Appendix B Simulations for Convergence of Beliefs

Figure 10 presents a simulation result for the dynamics of beliefs for a sequence of signals that are drawn randomly. In this simulation, there are two possible states $\{H, L\}$ and the agent begins with a prior of 0.5 that the state is H. In each period, the agent observes an independently drawn signal that predicts the state accurately 75% of the time. Supposing the drawn state is H, we randomly draw a sequence of signals and plot the beliefs of various updating rules over time.

Overreaction and base rate neglect are modeled using the Grether (1980) model, with $a = 0.5$ for the agent with base rate neglect and $b = 1.5$ for the agent with overreaction. For the incorrect prior, the starting prior was 0.25 instead of 0.5, and beliefs are updated with Bayes' rule. For the asymmetric updating pattern, we use the Hagmann and Loewenstein (2017) model, setting the reference belief to 0 and λ to 0.75. The weight attached to the reference belief is 0.25.

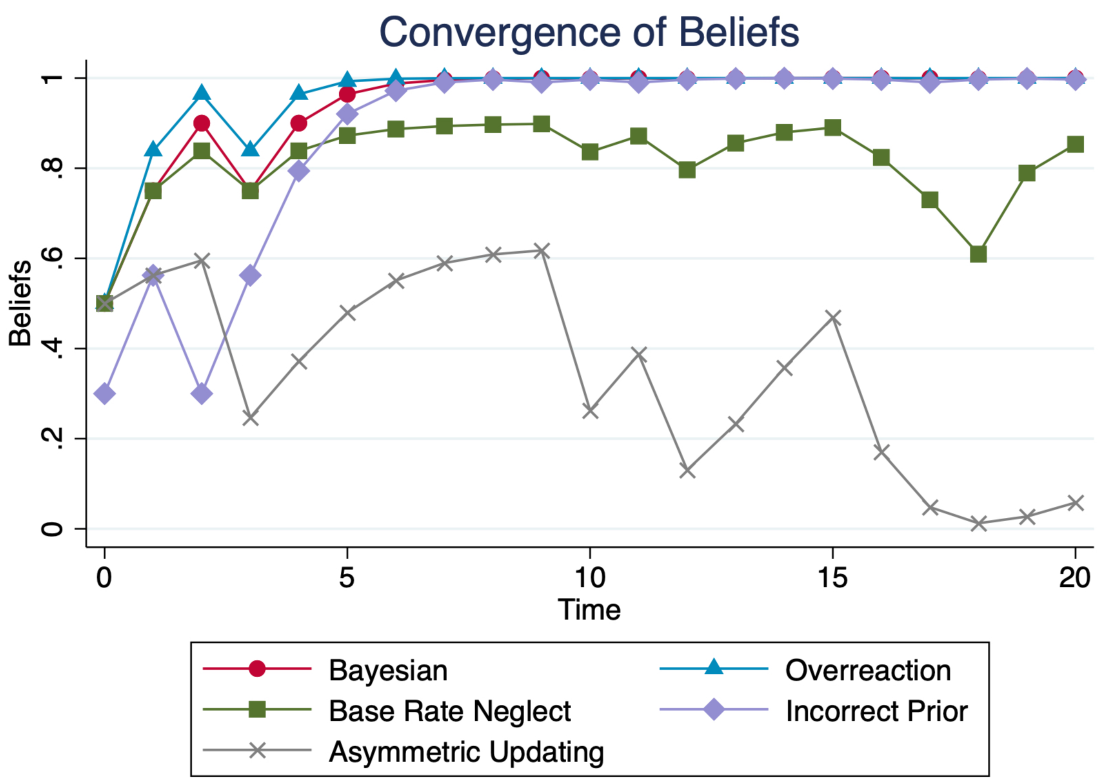

Figure 10: Simulation result showing the dynamics of beliefs.

Appendix C Alternative Posterior Fittings

C.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.

C.1.1 Log Normal Fitting, Recovered Individual Offer Distribution

The first alternative 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 nth 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.

C.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.

C.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.

C.2 Model Fit and Estimate Results by Method

The table below gives the average fitting error (mean squared 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. The kernel-based methods have worse fits than the other methods by a noticeable margin. The fitting errors are very similar whether we consider only the first updates or the full dataset.[50] Therefore, our ability to fit posteriors seems about the same for both second and third surveys.

Distribution	Update	MSE Average	MSE Standard Deviation
Log normal,	1	.0055	.0106
Best Wage Offer	2	.0053	.0106
Log normal,	1	.0052	.0102
Recovered Single Wage Offer	2	.0050	.0101
Gumbel	1	.0056	.0129
	2	.0053	.0131
Midpoint Kernel,	1	.0714	.0421
Restricted	2	.0721	.0412
Midpoint Kernel,	1	.0630	.0456
Unrestricted	2	.0638	.0446

Table 9: Average mean squared error by data fitting method and update number.

