly, the agent’s prior has to be correct. If the agent reported an incorrect prior and updated this incorrect prior in a Bayesian manner, the equality relationship between belief movement and uncertainty reduction will not hold. Consider a simple two-state model with a fully revealing signal. Suppose the correct prior that state 1 occurred is 0.5 but the 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. We may incorrectly conclude that the agent is non-Bayesian. Unfortunately, we are unable to validate this assumption, however, [Augenblick and Rabin \(2021\)](#) showed that their test has low power in detecting incorrect prior beliefs in a calibration exercise. This means that even when respondents begin with a wrong prior but are updating in a Bayesian manner, it is unlikely for the test to reject the hypothesis that people are non-Bayesian when they have wrong priors.

Secondly, the individuals’ wage distribution has to remain stable over time. If the wage distribution is changing over time, the equality of the expected belief movement and

uncertainty reduction will not hold. Intuitively, when the state is ever-changing, we expect beliefs to change constantly and the uncertainty will not be resolved. While we cannot test how individuals' offer distributions are changing, we can see whether the aggregate distribution of all reported wage offers changes from period to period for individuals who report offers. Qualitatively, it is unlikely that the wage distribution changes in a duration of month. To test this more rigorously, we perform a KL divergence test to determine if the offers received by the survey respondents are the same across time.

We have some subjects who reported wage offers they received in consecutive surveys. The offers on the earlier survey are shown here as "period 0." These here occurred before the time of the earlier survey. For individuals who only took two surveys, the period 0 wages will thus be for a time not covered by our main analysis. For individuals who took three surveys, if they report wages on the second and third survey, the period 0 wages would be the same as those following the first survey and thus for a time period covered by our main analysis. We test whether the (log) offer distribution differs between the available period 0 and period 1 responses to get a sense of wage offer distribution stability in our sample. The graph below plots both offer distributions.




Figure 5: Period Wage Distributions

All offers here are considered separate and given equal weight in the graphed distribution. That is, for each observation (consecutive pair of surveys), an individual can contribute up to three offers to each period’s distribution. The “Period 1” distribution here differs from that of figure 2, because this graph only includes a wage reported on a pair’s second survey if there is at least one corresponding wage reported from the pair’s first survey (as the periods are defined in relative terms, rather than being absolute time periods).

We now formally test whether the two distributions differ using the Kullback–Leibler divergence.

Statistic	Value
Entropy	0.0163442
Standard Error	0.0106287
Histogram Bins	20
Reference Observations (Period 0 Wage Offers)	626
Comparison Observations (Period 1 Wage Offers)	625

Table 3: Kullback–Leibler Divergence Statistics for Period Log Wage Offer Distributions.

We note the 95% confidence interval includes 0, so we do not find evidence that the aggregate wage distribution underwent major changes between periods for the available data.

We also perform an individual fixed effect regression to test if the average wage offer of survey respondents changed over time for those reporting offers.¹⁵ As shown in table 4, we see that wage offers are not statistically different across the survey period. While not directly covering the period used by our main analysis, these tests do not find evidence of wage distribution instability over an adjacent time period of equal length.

Variable	Coefficient
Time Period	1920.13 (1581.013)
Constant	71,365.16 (789.8748)
Observations	1,251
Individuals	312
R-squared	0.003

Table 4: Fixed Effect Regression of Wage Offers over Time. Dependent variable is reported wages, with each observation being a different reported wage and each individual having a fixed effect. Standard errors are clustered at the state level.

¹⁵We do this in non-logged terms for easier interpretation.

6 Bayesian Updating (Martingale) Test Results

6.1 Martingale Test Statistics

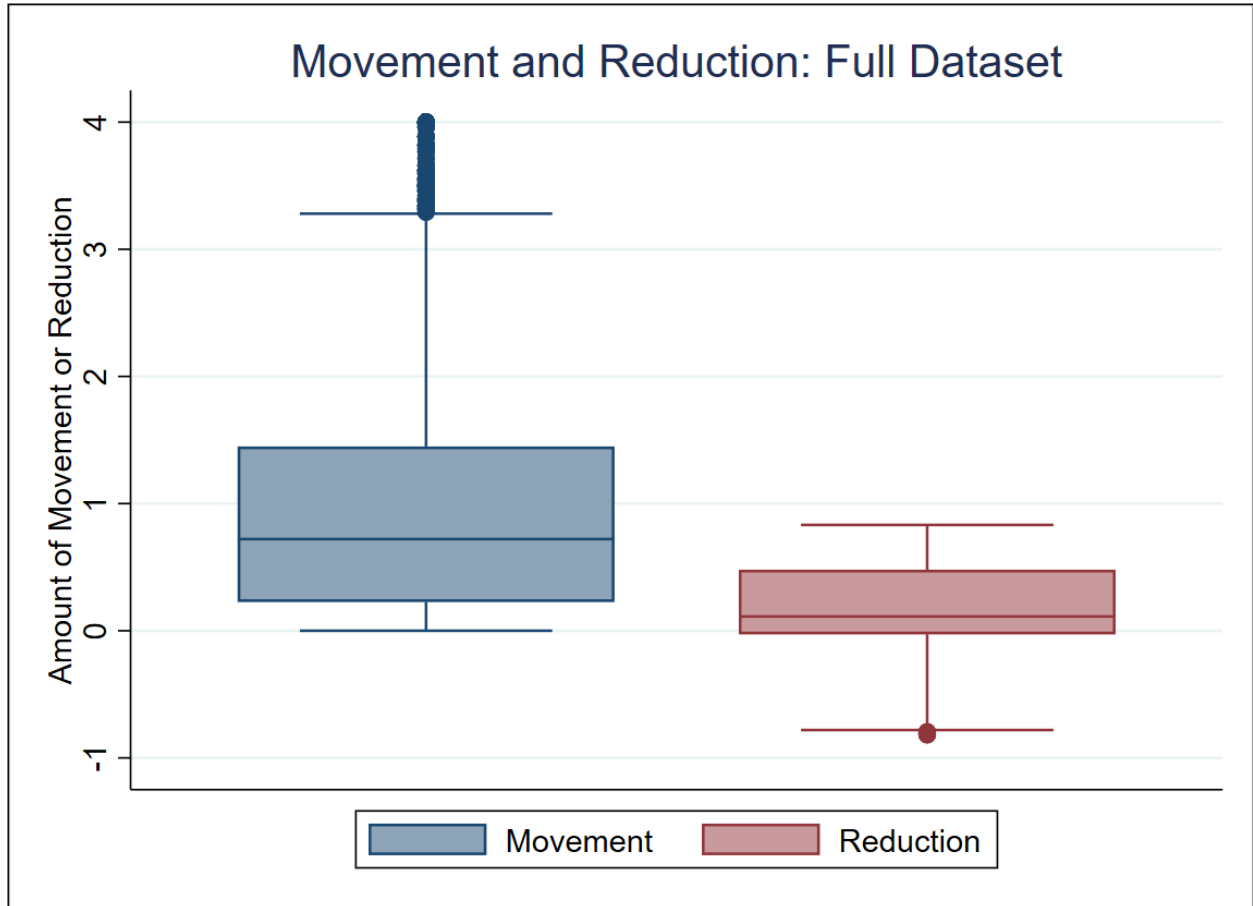


Figure 6: Movement and reduction summary statistics.

