Discrete Adjustment to a Changing Environment: Experimental Evidence*

Mel Win Khaw
Columbia University

Luminita Stevens
University of Maryland

Michael Woodford
Columbia University

November 9, 2016

Abstract

We conduct a laboratory experiment to shed light on the cognitive limitations that may affect the way decision makers respond to changes in their economic environment. The subjects solve a tracking problem: they estimate the probability of drawing a green ring out of a box with green and red rings, which changes stochastically. The subjects draw rings from the box and indicate their draw-by-draw estimate. Our subjects depart from the optimal Bayesian benchmark in systematic ways, but these deviations are not simply the result of some boundedly rational, but deterministic rule. Rather, there is a random element in the subjects' response to any given history of evidence. Moreover, subjects adjust their forecast in discrete jumps rather than after each new ring draw, even though there are no explicit adjustment costs. They adjust by both large and small amounts, contrary to the predictions of a simple *Ss* model of optimal adjustment subject to a fixed cost. Finally, they prefer to report "round number" probabilities, even though that requires exerting additional effort. We develop a model of inattentive adjustment and compare its quantitative fit with alternative models of stochastic discrete adjustment.

*We would like to thank Fernando Alvarez, Marty Eichenbaum, Cosmin Ilut, Filip Matejka, Alex Wolman, and conference and seminar participants at the Barcelona GSE Summer Forum, the NBER Summer Institute, the University of Maryland, and the JME-SNB-SCG conference for useful comments, and the National Science Foundation for support of this research.

1 Introduction

A central problem in macroeconomics is understanding the ways in which households and firms respond to changing market conditions, with a particular emphasis on how the immediate (or relatively short-run) effects of new developments differ from the adjustments that eventually occur. Often, the immediate responses to disturbances seem inconsistent with optimization under full information, in the absence of some impediments or costs of more rapid adjustment of behavior. As a consequence "micro-founded" macro models often assume costs of adjustment of a variety of types (or "habit-persistence" in preferences) in order to better fit the observed adjustment dynamics, even though direct evidence for the existence of external costs (i.e., costs unrelated to costs of information processing or deliberation on the part of the unit making the decision) is often lacking. When behavior is observed at the micro level, continuous decision variables (such as the price that a firm charges for a given product) are often observed to change only at discrete points in time, even though relevant market conditions are changing continuously; this is commonly interpreted as a response to the existence of non-convex costs of adjustment (such as "menu costs" in the case of prices), though again there is often little direct evidence on the magnitude of such costs.

An alternative potential explanation attributes delays in adjustment to imperfect attention by the decisionmakers, even when the information about the changing conditions is in principle available to the agents. A variety of models of noisy information (Woodford, 2003), sticky information (Mankiw and Reis, 2002; Reis, 2006), or rational inattention (Sims, 2003, 2011; Mackowiak and Wiederholt, 2009; Woodford, 2009) have been proposed (to cite only examples from the literature on price adjustment). In these cases, given that the postulated costs or constraints are internal to the minds of the decisionmakers, there are even harder questions about how to independently validate the realism of such assumptions.

This paper attempts to shed light on the empirical relevance of cognitive limitations as a source of both adjustment delays in general, and discreteness of adjustment more specifically, by conducting a laboratory experiment. Our experimental setting is very simple, and there are obvious questions about the similarities between the task faced by our subjects and those faced by decisionmakers in settings of relevance for macroeconomic modeling (such as pricing

decisions on the part of firms). Nonetheless, a laboratory experiment also has important advantages. We have significant control over the degree to which any non-trivial external costs of adjustment exist, and we can treat the decisionmaker's objective as known (that is, as specifiable as part of the experimental design), to the extent that we assume that our subjects care only about maximizing their monetary payment from the experiment and understand the consequences of their actions.[1]We can also be quite certain about exactly what information the decisionmaker has at each point in time, to the extent that this is a question about her objective circumstances; not only do we know everything that the experimental subject has had an opportunity to observe, but we know what she ought to understand about the data-generating process, and hence what prior probabilities over possible situations it would be rational to use in drawing inferences from what is observed during the experiment. Thus we need not make (and have no warrant for making) arbitrary conjectures either about decisionmakers' objectives or about the information that is (objectively) available to them, and we can focus on what our results indicate about apparent cognitive limitations.

Our experiment (described in more detail in the next section) is a forecasting exercise in which one of two outcomes can occur on each trial, and the subject must estimate the probability of one outcome (drawing a green ring rather than a red one). It is explained that the underlying probability will shift from time to time, but also that it is likely to remain the same for many successive trials, offering an opportunity to estimate the current probability from observation of past draws. This kind of exercise allows us to test alternative theories of expectation formation, including the familiar benchmark of "rational expectations" (optimal Bayesian inference from the available information).

At the same time, the subjects' problem can be viewed as an example of a more general class of problems, in which there is a continuous decision variable (here, the announced probability estimate), with the payoff from a given action depending on the current value of continuous state variable (here, the true probability of drawing a green ring), which varies over time (so that it is necessary to continue monitoring), but is sufficiently persistent to make attempts to monitor the changing state variable worthwhile. Viewed in this way, it is

[1]In the experiment reported here, the subjects are paid for their performance, and we take care to explain the strategy that would maximize the expected monetary reward.

an example of the same basic kind of decision problem faced by a firm that must set a price for its product, where the profits obtained by charging a given price depend on other state variables (demand conditions and factor costs) that fluctuate over time.

In fact, despite the differences in the setting (our subjects are not asked to set a price, there are no customers, etc.), our experimental data exhibit some notable (and puzzling) features of data on individual prices (as we discuss further in Section 3 below). In particular, our subjects usually leave their decision variable (the slider position indicating their probability estimate) unchanged for periods of time, despite the receipt of many new pieces of information in the meantime, after each of which their optimal estimate should have changed somewhat. This discreteness of the observed adjustments occurs (and is common to all of our subjects, albeit to varying extents) despite the fact that there is no discreteness of the set of available actions, nor any fixed cost of adjustment (other than a possible psychic cost). Since there are no true (external) adjustment costs in our problem, and subjects have (and we believe, are able to understand) all of the information needed to calculate the optimal Bayesian estimate at each point in time, we conclude that their failure to track the optimal Bayesian estimate more closely reflects imperfect attention, limited memory or related cognitive limitations.

The pattern of discrete adjustments observed also differs greatly from the one that would be predicted by an "*Ss* model" of optimal adjustment subject to a fixed cost. We observe adjustments of all sizes, including many small ones as well as a certain number of very large ones. Instead, the "*Ss* model" would predict that no adjustments occur until a particular threshold is crossed (so that all adjustments should be of a certain minimum size), and that adjustment always occurs as soon as the threshold is crossed (so that adjustments larger than a certain magnitude should never be needed). In addition, we observe that some adjustments occur very soon after the last adjustment, even though the median time lag between adjustments is much longer, an observation that is also relatively unlikely in data generated by an "*Ss* model." All of these unusual features of our data (relative to the predictions of the "*Ss* model") are also features of distributions of individual price changes and of the time intervals for which individual prices remain fixed.

Because in our case, we also know the optimal Bayesian estimate at each point in time,

we can check the prediction of the "*Ss* model" (along with many other models of discrete adjustment) that when adjustment occurs, the decision variable is set to its optimal value at that time (even though it soon ceases to be optimal, as additional information arrives). In fact, we find that the estimates chosen by our subjects are quite inaccurate even at points in time when they adjust their reported estimates. Nevertheless, the discrepancy between the currently optimal estimate and their own existing estimate does help to predict both how they will adjust their estimate if they adjust it, and the likelihood that they will choose to adjust it. This means that even when adjustments occur, subjects' choices continue to reflect some degree of inattention or related cognitive limitations.

We explore the ability of a variant of the theory of "rational inattention" proposed by Sims (2003, 2011) to account quantitatively for our subjects' behavior. Rational inattention provides an optimization-based framework in which to generate the kind of noisy actions observed in our data. Our model generalizes the model of inattentive discrete adjustment developed in Woodford (2009), most importantly by allowing not only the timing of adjustments but also the choice of where to set the decision variable conditional on adjustment to be inattentive; this makes our model a close cousin of the model of "logit dynamics" proposed by Costain and Nakov (2014). We compare the quantitative fit of our rational inattention model with other models of discrete adjustment that also allow the timing of adjustments to be probabilistic, namely a version of the "generalized *Ss* model" of Caballero and Engel (1993, 2007) and an optimizing model with a random fixed cost of adjustment of the kind proposed by Dotsey, King and Wolman (1999) and Dotsey and King (2005). While the latter models can describe the adjustment dynamics that we observe better than a classic "*Ss* model," we find that the rational inattention model outperforms these alternatives significantly.

Our experimetal design is complementary to that of Magnani, Gorry and Oprea (2016), who study how well Ss models versus time-dependent models can fit laboratory data on the decision of whether or not to adjust. Their experimental setup seeks to mimic the conditions assumed in Ss models of adjustment: subjects are shown the optimal action at each point in time, and they can adjust their estimate to that optimum by simply pressing a button; however, each time they do, they incur a fixed cost. Subjects' behavior departs from the

predictions of the Ss model, and is instead consistent with a control cost model for the decision of whether or not to adjust.