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. Finally, we note that the main results we get are very similar whether we use direct survey responses from OO2b or recover the individual wage distribution. Thus, our result of non-Bayesianism appears robust to concerns over the number of offers expected by subjects in our sample.

[50] First updates include individuals without second updates.

Statistic	Log Normal, Best Wage Offer	Log Normal, Recovered Single Wage Offer	Gumbel	Midpoint Kernel, Restricted	Midpoint Kernel, Unrestricted
$\overline{m}$	.9363	.9338	.9417	.4220	.4717
	(.0193)	(.0215)	(.0210)	(.0055)	(.0069)
$\overline{r}$	.1807	.1656	.1898	.3146	.2991
	(.0079)	(.0076)	(.0080)	(.0047)	(.0055)
X	.7556	.7682	.7519	.1073	.1726
	(.0209)	(.0224)	(.0216)	(.0101)	(.0118)
X_{norm}	5.1824	5.6383	4.9604	1.3412	1.5771
	(.2501)	(.2825)	(.2253)	(.0371)	(.0498)

Table 10: Excess belief movement test statistics using alternate distribution fittings. Standard errors clustered by state are in parentheses.

Appendix D Additional Figures and Tables

Figure 11 below shows the relationship between the excess movement statistics and the absolute error in the prior beliefs. We can see that the absolute error is increasing and convex in the excess belief movement statistics. This tells us that by assuming homogeneity in the measurement error, this gives us the upper bound for the amount of measurement error needed to be consistent with the data.

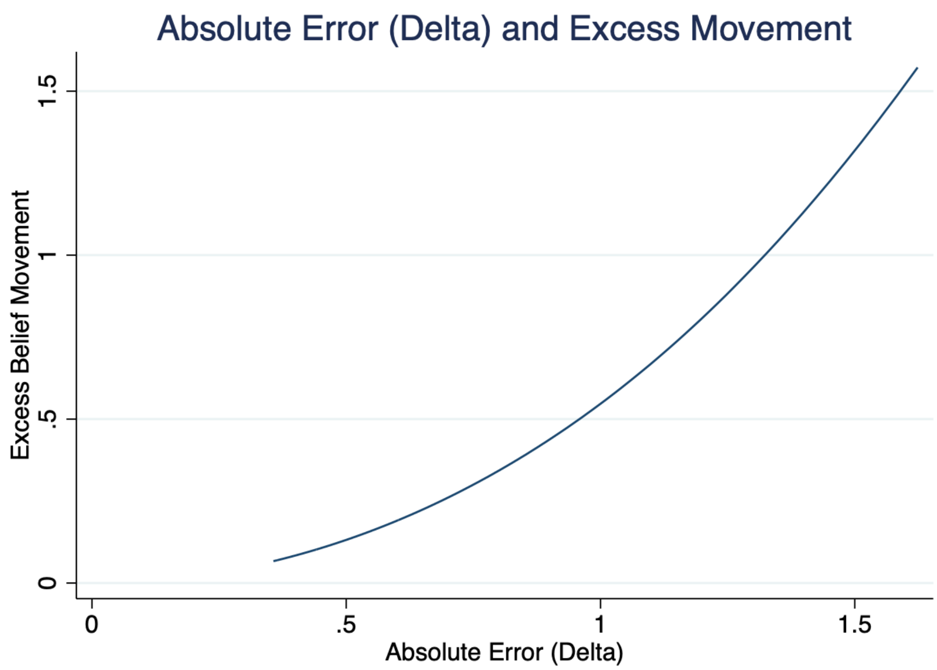

Figure 11: Relationship between Δ for a Dirichlet distribution centered around the uniform distribution and excess belief movement statistics.

Figure 12 below plots the normalized change in beliefs for the survey respondent's best wage offer. The normalized change in belief is defined as

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

y_1 and y_2 is the subject's expected the best wage offer they could earn over the next four months on survey 1 and 2 respectively (that is, the response to survey question "OO2a2"). $\overline{x}_2$ is the best wage offer received on survey 2.