We first present the result from the excess belief movement. Overall we found excess belief movement relative to the amount of uncertainty reduction, as shown in column 1 of table 5 and the box plot in figure 6.¹⁶ The normalized excess belief movement of $X_{norm} = 5.1798$ means that beliefs are moving 418% more than the amount of uncertainty reduction. Excess

¹⁶The observation counts in table 5 are slightly less than in the summary statistics table because there were 3 individuals missing state (unable to assign to cluster) and 46 individuals with two updates for which we had enough information to use their second update, but not their first. In the latter case, we still used their second update in the analyses based on single updates, but not for the movement and reduction-based analyses. This is because the movement and reduction-based analyses are meant to use all measured periods for an individual.

belief movement suggests that people are over-updating relative to the Bayesian benchmark, and this is consistent with updating biases such as base rate neglect and overreaction to signals. Due to limitations of this dataset, we only observe individuals three times at most, and are thus unable to effectively distinguish between these two updating biases. In the later section, we do provide some suggestive evidence of perfect base rate neglect in our sample.

To get a sense of the magnitude of our estimate, we compare our estimate to [Augenblick and Rabin \(2021\)](#) where they study forecasters predicting the likelihood of geopolitical events, as well as beliefs in a British prediction market. For the geopolitical forecasters, they estimated an excess belief movement of 0.0137, with a normalized excess movement of 1.20. For the British prediction market, they estimated an excess belief movement of 0.00036, with a normalized excess belief movement of 1.046. The estimate from our study is significantly higher either of these. This difference could stem from differences in the setting or the complexity of the belief updating problem.

Statistic	All Individuals	Got Offer?		Searched?		
		Yes	No	Yes	No	Unknown
\bar{m}	.9426 (.0165)	1.0716 (.0412)	.9064 (.0183)	1.0228 (.0334)	.9050 (.0182)	.9664 (.0476)
\bar{r}	.1820 (.0072)	.2093 (.0170)	.1743 (.0072)	.1970 (.0111)	.1712 (.0077)	.2198 (.0241)
X	.7607 (.0175)	.8623 (.0427)	.7321 (.0202)	.8258 (.0325)	.7337 (.0198)	.7466 (.0466)
X_{norm}	5.1798	5.1191	5.2002	5.1910	5.2855	4.3972
Observations	3,003	658	2,345	848	1,940	215

Table 5: Excess movement statistics: Log normal-fitted results. Clustered errors by state in parentheses.

We hypothesized that people who received a job offer would update in a more Bayesian manner because job offers are individualized feedback about one’s wage distribution, while other information like labor market news are at the population level. But surprisingly, while those who got at least one offer have a slightly lower normalized statistic, the statistic is still very large and of a similar magnitude as the statistic for those who did not receive any

offers.

We also examine the heterogeneous effect between respondents who are searching for a job and not searching. Intuitively, people who are actively searching for a job are likely to put in more thought in forming their wage expectation. This could mean job searchers are more Bayesian than non-job searchers. However, again we found that both groups had very large normalized statistics. This could be suggestive that if increased mental focus in job seeking helps with becoming more Bayesian, it is not major enough to overcome the biasing effects of other factors.

Figure 7 below provides an alternative look into the martingale test results. The provided box plot compares belief movement and uncertainty reduction measures of four different groups of individuals. The graph shows the surprising result that movement and reduction remain fairly similar across both categories, contrary to our prediction that those with personal wage and those who are searching would behave in a more Bayesian manner.




Figure 7: Movement and reduction summary statistics.

6.2 Heterogeneous Effect by Demographics

We also examine heterogeneous effects across different demographic groups. These differences, while not causal, might suggest labor market updating phenomena for further study.

First, when we divide the population by age, we find that the oldest group (individuals 60 and older) had the lowest normalized excess movement of any demographic group (not just compared to other age groups). This is because the oldest group had the highest level of uncertainty reduction among all demographic groups.

When we divide the sample by education level, individuals with at least some college education are closer to the Bayesian benchmark compared to individuals with only a high school diploma or lower. We found a normalized excess belief movement that is 0.76 points

higher for those with a lower level of education. Unlike in the previous comparison, here the difference in the normalized statistic between the demographic groups is mostly because those with a higher level of education are less reactive in changing where they think the mean is, not because they become more sure about their guess than the lower education group as the older group did compared to the other age groups.

The difference in normalized movement statistics between the gender groups was the smallest of the group differences, with the female group having a normalized statistic 0.34 higher than the male group. For the demographic race variable, the difference between the minority and majority group statistics was 0.59.

Statistic	Age			Some College		Gender		Race	
	18 to 39	40 to 59	60 and Above	Yes	No	Female	Male	White	Non-white
\bar{m}	0.8224 (0.0223)	1.0251 (0.0243)	0.9842 (0.0376)	0.9279 (0.0188)	1.1401 (0.0648)	0.9723 (0.0253)	0.9208 (0.0190)	0.9339 (0.0155)	0.9921 (0.0534)
$\bar{\tau}$	0.1673 (0.0117)	0.1687 (0.0117)	0.2348 (0.0132)	0.1811 (0.0079)	0.1939 (0.0183)	0.1810 (0.0109)	0.1829 (0.0080)	0.1833 (0.0079)	0.1746 (0.0133)
X	0.6552 (0.0186)	0.8564 (0.0223)	0.7494 (0.0380)	0.7468 (0.0199)	0.9462 (0.0673)	0.7913 (0.0294)	0.7379 (0.0184)	0.7506 (0.0164)	0.8175 (0.0490)
X_{norm}	4.9173	6.0777	4.1913	5.1237	5.8797	5.3710	5.0344	5.0956	5.6814
Observations	1,095	1,280	628	2,794	209	1,288	1,714	2,553	450

Table 6: Excess movement statistics: Demographic groups. Standard errors clustered by state in parentheses.

6.3 Test Implications for Updating Rules

Given the result that people are non-Bayesian, what are some updating rules that can generate this data? Firstly, we can reject all updating rules with the martingale property. This includes Bayesian updating and the [Epstein, Noor and Sandroni \(2010\)](#) model.

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 suggests the agent exhibits base-rate neglect and 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 which results in 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, 2020](#)).