Section 2 describes our experiment. Section 3 gives an overview of some of the salient features of the behavior that we observe, focusing on ways in which subjects' behavior resembles or differs from the predictions of the ideal Bayesian (or rational expectations) benchmark. Section 4 discusses the extent to which models of discrete adjustment that have been common in the macroeconomic literature, especially the literature on price adjustment, can account for our data. Section 5 then develops our rational inattention model, and tests its ability to explain the features of our data. Section 6 offers a concluding summary of what we can learn from the experiment.

2 The Experimental Design

Our laboratory experiment follows the setup of Gallistel, Krishan, Liu, Miller & Latham (2014), who study how well subjects can predict probabilities. We adapt their experiment to our setting as detailed further below. The subjects' task is to estimate the probability that the next draw is a green ring from a black box that contains both green and red rings. The subjects draw rings with replacement from the box and indicate their draw-by-draw estimate of this probability.

Figure 1 shows a screenshot of the experiment. The screen displays the box with a hidden number of red and green rings on the left. The slider at the bottom of the screen indicates the subject's estimate on the current trial, $\hat{p}_t$, with the number in percentage points displayed above the bar. The box on the right side of the screen displays 1,000 rings that also depict the subject's estimate visually. Whenever the subject moves the slider, this box is adjusted in real time to reflect the probability indicated by the slider position. The subject begins with a guess and adjusts the position of the slider to indicate his or her estimate. Each time the subject clicks the NEXT button, a new ring is drawn randomly from the box and displayed on the screen. The subject may then adjust his or her forecast or leave it unchanged. After each ring draw, the subject's cumulative score is updated and displayed at the top of the screen. This process is repeated until the session ends, after 1,000 ring draws.

For each session, the subject is rewarded with a fixed payment of $10 plus a variable

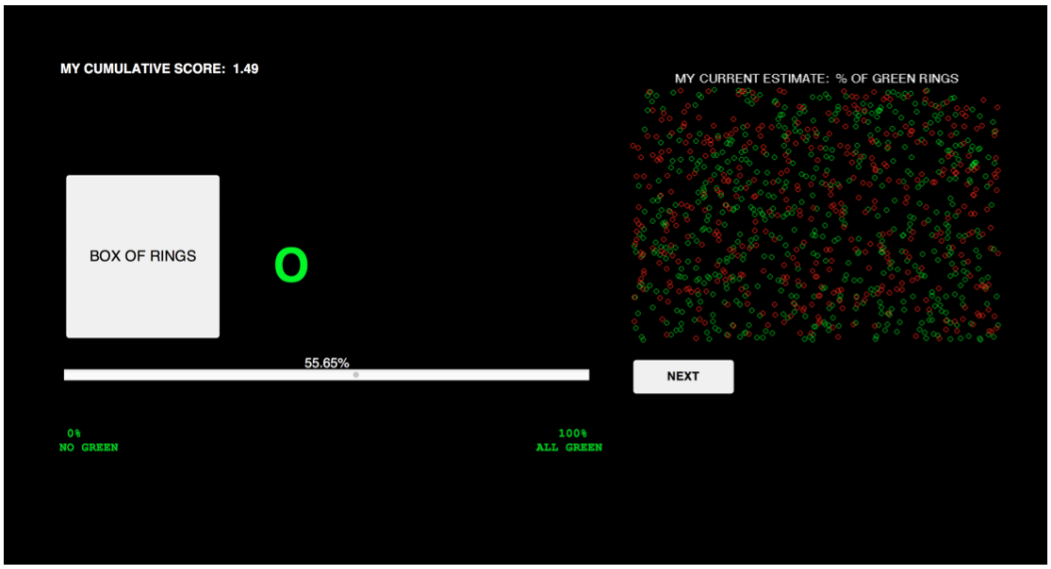


Figure 1: Experiment screenshot. The box of rings contains an unknown number of green and red rings. Subjects use the mouse to adjust the slider to their current estimate of the probability that the next draw from the box is a green ring. Their estimate is displayed numerically above the slider and visually both using the slider and in the box on the right side of the screen. Subjects click the NEXT button to draw another ring from the box of rings on the left. After each draw, the subject's cumulative score is updated and displayed at the top of the screen. Each session consists of 1,000 ring draws.

payment, equal to 2¢ × the subject's cumulative score. The score is cumulated as follows: for each trial t, if the subject sets the slider at $\hat{p}_t$ and the ring drawn is s_t (equal to 1 for a green ring and 0 for a red ring), then the reward is

$$r(\hat{p}_t; s_t) = 1 - (s_t - \hat{p}_t)^2. \tag{1}$$

The total score is the sum of the rewards for all trials in a session. This payoff structure generates a non-negative score that is increasing over the life of the session, to avoid loss aversion. If one believes that a green ring will be drawn with probability p_t (however this forecast might be formed), then the expected reward, $p[1 - (1 - \hat{p})^2] + (1 - p)[1 - \hat{p}^2]$, is maximized by setting the slider at $\hat{p}_t = p_t$.[2] The monetary reward is only a function of the history of the subject's behavior and the sequence of draws thus far. It reveals no information about the number of green rings in the box, either past or present. Before the start of the sessions, we explained to the subjects both the reward function and the optimal decision

[2]We assume that the subjects are risk neutral for the small monetary rewards involved in the experiment.

rule, using both written and verbal instructions. The subjects were told that to maximize their expected reward, the best thing that they can do on each trial is to always set the slider to coincide with their best guess about the probability that the next draw will show a green ring.

From time to time, the box of rings is replaced by a new box with a different probability of drawing a green ring. The probability is initially drawn from a uniform distribution between 0 and 1. After each ring draw, there is a constant probability $\delta = 0.5\%$ of a regime change. If there is a regime change, the new proportion of green rings in the box is an independent draw from the uniform distribution on the unit interval. The subjects are not told when a change in the box occurs, but they are told in advance that the box might change, and they know the probability of a regime change and the distribution from which the new box is drawn. The frequency of regime changes is small enough relative to the number of noisy observations that subjects receive, to enable them to learn the underlying probability. But it is nonetheless changing, so that the subjects should continuously take into account the fact that there may have been a regime change.

It is important to note that the timing of slider adjustments is completely up to the subject. We also allow the subjects to control the speed at which they draw rings so that they may choose how much attention to give to the task, rather than forcing them to make decisions within an externally imposed time limit. This setup allows us to investigate the extent to which the subjects choose to update their forecast every time they obtain a new piece of information. There is no explicit penalty for constant adjustment and new information arrives frequently, which puts the finding of discrete adjustment at odds with a large set of models, as we discuss in the next section.

We modified the experiment of Gallistel et al. (2014) along five dimensions, to make it more suitable to the focus of this paper: (1) we introduced a monetary reward that varies with performance, (2) we simplified the data generating process and the way we elicit responses from the subjects, to ensure the tractability of the Bayesian inference problem, (3) we added a numerical representation of the subject's estimate, in addition to the visual one, (4) we eliminated the possibility of discrete step sizes of the slider, ensuring smooth adjustment of the slider, and (5) we explained all features of the setup and the data generating

process to the subjects. The reward function allows us to have more control over the objective of our subjects, rather than having to make assumptions about what they might be maximizing. We incentivized the subjects to do a particular task and we also displayed their cumulative score in real time, so they could see how their strategies are affecting their reward. We showed the subjects both a visual and a numerical representation of their current estimate for multiple reasons. On the one hand, some subjects may have lower numeracy than others. Additionally, it has been shown that probabilities are notoriously difficult to interpret and communicate without some visual representation. The box showing them their current estimate helps the subjects visualize exactly what they are trying to estimate. On the other hand, subjects may want to know what probability they are setting with more precision than that allowed by the visual representation. In fact, displaying the numerical representation allows us to uncover evidence of a cognitive bias that we discuss in the next section. We also explained all features of the setup to the subjects, so that to the extent they have rational expectations, they have all the information they need in order to form the optimal RE forecast. This enables us to compare the subjects' performance to that of the ideal Bayesian decisionmaker, which we do in the next section.

Our setup is designed to minimize adjustment costs and to eliminate reasons for habit preferences, such that there would be no reason to not have constant adjustment to the currently optimal forecast in the absence of cognitive limitations. Additionally, we control fully the information that subjects have at each point in time, which allows us to test to what extent subjects are processing the information that is available to them. Finally, by controlling the true data generating process and what our subjects know about it, we control what the correct prior is. Thus we do not need to make conjectures about either the decisionmakers' objectives or the information that is available to them, and we can focus instead on what our results indicate about potential cognitive limitations. The setup allows us to test alternative theories of expectation formation, including the rational expectations benchmark (optimal Bayesian inference using all the available information).