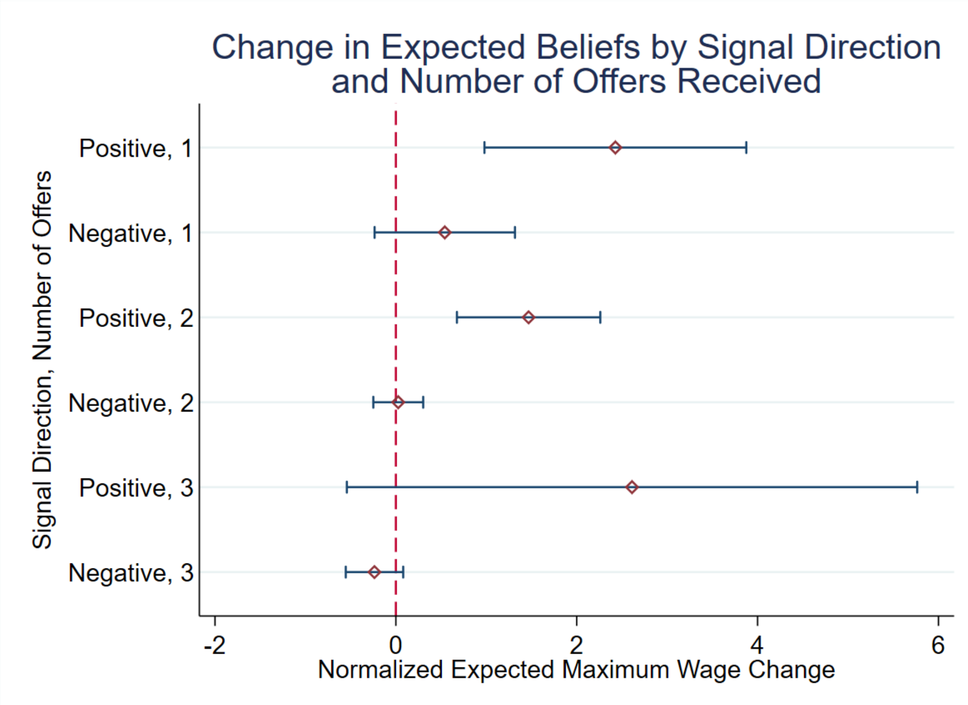

Figure 12: Normalized Expected Maximum Wage Changes by Signal Direction. Belief variable here is the maximum annual salary the survey respondent expects. Error bars display 95% confidence intervals around the mean.

Appendix E Excess Belief Movement Test

The table below replicates the excess belief movement test using only the responses from only the first two surveys. Although the excess movement statistics are lower than when using all three surveys, they remain large and statistically significant. Therefore, our main results are robust to using a shorter time horizon where it is more plausible that an individual's underlying wage offer distribution did not change.

Statistic	All Individuals	Got Offer?		Searched?		
		Yes	No	Yes	No	Unknown
$\overline{m}$	.6642	.7357	.6486	.6789	.6530	.7150
	(.0149)	(.0319)	(.0162)	(.0370)	(.0136)	(.0427)
$\overline{r}$	.1719	.2000	.1658	.1856	.1645	.1941
	(.0061)	(.0175)	(.0062)	(.0128)	(.0066)	(.0254)
$X = \overline{m} - \overline{r}$	.4923	.5357	.4828	.4933	.4886	.5209
	(.0161)	(.0379)	(.0173)	(.0352)	(.0155)	(.0488)
$X_{norm} = \frac{\overline{m}}{\overline{r}}$	3.8634	3.6783	3.9122	3.6571	3.9705	3.6833
	(.1619)	(.3733)	(.1754)	(.2672)	(.1843)	(.5210)
p-value of t-test: $X = \overline{m} - \overline{r} = 0$	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
Observations	2,456	440	2,016	588	1,671	197

Table 11: Excess movement statistics: Surveys 1 and 2 only. *Notes: Clustered errors by state in parentheses. Some of the subjects in our sample did not answer the question of whether they searched, but we still include them here as they are included in our full sample estimate.*

The following table presents our main Martingale test results including observations which were dropped from our final sample due to data quality or offer distribution stability concerns. The excess movement statistics remain large and significant under either specification.

Statistic	All Individuals	Got Offer?		Searched?		
		Yes	No	Yes	No	Unknown
$\overline{m}$	.9710	1.0863	.9309	1.0524	.9288	1.0177
	(.0166)	(.0328)	(.0195)	(.0315)	(.0200)	(.0431)
$\overline{r}$	.1950	.2225	.1855	.2069	.1879	.2103
	(.0065)	(.0146)	(.0070)	(.0104)	(.0071)	(.0201)
$X = \overline{m} - \overline{r}$	.7759	.8638	.7454	.8455	.7409	.8074
	(.0179)	(.0315)	(.0217)	(.0317)	(.0217)	(.0443)
$X_{norm} = \frac{\overline{m}}{\overline{r}}$	4.9784	4.8822	5.0194	5.0876	4.9425	4.8389
	(.1876)	(.3086)	(.2305)	(.2779)	(.2217)	(.4709)
p-value of t-test: $X = \overline{m} - \overline{r} = 0$	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
Observations	3,450	879	2,570	992	2,200	258

Table 12: Excess movement statistics: All observations included. *Notes: Clustered errors by state in parentheses. These results include individuals who listed annual wages or expectations under $10,000, individuals who had one of their surveys in 2020, and individuals who moved (between states) between surveys.*