The [Hagmann and Loewenstein \(2017\)](#) updating rule and the multiple prior model ([Ortoleva, 2012](#); [Ba, 2022](#)) can produce excess belief movement in the beliefs with an appropriate combination of model parameters.

7 Robustness Check: Measurement Error

In surveys, it is possible to potentially misreport their 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.

7.1 Theoretical Analysis

The normalized excess belief movement statistics is significantly larger than the estimates [Augenblick and Rabin \(2021\)](#) got in their empirical application in their paper. To further understand why we have such a large estimate we perform a Monte Carlo Simulation where we allow for measurement error.

Consider that 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 uncorrelated with recent belief and error realizations ($\mathbb{E}(\epsilon_t \pi_t) = \mathbb{E}(\epsilon_t \pi_{t-1}) = \mathbb{E}(\epsilon_t \epsilon_{t-1})$), [Augenblick and Rabin \(2021\)](#) showed that the excess belief movement will be equal to $2\sigma_{\epsilon_t}^2$.

Generalizing this to n states, we show that the excess belief movement will be equal to $\sum_{i=1}^n 2\sigma_{\epsilon_t^i}^2$.¹⁷ With measurement error, the excess belief movement to be equal to the total variance of the measurement error in the prior belief multiplied by 2. One interesting feature of this test is only measurement error in the prior affects the test statistics. This motivates our method of fitting a log-normal distribution to the posterior to compare it to the prior beliefs, which minimizes the error term in the prior.

This result tells us that only the measurement error in the first period will affect the excess belief movement statistics. In the calibration exercise, we only have to introduce the measurement error to the prior belief.

7.2 Monte Carlo Simulation

We perform a Monte Carlo simulation to determine how much measurement error is required to produce the result we obtained. The Monte Carlo simulation is designed to mirror the setting 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 3,003 observations for which we can calculate movement and reduction statistics, we generated 3,003 pairs of priors and posteriors in each simulation.

Since our statistics of interest only depend on measurement error 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.¹⁸

¹⁷Refer to appendix [A.1](#) for the poof.

¹⁸The draws from the Dirichlet distributions are a valid probability distribution that has to sum to 1. This

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 0.9426 and 0.1820, 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 posterior belief is then selected to match the supposed Bayesian belief movement and reduction statistics. One of the possible posterior beliefs that is selected in this calibration exercise is $(0, 0.838, 0.162, 0, 0, 0)$. 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.

Finally, the variance of the error term is not an intuitive measure of the amount of measurement error. To quantify measurement error in a more intuitive manner, 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.

We simulate this process 10,000 times. In Table 7, we report the average statistics computed from the 10,000 Monte Carlo simulations and the interval that contains 95% of our simulation result.

approach is better than adding an error that is normally distributed which may cause the probability measure to be negative or exceed 1.

We see that in order 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 about 117 percentage points across the six bins. The measurement error required to rationalize the behavior is very large and is unlikely to explain the result we have obtained.

Simulated Statistic	Uniform Prior with Matched Parameters	Target Parameter from Martingale Test
\bar{m}	0.9421 [0.9257, 0.9584]	0.9426
\bar{r}	0.1795 [-0.2297, 0.4607]	.1820
X	0.7606 [0.7401, 0.781]	.7607
X_{norm}	5.1933 [4.947, 5.4553]	5.1798
Δ	1.1669 [1.158, 1.1759]	

Table 7: Measurement error calibration statistics. Dirichlet priors were chosen to match sample average movement and reduction statistics from Martingale test results. Excess movement and normalized excess movement are also listed here for comparison, although not used directly in choosing parameters for Dirichlet draws.

8 Individual Level Analysis

We now present a few more analyses designed to distinguish between updating rules without the martingale property. We will be looking at two belief updating patterns: (1) Updating zero probability events in the prior to a positive probability, and (2) asymmetric updating between positive and negative news.

8.1 Zero Probability Analysis

The [Grether \(1980\)](#) model makes a stark prediction where once an event is thought to be impossible, no signal can change the agent’s mind about this. One important caveat is in the [Grether \(1980\)](#) model, there is a special and non-trivial case of perfect base rate neglect

($a = 0$). In this situation, it may be possible for a zero probability event to be updated into positive probability by completely ignoring prior.¹⁹ We also do acknowledge that when survey respondents report zero probability about receiving a certain wage offer, they may not mean zero probability but an infinitesimal chance. A more conservative interpretation of our result here would be we need extreme bias in the [Grether \(1980\)](#) model or they receive a very informative signal to rationalize this behavior.

However, there are some updating rules in the literature like the [Hagmann and Loewenstein \(2017\)](#) model and the multiple prior model ([Ortoleva, 2012](#); [Ba, 2022](#)), allows the agent to update a zero probability event to a positive probability. In this dataset, we observe numerous cases where individuals who initially deemed it impossible to receive wage offers within a specific range but altered their views upon subsequent surveys.

The main idea of this analysis is to find beliefs which were initially reported to be zero, but which become positive in the second survey. In general, this cannot be done directly, because the bin definitions change for each distribution, as described previously in the empirical strategy section. We use two different approaches for detecting this behavior.

8.1.1 Subset Method

The first method (hereafter referred to as the “subset method”) relies on identifying wage bins from the posterior distribution that are a subsets of wage bins on the prior distribution. We attempt to identify behavioral patterns as shown in the tables below which shows a survey response.

¹⁹We are grateful to Sevgi Yuksel for pointing this out.

Period 1 Belief

	$w < 52,000$	$52,000 \leq w < 58,500$	$58,500 \leq w < 65,000$	$65,000 \leq w < 71,500$	$71,500 \leq w < 78,000$	$w \geq 78,000$
$p(\cdot)$	0	0	0.7	0.2	0.1	0

Period 2 Belief

	$w < 44,000$	$44,000 \leq w < 49,500$	$49,500 \leq w < 55,000$	$55,000 \leq w < 60,500$	$60,500 \leq w < 66,000$	$w \geq 66,000$
$p(\cdot)$	0	0.3	0.4	0.3	0	0

Table 8: Example of a survey respondent’s belief

Note that the individual places zero probability on the wage bin $w < \$52,000$ in period 1’s belief. When the individual updates their beliefs, the second bin, $\$44,000 < w < \$49,500$ is a subset of the bin from the first period distribution, and there is a probability weight of 30% in this bin.

The strength of this approach is that it is nonparametric and can be used directly with the data without any additional assumptions. However, this method is restrictive as it requires zero probability weight and a posterior bin to be a subset of the prior bin to identify such behavior. As such, we should interpret the results as a lower bound of the number of the number of people who are inconsistent with the [Grether \(1980\)](#) model.