An important question that arises in any experimental study is that of external validity: To what degree can we expect the behavioral regularities that we uncover in the experimental data to persist in other economic environments? The answer to this question depends

on the degree to which the conditions that are relevant remain relatively unchanged across environments. In our particular setting, even though we place our subjects in a simple environment, the task that they face is an example of a very general class of *tracking problems*. Formally, it is analogous to the problem of a pricing manager who sets the price of a product and whose task is to track the profit-maximizing price as closely as possible, subject to any frictions and limitations that might impede adjustment. There are some key features of that problem that also apply to our experimental task: it features a continuous decision variable, the resulting profits depend on other state variables (demand, costs) that fluctuate continuously over time, and finally, the manager making the decision receives a steady stream of information about product demand, competitive pressures, cost conditions, the macroeconomy, etc., and incorporates these factors in the price of the product.

We conducted the experiment with 11 subjects, who each completed 10 sessions over the course of several days. On average, each session lasted 27 minutes. The subjects were undergraduate and Master's students at Columbia University. At the start of each session, the subjects received both written and verbal instructions and they completed a 15-minute practice session to familiarize themselves with the task.

3 Features of Observed Behavior

Figure 2 shows an example of a typical session for a particular subject. The figure plots the true hidden probability and the subject's forecast, compared to that of a Bayesian decisionmaker who makes full use of all the information available.[3] There are 1,000 ring draws and several regime changes over the course of the session.

The figure illustrates some of the key features of the experimental data that we document more systematically below. In particular, both the subject and the benchmark Bayesian decisionmaker track the hidden probability quite well, and they detect changes in this probability relatively quickly. However, one major distinction is that while the Bayesian forecast changes by a small amount after each ring draw, the subject keeps his estimate unchanged for long periods of time. This discreteness arises despite the fact that there is no discreteness in the set of available actions, nor any explicit cost of adjusting the slider in response to the stream

[3]The appendix presents the Bayesian inference problem.

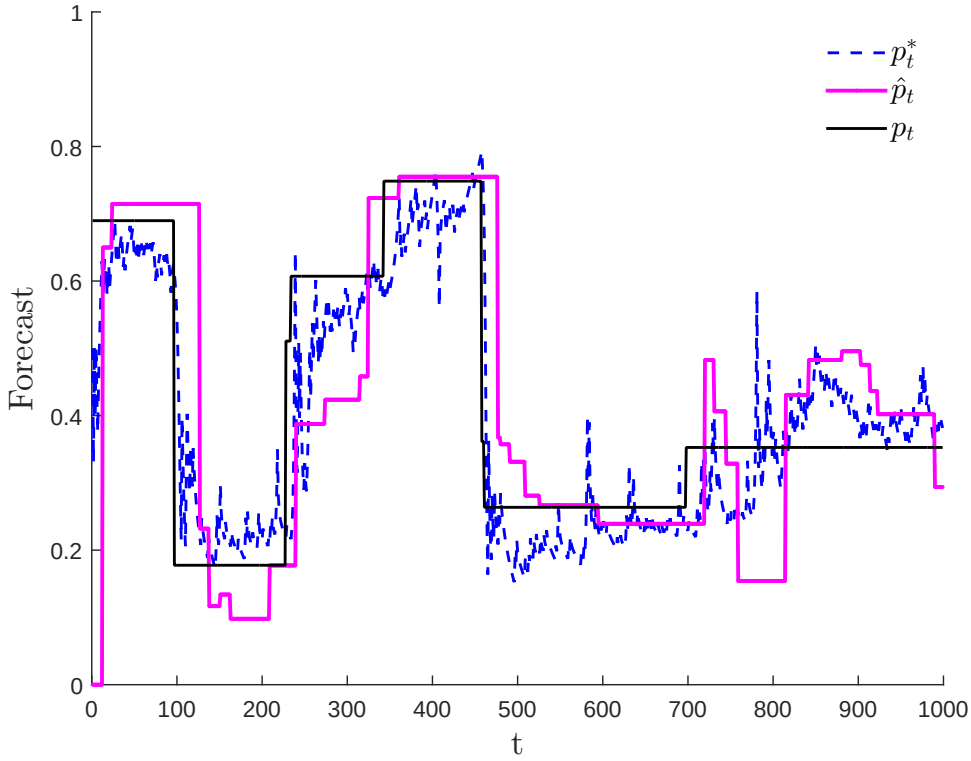


Figure 2: Experiment data for subject 11, session 10. The benchmark Bayesian forecast (blue dashed line) and the subjective estimate (magenta solid) are plotted against the true hidden probability (in black).

of new information. Despite the relatively infrequent adjustment, the figure shows both big and small slider movements. Overall, the subject makes much larger adjustments on average than the Bayesian decisionmaker. In addition, adjustments both big and small occur very soon after the last adjustment, even though the typical time lag between adjustments is much longer, a feature that is unlikely in data generated by models of optimal adjustment subject to some fixed cost. With these observations in mind regarding lumpiness in both the timing and the size of adjustments, we next assess the subjects' performance more systematically.

3.1 Overall Accuracy

Our subjects demonstrate the ability to track the varying state, although less well than would be possible given the available information. As described in the previous section, the subjects are rewarded using a function that depends on their mean square error. Figure 3 compares the scores of the subjects and the Bayesian decisionmaker to two polar benchmarks: the *complete information* case in which the subject knows the hidden probability at all times,

and the *no information* case in which the subject is fully rational but does not see any rings and therefore sets the slider to the unconditional estimate of 0.5 on all draws. For each observer type, the figure shows 11 bars, each corresponding to the realized draws seen by each of our 11 subjects across all sessions.

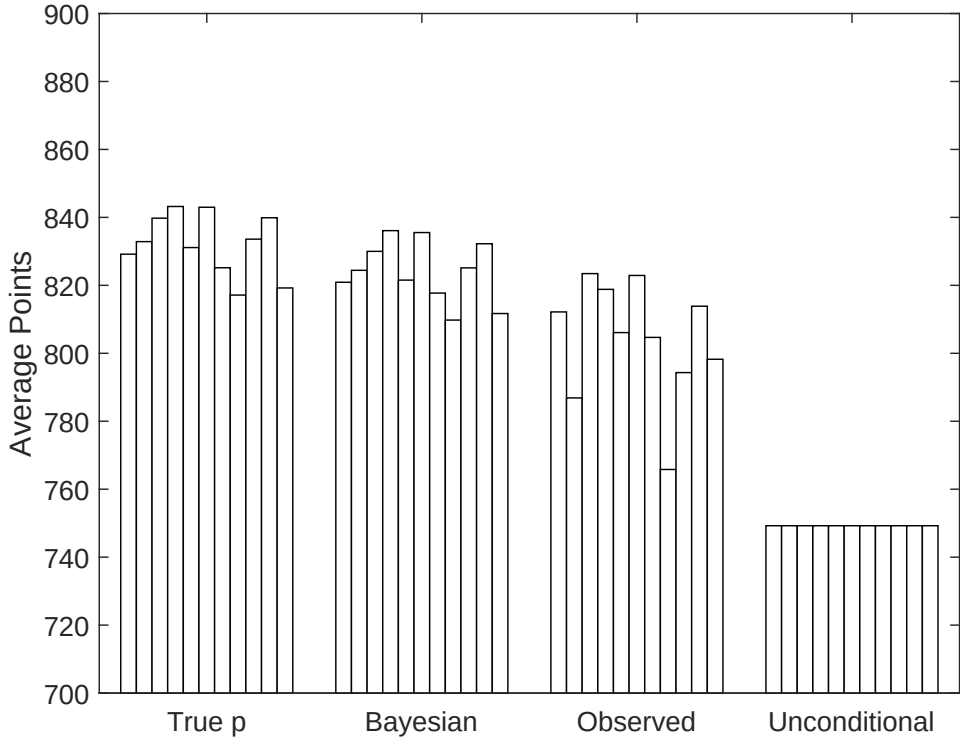


Figure 3: Experiment data for all subjects and sessions. The scores are broken down by subject. "True p" refers to the score that each subject would have received, given the realized ring draws, had they known the underlying probabilities at all times. "Bayesian" refers to the scores that would have been received by each subject, given the realized ring draws, had they acted like the fully rational optimal Bayesian decisionmaker. "Unconditional" refers to the scores received under a constant forecast equal to the unconditional prior of 0.5.

All the subjects outperform the *no information* benchmark, indicating that they are tracking the state at least to some extent. They all underperform the Bayesian decisionmaker: on average, the subjects' scores are 2.5% lower than that of the Bayesian observer. There is considerable heterogeneity in performance. One driver of this heterogeneity is the noise in the realized draws, which determines the difficulty of the forecasting problem. This can be seen in the differences in scores in the full information benchmark. Another driver of the heterogeneity in scores is heterogeneity in behavior across the subjects. Some subjects (such as subject 3) are much closer to the fully rational Bayesian benchmark than others

(for example, subject 8).