8.1.2 Simulated Cutoff Method

The second method (hereafter referred to as the “simulated cutoff method”) relies on the belief fitting process described in section 5. The following table shows the fitted period 2 beliefs for the same individual from the earlier table.

	$w < 52,000$	$52,000 \leq w < 58,500$	$58,500 \leq w < 65,000$	$65,000 \leq w < 71,500$	$71,500 \leq w < 78,000$	$w \geq 78,000$
$p(\cdot)$	0.4095	0.5880	0.0026	0	0	0

Table 9: Period 2 Fitted Beliefs

Since the fitting the beliefs gives us a full distribution, we are able to directly compare the fitted posterior beliefs to the raw prior beliefs using the wage bins from the prior. However, since the fitted distribution is a log normal distribution, it will technically have positive weight over any positive wage values.²⁰ We label the individual as being inconsistent with the [Grether \(1980\)](#) model if the estimated posterior distribution has a probability weight that exceeds a certain threshold in wage bins that were initially given a zero probability weight.

The lowest cutoff is for a fitted bin to have at least 0.5% probability to be considered non-zero, while the highest cutoff is for a fitted bin to have at least 10% probability. We choose 10% probability to provide some leeway for errors from estimating the posterior belief.

8.1.3 Results

The results of the two methods for finding people who are inconsistent with the [Grether \(1980\)](#) model is shown in the chart below. Overall, we estimate that about 20-50% of the observations in our data that updates in a way that is inconsistent with the [Grether \(1980\)](#) model where $a > 0$.

²⁰The last three bins shown are exactly 0, but this is only because the table values are rounded to four decimal places.




Figure 8: Fraction of observations which are inconsistent with Grether (1980).

Using the most conservative subset method, about 20% of the individuals update a zero probability event into a positive probability event. When we make some parametric assumption about the updated belief distribution, we find that about 37-49% of the subjects are inconsistent with the Grether (1980) model where $a > 0$.

There are two possible interpretations of this result. Firstly, individuals are updating their beliefs with the Grether (1980) model but with perfect base rate neglect. If people exhibit base rate neglect in this setting, it might be impossible for individuals to learn about the distribution from which they are drawing their wage. This may justify a need for some policy intervention to help individuals learn about the wage distribution. Secondly, they are updating their beliefs with rules where zero probability events in the prior can be updated to a positive probability in the updated beliefs.

8.2 Other non-Bayesian Updating Patterns

For individuals who reported receiving a job offer, we have some information about the signals they have received during the job search. This allows us to examine other updating patterns as well.

Under a standard Bayesian updating job search model, when an agent receives a wage offer that is lower than the prior mean, the agent would adjust his wage expectation downwards. Conversely, when the agent receives a wage offer that is higher than his prior mean, he will adjust his expectation upwards. In a Gaussian updating framework, the mean of the posterior belief should always be a convex combination between the prior mean and the wage offer.

Motivated by this framework, we look at individuals whose mean updated beyond the wage from the reported job offers (“overshooting”) and those whose expected best wage (response to survey question oo2a2) moved in the opposite direction of any wage offer (“wrong direction”). For example, suppose an individual expected their best wage offer to be \$30,000 in period 1 and \$40,000 in period 2. We would classify this update as “overshooting” if they reported having only been offered wages of \$34,000, \$35,000, and \$36,000. On the other hand, we would consider the update “moving in the wrong direction” if they reported only having been offered wages of \$27,000, \$28,000, and \$29,000. Importantly, we would not classify this update as “overshooting” if they were to report wages of \$29,000, \$35,000, and \$36,000 - although the mean of such offers is above the initial expectation, we do not know how informative the agent considers each wage signal. Similarly, if they reported wages of \$27,000, \$28,000, and \$31,000, we would not consider this “moving in the wrong direction.”

We found individuals who exhibit these updating patterns, but we also see that there is an asymmetry in the direction of the update which is consistent with motivated beliefs (Bénabou and Tirole, 2002; Eil and Rao, 2011; Hagmann and Loewenstein, 2017). The fraction of updates with offers following each updating pattern is summarized in the figure

9.

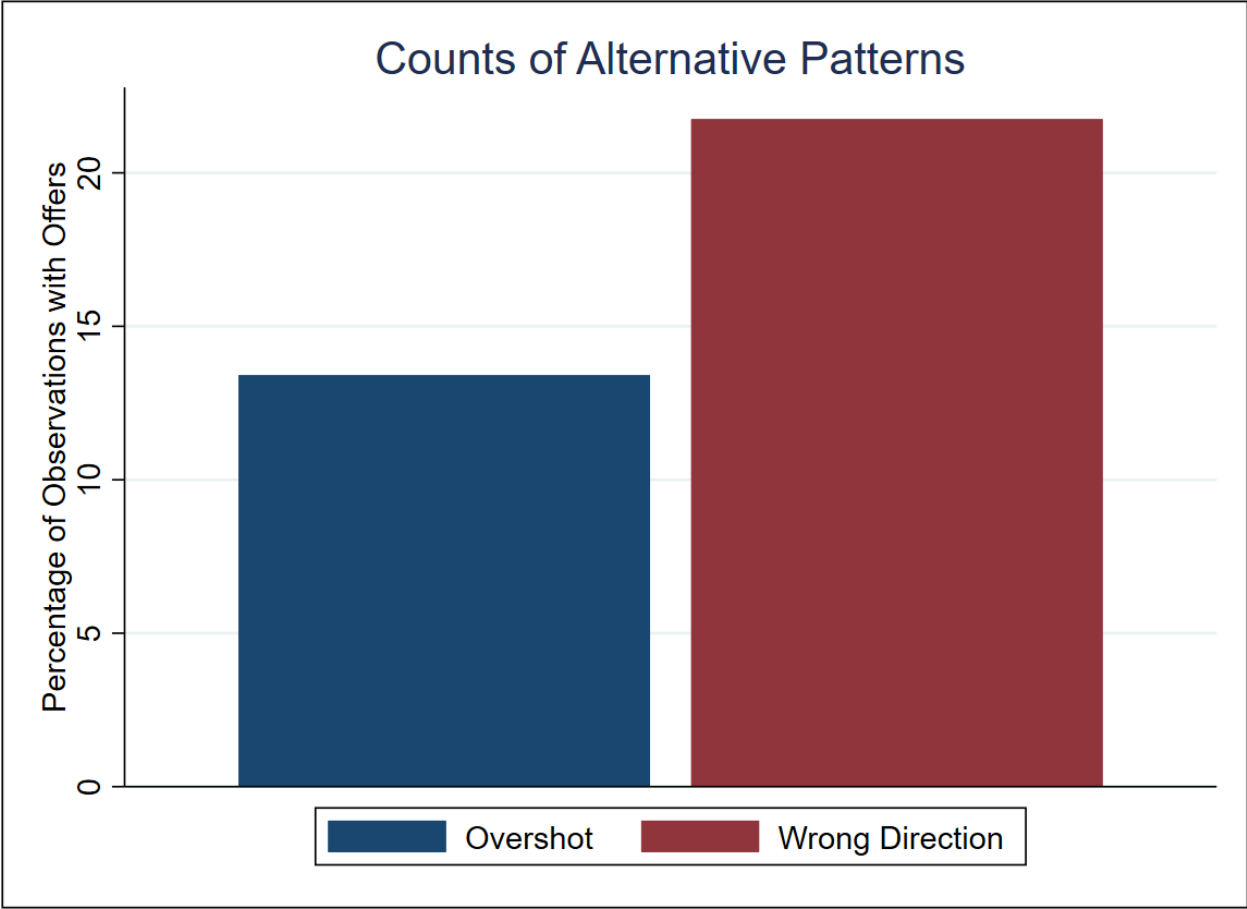


Figure 9: Counts of specific patterns in data as percentages.

We see that 98 updates (13.41% of all observations with offers) overshoot, while 159 updates (21.75% of all observations with offers) moved in the wrong direction.

We see a pattern when looking at updates by whether the wage information received should have raised or lowered expectations. In figure 10 below, we see that among subjects who received an upward signal, 44.44% of the observations (80 updates) overshoot. On the other hand, only 3.58% of observations (17 updates) overshoot when receiving a downward signal. When looking at moving opposite to signals received, 30.74% of observations (146 updates) that received only downward signals updated upward, compared to only 7.22% of observations (13 updates) that updated downwards when receiving only positive signals. We also note here that the total number of observations where the maximum wage was above the initial expectation was 180, while the number of observations where the maximum wage was

below the initial expectation was 475. This difference would be consistent with individuals being overly optimistic about wages in general, not just in how they update.²¹




Figure 10: Signal direction defined as relative to reported first period best offer expectation.

Here, we defined signal direction as relative to the maximum reported wage. That is, if someone had reported three wages: \$31,000, \$32,000, and \$33,000, we would consider this an “upward signal” to anyone with a response to oo2a2 of less than \$33,000. We did it this way in this analysis because survey question oo2a2 is explicitly about the “best” (most likely to be accepted) wage, which we interpret here as the maximum wage. However, we also tried an alternate analysis using question oo2a and defining signal direction relative to the average reported wage, which can be found in appendix C. We note that the decision whether to use the maximum or average wage to define signal direction will not affect a

²¹The remaining 76 had their initial expectation exactly equal to the maximum reported offer received.

signal direction category’s raw number of observations considered overshooting or moving in the wrong direction. This is because to be classified as having one of the two non-Bayesian patterns, reported wages must not be divided on either side of a reported expectation (as described by the descriptive examples earlier in this section). However, this difference in signal direction definition strategy could affect whether wages in an update are considered an “upward” or “downward” signal for those not showing one of the patterns, thus affecting the relative percentages in figure 10.

The behavior we see in the figure is consistent with motivated beliefs (Bénabou and Tirole, 2002; Eil and Rao, 2011) where individuals want to have an optimistic outlook on their job prospect. In a theoretical model like Hagmann and Loewenstein (2017), we set the reference belief to an “optimistic” belief and we can obtain this asymmetric updating pattern where beliefs would move upwards. If we set λ to be large it is possible for people to update in the wrong direction as well.

Another possible explanation for people updating in the wrong direction is the gambler’s fallacy (Rabin, 2002).²² When an individual sees a low offer, he thinks that the next offer is more likely to be high. In this dataset, we see that people exhibit this bias in a self-serving way as the gambler’s fallacy is more likely to manifest when they are supposed to update their belief downwards.

9 Discussion

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 Epstein, Noor and Sandroni (2010) updating rule. Theoretical models that wants to incorporate non-Bayesian updating to job search model, should avoid using updating rules with the martingale property as such models do not describe how people update their beliefs.

The updating patterns observed consistent with base-rate neglect and overreaction to

²²This is conjecture posited by Mueller, Spinnewijn and Topa (2021) as well who analyzed the SCE data.

signals. We also find a significant proportion of individuals who updates zero probability events into positive probability. This updating pattern is inconsistent with the [Grether \(1980\)](#) model without perfect base rate neglect. Moreover, we detected evidence of asymmetric updating. When providing information that can cause individuals to adjust their wage expectations downwards, we may need a stronger or more informative signal. This behavioral pattern can be explained by the [Hagmann and Loewenstein \(2017\)](#) updating model.

Our findings suggest that information intervention is a powerful tool for influencing people’s beliefs in a job search setting, due to over-updating relative to the Bayesian benchmark. However, in our information-rich world, where individuals are continuously exposed to a stream of new information, it is essential to recognize that the effect of providing a single piece of information may be short-lived as people may overreact to the new information they receive or neglect the older information. Further, since we did not find that the overreaction bias was reduced for individuals receiving a signal directly from the labor market personalized to them, it could be that not all sources of labor information may have the same level of impact on affecting beliefs. To ensure a more lasting impact on beliefs, it might be necessary to adopt a strategy of consistently providing new and relevant information over time, as well as experimenting with different types of labor market information to see which kind individuals use the most effectively.

Our results are limited in a few important ways that leave scope for future work. Firstly, the current dataset does not allow us to effectively distinguish between overreaction and base rate neglect. These two updating biases have different asymptotic properties which will result in different long run behavior. If individuals exhibit overreaction, it is possible for them to eventually learn about the wage distribution they are drawing from. Overreaction does not pose a problem in the long-run if there is enough feedback and information to allow their beliefs to converge. On the other hand, if individuals exhibit base rate neglect, it may be impossible for them to learn about the distribution they are drawing their wage distribution. In this case, there may be a stronger need for intervention to correct people’s beliefs as it

is unlikely they will ever learn about their underlying wage distribution through feedback alone.

Second, 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. Third, we find a large effect among individuals with no wage offers, implying that our main result may be the product of unobservable factors or survey design. Future analyses could test the sensitivity of question phrasing and propose possibly more effective survey techniques. Another limitation is that we observe most individuals for only two periods, which limits our ability to study learning over extended time periods. Future analyses could collect higher-frequency data, allowing for a better study of the change in learning biases over time and across different time horizons.

Appendix A Proofs

A.1 Measurement Error

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 ($\mathbb{E}(\epsilon_t \pi_t) = \mathbb{E}(\epsilon_t \pi_{t-1}) = \mathbb{E}(\epsilon_t \epsilon_{t-1})$).