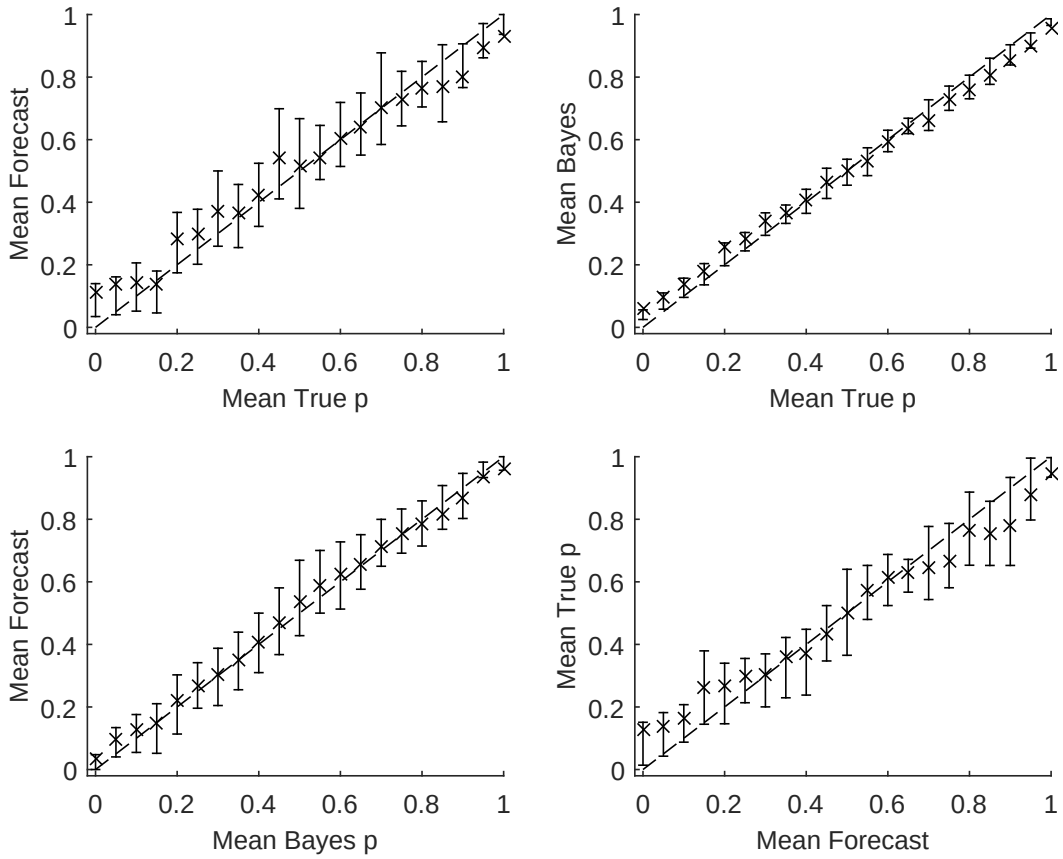


Figure 4: Measures of forecast bias. Upper left: $\mathrm{E}[\hat{p}|p]$ versus p. Upper right: $\mathrm{E}[p^*|p]$ versus p. Lower left: $\mathrm{E}[\hat{p}|p^*]$ versus p^*. Lower right: $\mathrm{E}[p|p^*]$ versus p^*.

3.2 Forecast Bias

A relatively simple way of characterizing the degree to which behavior resembles or differs from the rational (perfect Bayesian) benchmark is to ask to what extent forecasts correctly track the true state on average. One familiar diagnostic considers whether forecasts are *unbiased*, in the sense that $\mathrm{E}[\hat{p}_t|p_t]$ is equal to p_t, the true state. Gallistel *et al.* (2014) emphasize the extent to which their subjects are close to perfectly unbiased, though they compare the median forecast conditional on the true p, rather than the mean.[4] Robinson (1964) had similarly reported no significant forecast bias (using the more conventional definition based on the conditional mean) for any value of p, and this had been confirmed by

[4]Ricci and Gallistel (2016) report a similar finding in a variant of the experiment of Gallistel *et al.* (2014) with a different data-generating process.

Regression: $y = \alpha + \beta x + \epsilon$				
y	x	α	β	F-stat
$\hat{p}$	p	0.086	0.849	4056
		(0.001)	(0.002)	
p^*	p	0.054	0.894	6901
		(0.001)	(0.001)	
p	$\hat{p}$	0.089	0.818	6273
		(0.001)	(0.002)	
$\hat{p}$	p^*	0.033	0.952	647
		(0.001)	(0.002)	

TABLE I: Regression tests of forecast bias, with alternative choices of the variables y and x, corresponding to the four panels of Figure 4. Standard errors are shown in parentheses below the regression coefficients. The F statistic in each case is for a test of the null hypothesis that $\alpha = 0$ and $\beta = 1$, and under the null hypothesis should be distributed as $F(2, 109888)$. All null hypotheses have p-values of less than 10^{-279}.

additional studies summarized in Peterson and Beach (1967).

We examine this question in the case of our data in the upper left panel of Figure 4. We pool the data from all subjects and sessions, and sort the data into 21 bins centered on multiples of 0.05, according to the true probability p_t on that trial. For each bin, the plot shows the mean slider position $\hat{p}_t$ across those trials (as an x), together with the inter-quartile range of slider positions associated with that bin (shown as a vertical interval).[5] As in the corresponding figure in Gallistel *et al.*,[6] the diagonal (i.e., the predicted location of the means under the hypothesis of unbiasedness) is contained within the IQR in nearly every case, over the entire range of true probabilities. However, it is not quite right to conclude from this that the average slider position, conditional on any true state, is exactly that value. As indicated in Table I, a linear regression of p_t on $\hat{p}_t$, using our pooled data, yields a regression line with a slope slightly less than 1; and the null hypothesis that the regression line is the diagonal can be rejected with a p-value that is indistinguishable from zero (at the precision allowed by Matlab).

Unbiasedness in this sense is not, however, the only hypothesis of interest, and is not even a prediction of rationality. Even an ideal Bayesian observer would not be able to track

[5] The horizontal coordinate of the x that is plotted for each bin is the mean value of p_t for the trials in that bin. Under the hypothesis that $\mathrm{E}[\hat{p}_t|p_t] = p_t$ for all values of p_t, we should also observe that $\mathrm{E}[\hat{p}_t|p_t \in b_j] = \mathrm{E}[p_t|p_t \in b_j]$, for any bin b_j, i.e., that the X for bin b_j should lie on the dashed diagonal line.

[6] See their Figure 6.

the hidden true probability perfectly, and as shown in the upper right panel of the figure, the forecasts of an ideal Bayesian observer would not be "unbiased" in the sense just mentioned: we should not expect $\mathrm{E}[p_t^*|p_t]$ to equal p_t. Because of noise in the series of ring draws as an indication of the true probability, the ideal observer would on average over-estimate the probability when p_t is low, and under-estimate it when p_t is high. Bayesian rationality instead would require that (under the prior implied by our data-generating process) $\mathrm{E}[p_t|p_t^*] = p_t^*$.

In fact, even if subjects do not produce fully optimal forecasts (due for example to memory limitations), if their forecasts are Bayes-rational *conditional on the information used to make the forecast,* they should satisfy the property that $\mathrm{E}[p_t|\hat{p}_t] = \hat{p}_t$ for all slider positions $\hat{p}_t$. We check for this property in the lower right panel of Figure 4, where the trials are now binned according the value of $\hat{p}_t$ and the mean value of p_t associated with each bin is indicated by the vertical coordinate of the x. This rationality condition is not grossly rejected; the diagonal falls within the IQR (if only barely) for each of the bins. However, we note a fairly clear pattern in the figure, with the mean true p above the forecast whenever $\hat{p}$ is below 0.3, and below the forecast whenever $\hat{p}$ is above 0.65. And indeed, a regression of p_t on $\hat{p}_t$ using our pooled data yields a slope coefficient significantly below 1, as shown in Table I.

The pooled data from our subjects is instead more consistent with an alternative hypothesis, namely that subjects' forecasts are distributed around the rational forecast, and equal to it on average — that is, that $\mathrm{E}[\hat{p}_t|p_t^*] = p_t^*$ for all values of p_t^*. This is related to the hypothesis of unbiasedness that Gallistel *et al.* test, but recognizes that we can at best expect subjects' forecasts to reflect the optimal Bayesian estimate of the state, and not the hidden state itself. It is also the hypothesis of "rational expectations" as originally proposed by Muth (1960), though that term has since come to be associated with the hypothesis of full Bayesian rationality, since the work of Lucas (1972).

This hypothesis is tested in the lower left panel of Figure 4. We see that when the data are binned according to the value of p_t^* (rather than by the value of p_t, as in the upper left panel), the conditional mean indicated by the x is close to the diagonal for all bins. We can nonetheless reject that the hypothesis holds exactly in our data through a regression test, as shown in Table I; but note that in this case the F statistic is less gigantic than for the other null hypotheses considered in the table.

Note that even to the extent that this last hypothesis is accepted as an approximate characterization of our data, this does not mean that our subjects' forecasts are Bayes-rational; if they (always) were, not only would the x marks be exactly on the diagonal, but the IQR would be an interval of zero length for each bin in the lower left panel of Figure 4, which is not the case. Subjects' forecasts can be more accurately characterized as equal to the rational forecast plus random noise (a characterization that also explains the slope less than 1 in the lower right panel). The extent to which forecasts are consistent with this simple hypothesis is discussed further below.

3.3 Departures from Bayesian Rationality

We next look more closely at some of the ways in which our subjects' forecasts are different from the predictions of the fully rational benchmark. Here we discuss general features of our data, without relying on any specific structural models of subject behavior. We then use these observations to motivate our subsequent consideration of a variety of specific models.