The belief movement will be larger by σ_ϵ^2

$$\begin{aligned} \mathbb{E}(M_{t,t+1}) &= \mathbb{E}(\pi_{t+1} + \epsilon_{t+1} - \pi_t - \epsilon_t)^2 \\ &= \mathbb{E}[(\pi_{t+1} - \pi_t)^2 + 2(\pi_{t+1} + \pi_t)(\epsilon_{t+1} - \epsilon_t) + (\epsilon_{t+1} - \epsilon_t)^2] \\ &= \mathbb{E}[(\pi_{t+1} - \pi_t)^2 + \epsilon_{t+1}^2 - 2\epsilon_{t+1}\epsilon_t + \epsilon_{t+1}^2] \\ &= \mathbb{E}[(\pi_{t+1} - \pi_t)^2] + \sigma_{\epsilon_t}^2 + \sigma_{\epsilon_{t+1}}^2 \end{aligned}$$

The uncertainty reduction will be smaller by σ_ϵ^2

$$\begin{aligned}
\mathbb{E}(R_{t,t+1}) &= \mathbb{E}[(\pi_t + \epsilon_t)(1 - \pi_t - \epsilon_t) - (\pi_{t+1} + \epsilon_{t+1})(1 - \pi_{t+1} - \epsilon_{t+1})] \\
&= \mathbb{E}[(\pi_t)(1 - \pi_t - \epsilon_t) + \epsilon_t(1 - \pi_t - \epsilon_t) - (\pi_{t+1})(1 - \pi_{t+1} - \epsilon_{t+1}) + \epsilon_{t+1}(1 - \pi_{t+1} - \epsilon_{t+1})] \\
&= \mathbb{E}[(\pi_t)(1 - \pi_t) - \pi_{t+1}(1 - \pi_{t+1})] - \sigma_{\epsilon_t}^2 + \sigma_{\epsilon_{t+1}}^2
\end{aligned}$$

The expected excess belief movement statistics 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,

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

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

The excess belief movement is

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

Appendix B Alternative Posterior Fittings

B.1 Fitting Method Descriptions

In this section, we describe other methods we tried to fit the posteriors to the data besides fitting log normal distributions directly to the oo2b responses. The effects of these alternate fitting methods are given in the section after.

B.1.1 Log Normal Fitting, Recovered Individual Offer Distribution

The first other way we tried was to back out the offer distribution for individual wage offers rather than directly fitting the distribution implied by oo2b. In the main paper, we fit the data with a log normal distribution because observed wages have been found to be generally log normal. Since the question asks about offers the agent is most likely to accept, beliefs about these offers are likely to more closely reflect beliefs about accepted wages and thus also follow a log normal distribution. However, the best wage offer distribution is a function of the individual wage distribution and the number of offers the respondent expects. The latter is likely to be a function of the respondent’s search effort. For instance, if the respondent is actively searching for a job, the worker may expect to receive more job offers, and the respondent will report a “better” best-wage distribution. This makes it challenging to deduce if the respondents are Bayesian if the search effort differs across time. To alleviate this concern, we estimate their single-wage distribution using the data and fit it as a robustness check.

We recover the single-wage offer using the following procedure. Assuming that all wages are drawn from the same single wage offer distribution independently, we let the CDF from a single wage offer be $F(w)$, the CDF of the maximum wage distribution from n offer is $F^n(w)$. From the maximum wage distribution given in “oo2b,” we take the n th root to obtain the CDF of the individual wage offer. For individuals who expect to receive zero wage offers, we assume the distribution they report in this question is the distribution of the single wage

offer distribution.

B.1.2 Extreme Value Fitting

We next tried to fit an extreme value distribution to the data. The motivation for this was that question oo2a2 was about the “best” offer an individual received. Our intuition is that this would be the highest (maximum) offer, on average. We therefore used SMM to fit location and scale parameters for a Gumbel distribution to the data, similar to how we fit mean and variance parameters for the log normal distributions.

B.1.3 Kernel-Fitted Posterior

Finally, we also estimated the probability density function of the posterior by considering the midpoint of each bin as representative of samples drawn from the bin. The reported bin probability from question oo2b was considered as if it were the percentage of samples drawn from the bin. The kernel density estimator was then applied to logged values of the six bin midpoints. That is, for a value x on the logged posterior distribution corresponding to question oo2b, probability density was estimated as

$$f_h(x) = \sum_{i=1}^6 w_i K_h(x - m_i)$$

where w_i was the probability assigned to each of the six bins, m_i was the logged value of the i^{th} bin’s midpoint, and h was the bandwidth. h was selected for each individual using the simulated method of moments to minimize the error between the cumulative fitted density over the bin and reported density over the bin. The Gaussian kernel was used.

Note, however, that the top bin is unbounded and that the bottom bin is much larger than the others in the survey question. To allow for assigning the top bin a point for use in the density estimation, the top bin was assumed to be from 120% to 130% of the response of question oo2a2, rather than from 120% to infinity as on the actual questionnaire. This

let the top bin cover as much of the distribution as the other bins.

For the bottom bin, we tried two different specifications. In the “restricted” method, we assumed the bottom bin was 70% to 80% of the oo2a2. This meant that each bin would have equal width (10%) and that the midpoint for the bottom bin would be assigned to 75% of oo2a2, something we expected would be closer to where individuals would place the actual weight of the distribution (since 40% offers could be very low or unrealistic, depending on the value of oo2a2). In the “unrestricted” method, we assumed the bottom bin was between 0% and 80% as explicitly defined on the survey, so that the midpoint of the bottom bin would be 40% of oo2a2.

B.2 Model Fit and Estimate Results by Method

The table below gives the average fitting error for each method described in the previous section, as well as for the main method used in the paper. The log normal distributions fit the best, but each of the fittings not based on kernel estimation is very close when minimizing the vertical distance rather than the horizontal distance.²³ The kernel-based methods have worse fits than the other methods by a noticeable margin.

Distribution	Error Minimization	Mean Total Error	Standard Deviation	Average Bin Error
Log normal, Direct Survey Responses	Horizontal	.4859	.3494	.0810
	Vertical	.2202	.2195	.0367
Log normal, Recovered Individual Offer	Horizontal	.5048	.3702	.0841
	Vertical	.2120	.2166	.0353
Gumbel	Horizontal	.8075	.5626	.1346
	Vertical	.2255	.2296	.0376
Midpoint Kernel	Restricted	.9309	.1013	.1551
	Unrestricted	.8321	.2332	.1387

Table 10: Average error by data fitting method. Total error is defined as the absolute differences between a survey posterior’s cumulative probability (in decimal form) over a wage bin and the cumulative probability for the same wage range on the fitted cumulative density function, summed over all six wage bins.

²³It is not surprising that using methods attempting to minimize the vertical distance between the points leads to smaller fitting error, since the reported error here is in vertical terms.