3.3.1 Stochasticity

Our subjects' deviations from the Bayesian benchmark are not simply due to the use of some boundedly rational but deterministic rule (as for example in typical "adaptive expectations" models). Instead, there is a random element in the subjects' responses to any given history of evidence. This stochasticity is at odds with traditional economic models in which optimization rules out randomization, as well as with many common models of behavioral heuristics.

Our observation above that certain statistics can be explained by a hypothesis that subjects' forecasts are equal to the optimal Bayesian forecast plus random noise has already suggested a model of behavior that incorporates random noise. However, the mere fact that the interquartile ranges shown in the lower left panel of Figure 4 are of non-zero length does not in itself prove that subjects' behavior is random; their forecasts might be deterministic functions of some aspect of the sequence of rings that they have seen that is not perfectly captured by the value of p_t^* (though highly correlated with it). More convincing evidence of stochasticity is provided by sessions in which subjects experienced precisely the same

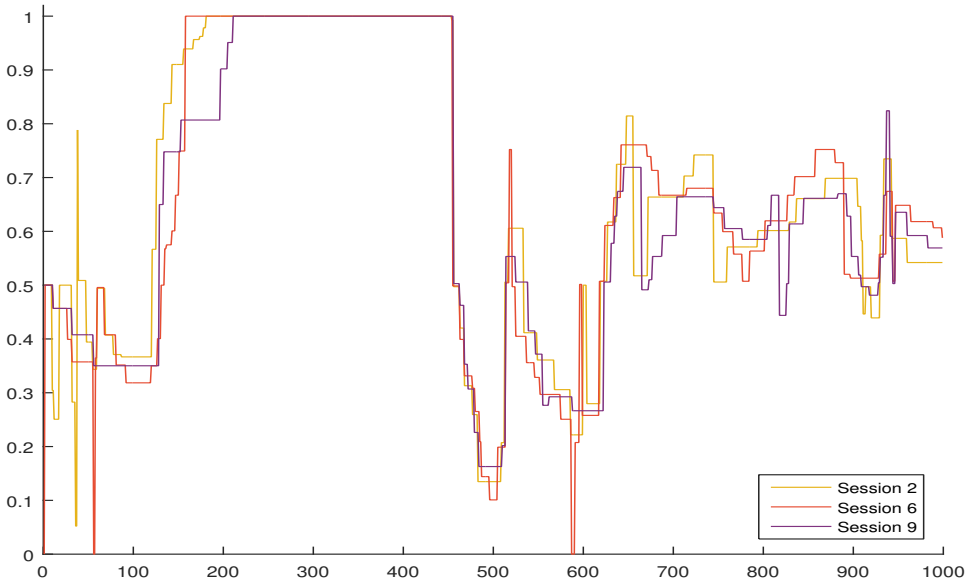


Figure 5: Experiment data, subject 1 forecasts, all repeated sessions.

sequence of ring draws as they had observed on a previous occasion.

In 20 out of our 110 sessions, we showed the subjects the same sequence of rings. These repeated sessions were never conducted on the same day for the same subject, and there is no evidence that subjects were able to recognize the repetition. We find substantial heterogeneity in different subjects' responses to the given ring sequence, but also in the responses of the same subject across the repeated sessions. Figure 5 shows the forecasts of subject 1 across the repeated sessions. The three forecast series are correlated, since they are tracking the same sequence of realized rings, but there is also considerable dispersion. The subject responds in different ways to the same history of ring realizations. Cross-sectional dispersion has also been documented in the pricing literature, using micro data for either the same retailer selling identical products in different markets, or for different customers purchasing the same good at different prices (see Kaplan and Menzio, 2015, for a recent addition to this literature). While the price setting literature has largely focused on external frictions (limits to arbitrage such as trade barriers and search costs) as the source of such dispersion, our experimental data generates this dispersion in the absence of any such frictions, which suggests that they might not be necessary in the presence of cognitive limitations affecting the behavior of pricing managers.

The top panel of Figure 6 shows the cross-sectional dispersion across all 20 repeated

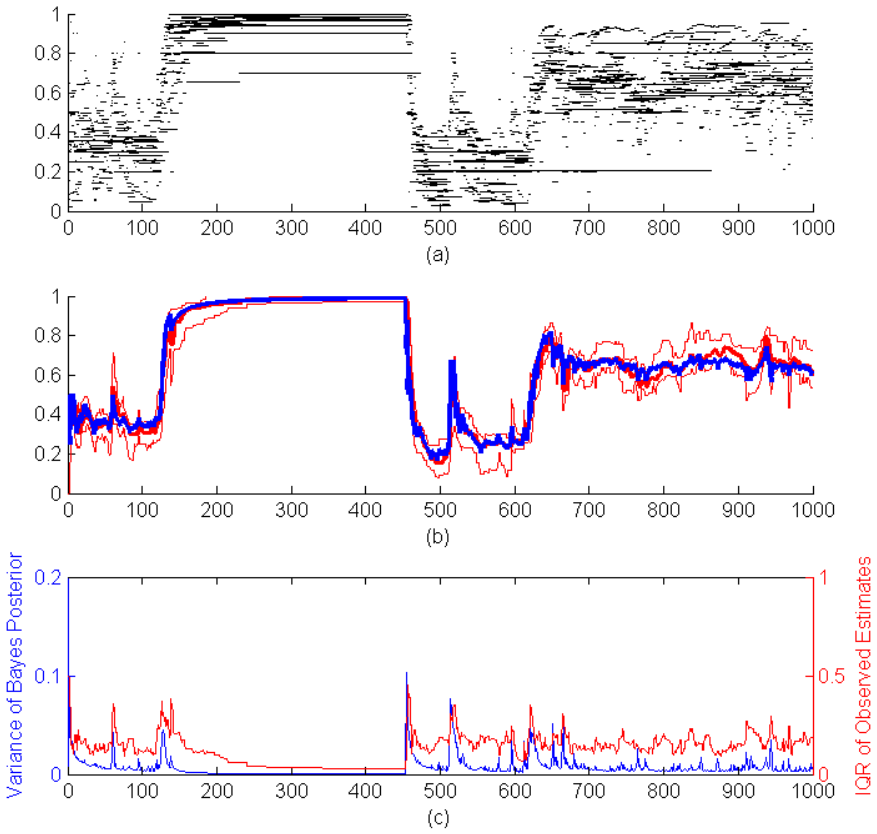


Figure 6: Experiment data, all subjects, all repeated sessions. (a) Individual forecast series, (b) Bayesian posterior mean in blue, median subject forecast and interquartile range in red, (c) Bayesian posterior variance in blue (left axis) and subjective forecasts interquartile range in red (right axis).

sessions. Throughout the session, there is a wide range of forecasts that depart – at times substantially – from the rational expectations benchmark. Nevertheless, the aggregate behavior over the 20 repeated sessions tracks the RE prediction fairly well, as shown in the middle panel of the figure, which plots the Bayesian forecast against the median slider position. This evidence, while at odds with the classic definition of rational expectations (Lucas (1972)), does provide some support for Muth's (1960) hypothesis that while individual behavior may deviate even substantially from rational expectations, it averages out to the truth. It is important to emphasize, however, that even though the aggregate appears to track the RE forecast, this finding does to mean that the RE forecast suffices for macroeconomic analysis. One reason is that the heterogeneity of responses implies greater population dispersion in beliefs, and therefore actions, than would be predicted by the RE model. In

turn, this dispersion often matters for welfare, as well as for statistics like the volume of credit or trade.

We also uncover evidence for endogenous dispersion. The heterogeneity in forecasts does not simply reflect exogenous noise in decisions, independent of the state. Instead, the dispersion of opinions is higher when objective uncertainty is higher, as measured by the Bayesian posterior variance. We show this relationship in the bottom panel of the figure: when the Bayesian posterior variance spikes, the cross-sectional dispersion also rises. This evidence suggests that variation in cross-sectional dispersion may be a reasonable proxy for uncertainty shocks, as it is sometimes used in the literature on uncertainty.

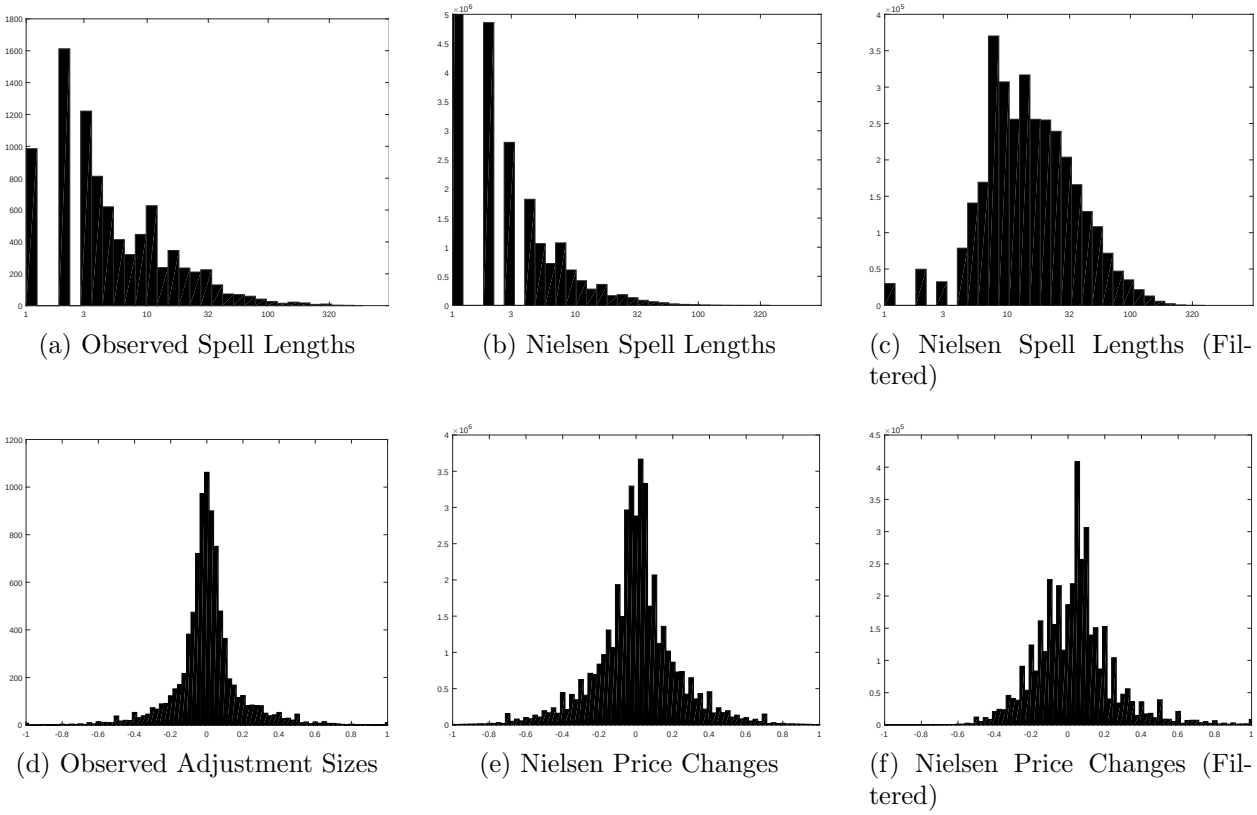
3.3.2 Discrete Adjustment

In our experiment, subjects should adjust their forecast after each additional ring draw, since each draw provides additional information about the hidden probability that subjects seek to track. Instead, we observe that subjects adjust their estimate of the hidden probability infrequently, even though there are no explicit fixed costs of adjustment. This discreteness is common to all of our subjects, albeit to varying degrees.

As shown in Figure 7, subjects exhibit a wide range of wait times between adjustments, and there is no evident pattern of time-dependence. There is also considerable variation in the size of the slider adjustments conditional on adjustment, as shown by the distribution of adjustment sizes in the lower left panel of the figure. Both of these features of our data are quite different from the predictions of the Bayesian benchmark which predicts that the slider should be adjusted on every trial, and that the adjustments should seldom be very large.

At the same time, the distribution of adjustment sizes is also inconsistent with the distribution that would be predicted by a simple "*Ss* model" of optimal adjustment. We observe adjustments of both small and large sizes, instead of a pattern in which no adjustments occur until some threshold is crossed. Such a model would predict no small adjustments and also no adjustments larger than a certain magnitude, since one would never drift that far away from the optimal forecast.

Instead, the distribution of step sizes and that of spell lengths resemble the distributions of



(a) Observed Spell Lengths (b) Nielsen Spell Lengths (c) Nielsen Spell Lengths (Filtered)

(d) Observed Adjustment Sizes (e) Nielsen Price Changes (f) Nielsen Price Changes (Filtered)

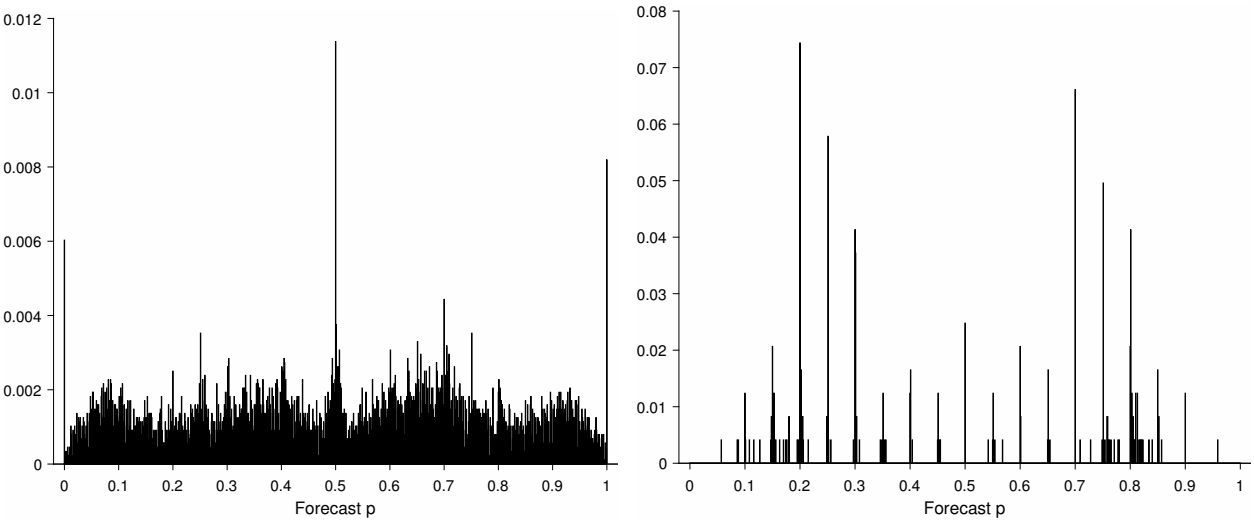
Figure 7: Discreteness. (a) Distribution of the number of ring draws between adjustments (all subjects and all sessions). (b) Distribution of the number of weeks between price changes (AC Nielsen Retail Scanner data; y-axis truncated). (c) The same distribution, but using filtered prices. (d) Distribution of the size of changes in the slider position, counting only trials on which the slider is moved (all subjects and all sessions). (e) Distribution of the size of price changes, conditional on adjustment (AC Nielsen Retail Scanner data). (f) The same distribution, but using filtered prices.

product-level price changes in micro price data. For comparison, the right panels of Figure 7 show the distribution of the size of price changes and of the periods between price changes for products in the AC Nielsen Retail Scanner data. Both our experimental data and the price data generate a leptokurtic distribution. However, in our experiment, this feature arises spontaneously, as a result of the subjects' behavior, rather than for the reasons proposed in the pricing literature, such as leptokurtic shocks or opportunities to change prices that depend on developments in other markets.

3.3.3 Round-Number Bias

We document a bias towards "round numbers" when subjects choose where to position the slider. As shown in Figure 8, subjects are more likely to estimate the hidden proportion

of green rings using numbers that are multiples of 0.05. This bias is common to all of our subjects, to varying degrees. Subject 2 is a particularly stark example of this behavior, as illustrated in the right panel of the figure.



(a) Full sample (b) Subject 2

Figure 8: Experiment data. Distribution of the subjects' slider position choices for (a) the full sample and (b) subject 2.

This discreteness in the chosen slider positions is another significant departure from the RE benchmark, which would predict a flat distribution. It occurs even though there is no discreteness in the distribution from which the hidden probability is drawn (which is continuous uniform between 0 and 1), or in the set of available actions (the mouse moves continuously[7]). There is also no sense in which the experiment makes it either easier or more profitable to choose certain slider positions over others. Nevertheless, the subjects appear to exert considerable effort to set the slider at or near these preferred numbers. We conclude that these round decimal numbers offer some cognitive advantage that compensates for the effort needed to select them.

This hypothesis is reinforced by the patterns in Figure 9, which shows the two repeated sessions for subject 2: during the first session, the subject repeatedly moves the slider to 0.2 and 0.8, while during the second session, the subject largely alternates between 0.3 and 0.7, even though seeing the same sequence of ring realizations. This pattern suggests that there

[7]The step size of the mouse is 10^{-7}, hence much finer than even the estimate that is displayed on the subjects' screen (which is on the order of 10^{-4})

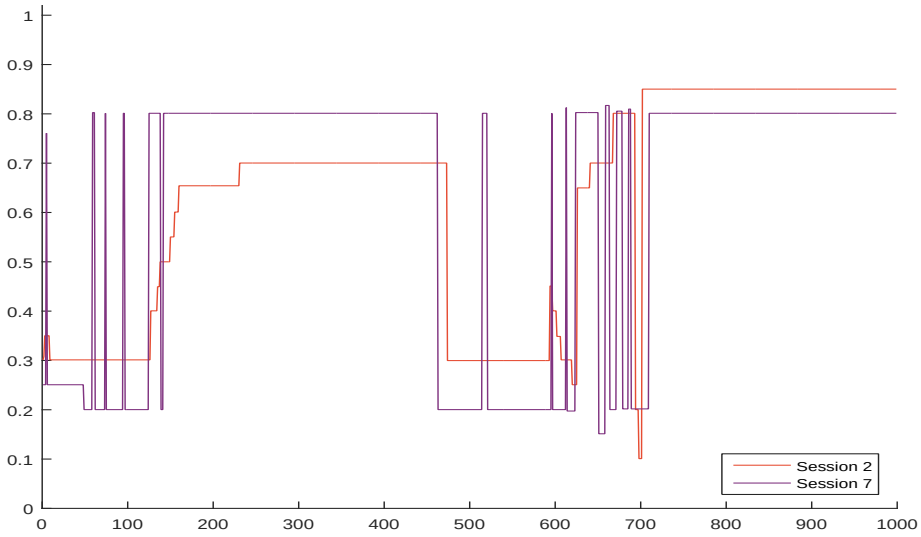


Figure 9: Experiment data, subject 2 forecasts for repeated sessions.

is some cognitive advantage to picking a small number of slider values to consider.