In the next table, we see whether changing the fitting method impacts our main results. We find large normalized statistics for all of the non-kernel methods, and the other statistics seem fairly close to each other, regardless of fitting method. The non-kernel methods have much lower normalized statistics, but it should be noted that they also have much worse fit. Therefore, these results suggest that our results are not very sensitive to which distribution we fit.

Statistic	Log Normal, Direct Survey Responses		Log Normal, Recovered Individual Offer		Gumbel		Midpoint Kernel	
	Horizontal	Vertical	Horizontal	Vertical	Horizontal	Vertical	Restricted	Unrestricted
\bar{m}	0.8466 (0.0161)	0.9426 (0.0165)	0.8546 (0.0169)	0.9406 (0.0175)	1.0226 (0.0174)	0.9463 (0.0183)	0.4258 (0.0051)	0.4768 (0.0084)
\bar{r}	0.1432 (0.0076)	0.1820 (0.0072)	0.1271 (0.0071)	0.1650 (0.0070)	0.2365 (0.0071)	0.1910 (0.0075)	0.3116 (0.0044)	0.2988 (0.0057)
X	0.7033 (0.0174)	0.7607 (0.0175)	0.7275 (0.0177)	0.7756 (0.0181)	0.7861 (0.0190)	0.7552 (0.0192)	0.1142 (0.0095)	0.1780 (0.0110)
X_{norm}	5.9099	5.1798	6.7221	5.7016	4.3243	4.953	1.3664	1.5957

Table 11: Martingale test statistics using alternate distribution fittings. Standard errors clustered by state are in parentheses.

Appendix C Alternative Asymmetry Analysis

In the main paper, we used the question oo2a2 as the basis of our asymmetry analysis to keep it consistent with the other analyses which required the distributional responses in oo2b (themselves defined using the responses recorded in oo2a2). However, our asymmetry analysis does not require any distributional variables. Hence, we could also have used question oo2a for this part of the paper (see section 3.2 for a description of the differences in these survey questions). We replicate figures 9 and 10 from section using question oo2a in figures 11 and 12 below.




Figure 11: Counts of specific patterns in data as percentages.

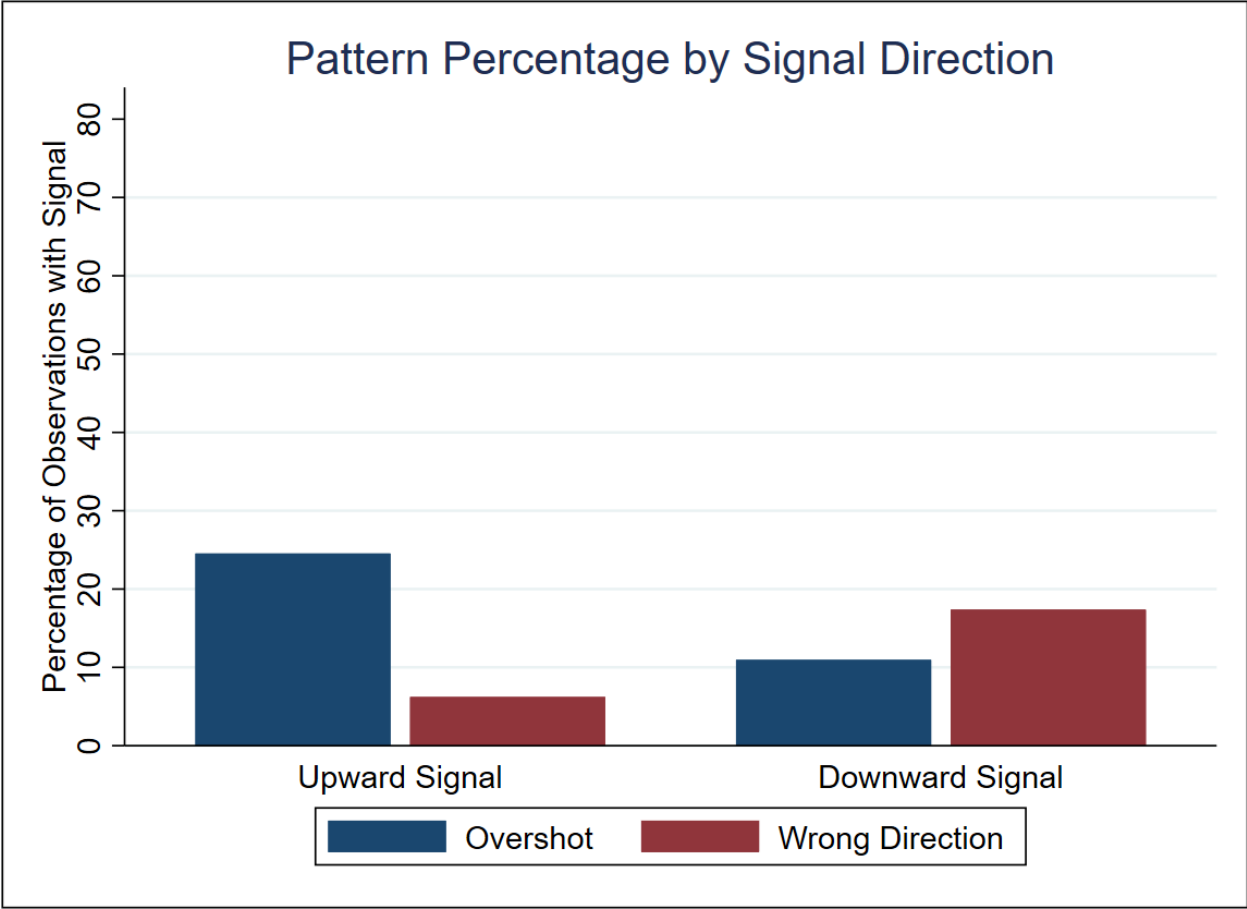


Figure 12: Signal direction defined as relative to reported first period average offer expectation.

The main methodological difference here is that for figure 12, we defined signal direction as relative to the average wage, rather than the maximum as we had done in figure 10. That is, if someone had reported three wages: \$31,000, \$32,000, and \$33,000, we would consider this an “upward signal” to anyone with a response to oo2a of less than \$32,000. We did it this way in this analysis because survey question oo2a is explicitly about the *average* wage, not the *best* wage like in oo2a2. Again, we note that this should not change the way those overshooting or moving in the wrong direction themselves are classified, but could change their relative percentages by changing whether observations not exhibiting either pattern are considered having “upward” or “downward” signals.

While less of the data is classified as belonging to one of the two patterns using this

alternative method, the direction of the asymmetric patterns remains the same: individuals are more likely to overshoot than to update in the wrong direction when receiving higher than expected wage offers and vice versa for lower than expected offers.