The evidence for round number bias mirrors the patterns found in product-level retail price data, where we see prices alternating among a small set of values for fairly long periods of time (e.g., Stevens, 2015), and where we see 9 as the most frequently used price-ending for the penny, dime, dollar and the ten-dollar digits, and multiples of dimes, dollars and ten-dollars as the most common price changes (Levy et al, 2011).

4 Modeling Discrete Adjustment

We study the degree to which our experimental data conform to the predictions of a variety of different possible models of discrete adjustment of the decision variable, with a particular focus on types of models that have been popular in the macroeconomic literature on price adjustment. Given the evidence that subjects do not adjust continuously, it is useful to decompose a model of discrete adjustment into two parts: first, a criterion that determines *when* the slider will be adjusted, and second, a criterion that determines *where* the slider position will be set, conditional on adjustment. There are a variety of models that have been proposed for each of these aspects of the adjustment dynamics; and to some extent, it is possible to combine a given model of the decision whether to adjust with any of the candidate models of the decision where to set the slider if one adjusts. We propose to test the empirical validity of alternative models of one aspect of the dynamics without

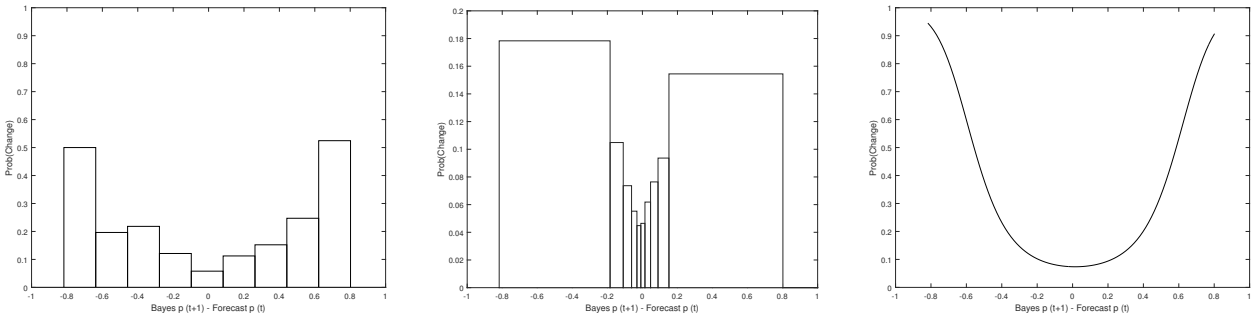
presupposing a particular answer to the other aspect.

For the adjustment decision, we consider three classes of models: (1) "gap-based" models, in which the decision whether to adjust the position of the slider depends on the distance between the existing slider position and the current optimal estimate of the underlying state (using all available information), as in the "generalized *Ss* model" proposed by Caballero and Engel (1993, 2007); (2) optimizing models with fixed costs of adjustment, in which the decision whether to adjust depends on a comparison between the value of (or expected payoff from) keeping the slider at its current position and the expected value of the position to which one would move if one were to adjust the slider, as in the price-adjustment models of Sheshinski and Weiss (1977, 1983), Dotsey, King and Wolman (1999), or Dotsey and King (2005); and (3) models in which the degree of sensitivity of the adjustment decision to the value of adjusting is limited by a cost of attention, as in the price-adjustment models of Woodford (2009) and Stevens (2015).

We also consider three possible models of the slider-position decision, conditional on adjustment: (1) models in which the slider is moved to the current optimal estimate of the underlying state (so that the "gap" is eliminated); (2) models in which the slider is moved to the position that maximizes the expected payoff, given the currently available information about the underlying state, as is usually assumed in optimizing models with fixed costs of adjustment; and (3) models in which the slider is set so as to maximize the expected payoff net of attention costs, as in the "rational inattention" theory of Sims (2003, 2010). We now take up each of these possibilities in turn.

4.1 "Gap-Based" Models of the Adjustment Decision

As explained above, under the fully rational (Bayes-optimal) benchmark, the slider position $\hat{p}_t$ should be set equal to the optimal Bayesian estimate p_t^* at all times; thus no "Bayesian gap" $\Delta_t \equiv p_t^* - \hat{p}_t$ should ever be allowed to exist. Since we can precisely determine p_t^* at all times in our experiments, this is plainly not the case, as we have documented in Section 3. One might however still hypothesize that it is the "gap" between the optimal Bayesian estimate and the current slider position (and nothing else) that should determine whether there is sufficient reason to adjust the slider. In fact, this would not generally represent optimal



(a) Equal-width bins (b) Equal-weight bins (c) Polynomial

Figure 10: Hazard functions for the slider adjustment probability as a function of the Bayesian gap. (a) Nonparametric empirical hazard plotting the fraction of trials on which adjustment occurs, for each range of values of the gap, with equal-width bins. (b) Similar plot, but with boundaries chosen so that each bin contains an equal number of observations. (c) The best-fitting logistic-polynomial hazard function. (Coefficients given in Table II.)

behavior subject to some fixed cost of adjustment, as we explain in the next section. But it is a simple rule of thumb that corresponds to a myopic form of rational choice. If a decision-maker cares *only* about the expected payoff on the current trial when deciding whether to adjust and where to adjust, and must pay a fixed cost to adjust, then it would be optimal to plan to set $\hat{p}_t$ equal to p_t^* conditional on adjusting, and to adjust if and only if the expected payoff reduction from non-adjustment exceeds the adjustment cost. The payoff reduction is proportional to the square of the Bayesian gap,

$$\mathrm{E}_t[(p_t - \hat{p}_{t-1})^2] \; - \; \mathrm{E}_t[(p_t - p_t^*)^2] \; = \; (p_t^* - \hat{p}_{-1})^2.$$

Thus myopic rational choice with a fixed cost would imply adjustment of the slider if and only if the Bayesian gap has a large enough absolute value.

More generally, we might hypothesize that adjustment should occur if and only if the gap is either above some upper threshold or below some lower threshold, as in an "*Ss* model." Even more generally, we might suppose, as in the "generalized *Ss* models" proposed by Caballero and Engel (1993, 2007), that there are no sharp thresholds, but that instead the probability of adjustment of the slider in any period is given by some continuous function of the Bayesian gap $\Lambda(\Delta_t)$, non-increasing for $\Delta < 0$ and non-decreasing for $\Delta > 0$. In the case that this "hazard function" $\Lambda(\Delta)$ is symmetric (i.e., $\Lambda(-\Delta) = \Lambda(\Delta)$ for all Δ), a

generalized *Ss* model could also be interpreted as resulting from myopic optimization, under the assumption of a cost of adjustment that is a random draw from a continuous distribution, as in the random-menu-cost models proposed by Caballero and Engel (1999) and Dotsey, King and Wolman (1999). Here, however, we also consider the fit of hazard functions that are not necessarily symmetric, as do Caballero and Engel (1993).

Figure 10 shows the fraction of trials on which the slider is adjusted, if the data are sorted into bins according to the value of the Bayesian gap at the time of the decision whether to adjust the position. In panel (a), the bins are chosen to be of equal width. This leads to a possibly misleading sense of the degree to which the hazard function is well-estimated in the case of more extreme Bayesian gaps, since the more extreme bins contain many fewer observations than do the central bins; panel (b) instead sorts the data into 10 bins each containing the same number of observations, with the boundaries of each bar indicating the boundaries of the corresponding bin. Both versions of the figure are consistent with the existence of a non-uniform hazard function, which is furthermore decreasing for $\Delta < 0$ and increasing for $\Delta > 0$, as assumed by Caballero and Engel. We note, however, that our data do not appear to support the further assumption, often made in the empirical work of Caballero and Engel, that the hazard falls to zero at least when $\Delta = 0$, if not over some larger "zone of inaction."