References

- Ahn, Hie Joo and Ling Shao**, “The Cyclicalities of On-the-Job Search Effort,” *The B.E. Journal of Macroeconomics*, 2021, 21 (1), 185–220.
- Armantier, Olivier, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar**, “An overview of the survey of consumer expectations,” *Economic Policy Review*, 2017, pp. 51–72.
- Arni, Patrick P**, “Opening the Blackbox: How Does Labor Market Policy Affect the Job Seekers’ Behavior? A Field Experiment,” *IZA Discussion Paper Series*, 2016.
- Augenblick, Ned and Matthew Rabin**, “Belief movement, uncertainty reduction, and rational updating,” *The Quarterly Journal of Economics*, 2021, 136 (2), 933–985.
- Ba, Cuimin**, “Robust Misspecified Models and Paradigm Shift,” 2022.
- Benjamin, Daniel J**, “Errors in probabilistic reasoning and judgment biases,” *Handbook of Behavioral Economics: Applications and Foundations 1*, 2019, 2, 69–186.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer**, “Diagnostic expectations and stock returns,” *The Journal of Finance*, 2019, 74 (6), 2839–2874.
- , —, **Yueran Ma, and Andrei Shleifer**, “Overreaction in macroeconomic expectations,” *American Economic Review*, 2020, 110 (9), 2748–82.
- Bénabou, Roland and Jean Tirole**, “Self-Confidence and Personal Motivation,” *The Quarterly Journal of Economics*, 2002, 117 (3), 871–915.
- Chan, Alex**, “Discrimination Against Doctors: A Field Experiment,” 2022. Unpublished.
- Conlon, John J, Laura Pilossoph, Matthew Wiswall, and Basit Zafar**, “Labor market search with imperfect information and learning,” Technical Report, National Bureau of Economic Research 2018.

- Cooper, Michael and Peter Kuhn**, “Behavioral job search,” *Handbook of Labor, Human Resources and Population Economics*, 2020, pp. 1–22.
- Cripps, Martin W**, “Divisible updating,” *Manuscript, UCL*, 2018.
- DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F. Schmieder**, “Reference-Dependent Job Search: Evidence from Hungary*,” *The Quarterly Journal of Economics*, 05 2017, *132* (4), 1969–2018.
- Eil, David and Justin M. Rao**, “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself,” *American Economic Journal: Microeconomics*, 2011, *3* (2), 114–138.
- Epstein, Larry G and Michel Le Breton**, “Dynamically consistent beliefs must be Bayesian,” *Journal of Economic theory*, 1993, *61* (1), 1–22.
- , **Jawwad Noor, and Alvaro Sandroni**, “Non-bayesian learning,” *The BE Journal of Theoretical Economics*, 2010, *10* (1).
- Esponda, Ignacio, Emanuel Vespa, and Sevgi Yuksel**, “Mental models and learning: The case of base-rate neglect,” Technical Report, Tech. rep 2020.
- Eyster, Erik and Matthew Rabin**, “Cursed equilibrium,” *Econometrica*, 2005, *73* (5), 1623–1672.
- Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa**, “Job Search Behavior Among the Employed and Non-Employed,” *Econometrica*, 2022, *90* (4), 1743–1779.
- Gee, Laura K**, “The more you know: Information effects on job application rates in a large field experiment,” *Management Science*, 2018, *65* (5), 2077–2094.
- Grether, David M**, “Bayes rule as a descriptive model: The representativeness heuristic,” *The Quarterly journal of economics*, 1980, *95* (3), 537–557.

- Haaland, Ingar and Christopher Roth**, “Beliefs about racial discrimination and support for pro-black policies,” *Review of Economics and Statistics*, 2023, 105 (1), 40–53.
- Hagmann, David and George Loewenstein**, “Persuasion with motivated beliefs,” in “Opinion Dynamics & Collective Decisions Workshop” 2017.
- Kahneman, Daniel and Amos Tversky**, “On prediction and judgement,” *ORI Research monograph*, 1972, 1 (4).
- Kudlyak, Marianna, Damba Lkhagvasuren, and Roman Sysuyev**, “Systematic Job Search: New Evidence from Individual Job Application Data,” Technical Report 12-03r, Federal Reserve Bank of Richmond 2014.
- Li, Shengwu and Ning Neil Yu**, “Context-dependent choice as explained by foraging theory,” *Journal of Economic Theory*, 2018, 175, 159–177.
- McCall, John Joseph**, “Economics of information and job search,” *The Quarterly Journal of Economics*, 1970, pp. 113–126.
- Milgrom, Paul and Chris Shannon**, “Monotone comparative statics,” *Econometrica: Journal of the Econometric Society*, 1994, pp. 157–180.
- Mortensen, Dale T.**, “Job Search, the Duration of Unemployment, and the Phillips Curve,” *The American Economic Review*, 1970, 60 (5), 847–862.
- Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa**, “Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence, and bias,” *American Economic Review*, 2021, 111 (1), 324–363.
- Ortoleva, Pietro**, “Modeling the change of paradigm: Non-Bayesian reactions to unexpected news,” *American Economic Review*, 2012, 102 (6), 2410–2436.
- Phillips, Lawrence D and Ward Edwards**, “Conservatism in a simple probability inference task.,” *Journal of experimental psychology*, 1966, 72 (3), 346.

- Potter, Tristan**, “Learning and job search dynamics during the Great Recession,” *Journal of Monetary Economics*, 2021, *117*, 706–722.
- Rabin, Matthew**, “Inference by believers in the law of small numbers,” *The Quarterly Journal of Economics*, 2002, *117* (3), 775–816.
- Rosenfield, Donald B and Roy D Shapiro**, “Optimal adaptive price search,” *Journal of Economic Theory*, 1981, *25* (1), 1–20.
- Rothschild, Michael**, “Searching for the lowest price when the distribution of prices is unknown,” in “Uncertainty in Economics,” Elsevier, 1978, pp. 425–454.
- Sadish, D, Achyuta Adhvaryu, and Anant Nyshadham**, “(Mis) information and anxiety: Evidence from a randomized Covid-19 information campaign,” *Journal of Development Economics*, 2021, *152*, 102699.
- Spinnewijn, Johannes**, “UNEMPLOYED BUT OPTIMISTIC: OPTIMAL INSURANCE DESIGN WITH BIASED BELIEFS,” *Journal of the European Economic Association*, 2015, *13* (1), 130–167.
- Talmain, Gabriel**, “Search from an unknown distribution an explicit solution,” *Journal of Economic Theory*, 1992, *57* (1), 141–157.
- Weitzman, Martin L**, “Optimal search for the best alternative,” *Econometrica: Journal of the Econometric Society*, 1979, pp. 641–654.
- Wiswall, Matthew and Basit Zafar**, “Determinants of college major choice: Identification using an information experiment,” *The Review of Economic Studies*, 2015, *82* (2), 791–824.