These plots suggest that there is at least some degree of "state-dependence" of the decision to adjust the slider. We can quantify the degree of state-dependence by computing how much the likelihood of our data is increased by allowing the adjustment hazard to depend on the Bayesian gap. If we assume that the adjustment decision α_t (taking the value 1 if the slider is adjusted, and zero otherwise) is an independent draw on each trial from a Bernoulli distribution with probability $\Lambda(\Delta_t)$ of adjustment, then the log likelihood of the data sequence $\{\alpha_t\}$ is equal to

$$LL \;=\; \sum_t \left\{\alpha_t \log \Lambda(\Delta_t) \;+\; (1-\alpha_t)\log(1-\Lambda(\Delta_t))\right\}.$$

We should not, however, be interested in the question of how large LL can be made through choice of an arbitrarily complex function $\Lambda(\Delta)$, as this is likely (with sufficiently

finely differentiated measures of Δ_t on the separate trials) to allow severe over-fitting of the particular adjustment sequence that happens to have been observed. In order to avoid such over-fitting, we need to penalize unduly complex hypotheses about the way that the adjustment hazard can depend on the gap.

We consider parametric families of possible hazard functions $\Lambda_k(\Delta;\theta)$, where θ is a vector of k parameters, and we choose both the family of models (the value of k) and a specific model within that family (the value of the parameter vector θ) so as to minimize the Bayes Information Criterion (BIC)[8]

$$BIC \equiv -2LL + k \log N, \tag{2}$$

where N is the number of observations.

In the absence of a more specific theory of the hazard function, we wish to consider a nested sequence of progressively more complex families of possible hazard functions that will allow us to approximate any continuous function arbitrarily well, if k is large enough. Here we suppose that the log odds of adjustment

$$\lambda(\Delta) \equiv \log \frac{\Lambda(\Delta)}{1-\Lambda(\Delta)} \tag{3}$$

is a polynomial function of Δ,

$$\lambda_k(\Delta;\theta) = \sum_{j=0}^{k-1} \theta_j \Delta^j, \tag{4}$$

with $k-1 \geq 0$ the order of the polynomial.[9]

[8]See, for example, Burnham and Anderson (2002, chap. 6) for discussion of this approach to model selection. Hansen and Yu (2001; section 2) propose an equivalent criterion for model selection (based on a statistic equal to the BIC divided by two) as an example of a minimum-description-length criterion. Half of the first term in (2) is a measure of the length of the codeword required to encode the data if an efficient code is used for data generated by the model in family k indexed by parameter vector θ, while half of the second term is a measure of the length of the codeword required to identify the particular model (and hence the code) used to encode the data. Here we emphasize the more familiar Bayesian interpretation of this model selection criterion, which allows us to make judgments about the relative likelihood of alternative models as explanations of our data.

[9]The choice of this particular nested sequence of parametric families is motivated by the fact that the more specific theory based on attention costs that we consider below corresponds to a particular family of models of this kind, though with a different measure of the relevant "gap." A further appealing feature of the assumption that $\lambda(\Delta)$ is a polynomial function — and hence that $\Lambda(\Delta)$ is a logistic function of a polynomial function of Δ — rather than simply assuming that $\Lambda(\Delta)$ is a polynomial function, is that the

TABLE II: Best fitting models

Gap Models							
Model	k	θ_0	θ_1	θ_2			BIC
Constant hazard	1	-2.33	–	–			59,325
Polynomial hazard	3	-2.53	-0.23	7.76			58,006.5
Symmetric poly.	2	-2.53	–	7.82			58,007.3
Model	k	ϵ	$\underline{\Delta}$	$\bar{\Delta}$			BIC
Ss with errors	3	0.089	-0.641	0.771			59,329
Symmetric *Ss*	2	0.089	-0.923	0.923			59,336
Value Models							
Model	k	θ_0	θ_1	ϵ_1	ϵ_2	$\bar{\Delta}$	BIC
Constant hazard	1	-2.33	–	–	–	–	59,325
Polynomial hazard	2	-2.49	0.95	–	–	–	58,169
Ss with errors	2	–	–	0.089	0.911	4.03	59,285
Step function	3	–	–	0.083	0.207	0.60	58,663
Inattention and Control							
Model	k	θ_0	θ_1	ψ_1	$\tilde{\Lambda}$		BIC
Rational inattention	2	-2.052	0.816	1.225	0.1139		57,598
No intrinsic costs	1	-2.326	0.299	3.335	0.0889		57,929
Control costs	2	-2.055	0.805	1.242	0.1141		57,672

Top panel: Polynomial coefficients and BIC for the best-fitting models within the Bayesian gap class of models. Middle panel: Coefficients and BIC for the best fitting models that allow adjustment to depend on the expected value of adjustment, assuming optimization subject to a random fixed cost of adjustment. Bottom panel: Models assuming that adjustment economizes on information costs, as in the "rational inattention" model with and without intrinsic costs, and assuming that adjustment is subject to control costs.

In the case of our data (pooling the data from the 99 unique sessions[10]), this statistic is minimized when $\lambda(\Delta)$ is a quadratic function, with coefficients $\{\theta_j\}$ given in the top panel of Table II. The implied hazard function $\Lambda(\Delta)$ is shown in the last panel of Figure 10. The best-fitting hazard function reaches its minimum near $\Delta = 0$, and approaches 1 for large enough positive or negative values of the gap; but even at its minimum, the probability of adjustment remains positive.

former assumption, unlike the latter, implies a hazard between 0 and 1 for all values of the gap, without any need to truncate the function implied by a given parameter vector. Moreover, this specification implies that the hazard will be greater than zero and less than 1 for all values of Δ, so that any sequence of observations will have a finite log likelihood, regardless of the parameter vector.

[10]The main text presents results for the unique sesssions only, while the Appendix presents results for all our data, including the repeated sessions.

The degree to which the data support these conclusions can be determined by comparing the minimum values of the BIC statistic that can be obtained under more restrictive families of models. Specifically, for two model families $\mathcal{M}_1$ and $\mathcal{M}_1$, the relative posterior probability that the true model belongs to one of these families rather than the other, after observing the data, is given by

$$\log \frac{P^{post}(\mathcal{M}_1)}{P^{post}(\mathcal{M}_2)} \;=\; \log \frac{P^{prior}(\mathcal{M}_1)}{P^{prior}(\mathcal{M}_2)} \;-\frac{1}{2}\left[BIC(\mathcal{M}_1) - BIC(\mathcal{M}_2)\right].$$

For example, if we assume a constant hazard function (as in the Calvo (1983) model of price adjustment), which amounts to considering only the $k = 1$ family of polynomials, then as indicated in Table II, the minimum BIC statistic is higher by 1,318 log points.[11] Thus the data increases the relative posterior probability that the quadratic log odds model is correct by a factor greater than 10^{286}, relative to the hypothesis that there is actually no dependence of the adjustment decision on the Bayesian gap.

As another example, we can consider the degree to which the data support the conclusion that the hazard function is asymmetric, by considering the more restrictive family of only symmetric polynomials — that is, only the families

$$\begin{aligned}
\lambda_1(\Delta;\theta) &= \theta_0 \\
\lambda_2(\Delta;\theta) &= \theta_0 + \theta_2\Delta^2 \\
\lambda_3(\Delta;\theta) &= \theta_0 + \theta_2\Delta^2 + \theta_4\Delta^4
\end{aligned}$$

and so on. In this case we again find that the best-fitting polynomial is quadratic (but now constrained to have a zero linear term). The fit (as measured by the likelihood of the data) is necessarily not quite as good, but the alternative model also has one fewer free parameter (as θ_1 is assumed to be zero *a priori*). In fact, as indicated in Table II, the BIC statistic is only 0.7 log points higher than in the unconstrained case, indicating very weak evidence in favor of an asymmetric hazard function.

While we are able to reject the hypothesis of no dependence on the Bayesian gap (the

[11] In all of the cases shown in Table II, the log likelihood and the BIC statistic are computed using natural logarithms.

Calvo model), our data are also inconsistent with the strong degree of state-dependence assumed in a simple *Ss* model, in which adjustment never occurs for gaps within some fixed thresholds, and occurs with certainty for gaps outside them. Any such model will have an unboundedly large BIC statistic, as it is impossible to choose the thresholds such that there are not some trials with adjustment inside the thresholds and/or some trials with non-adjustment outside the thresholds, resulting in a likelihood of zero.

This exact version of an *Ss* model may seem a straw man, so we also consider a generalization in which it is assumed that there will inevitably be some random error in the execution of any decision with regard to adjustment: that even when a decision to adjust is made, the slider will only be adjusted with probability $1-\epsilon$, and when a decision not to adjust is made, the slider will nonetheless be adjusted with probability ϵ, for some $0 < \epsilon < 1/2$. We can then consider an "*Ss* with errors" model, according to which the adjustment hazard is equal to ϵ for all $\underline{\Delta} < \Delta < \bar{\Delta}$, and equal to $1-\epsilon$ for all Δ outside this interval. This is a model with 3 free parameters, such that the log likelihood of the data is finite for all possible values of the parameters.

But even allowing for errors of this kind, the minimum value of the BIC statistic remains 1,323 log points higher under a model with sharp thresholds than under the best-fitting polynomial hazard function (the quadratic log odds model). This means that our data increase the posterior probability of the quadratic log odds model relative to that of the "*Ss* with errors" model by a factor of more than 10^{287}. If we assume that the thresholds must be symmetric (i.e., we impose the restriction that $\underline{\Delta} = -\bar{\Delta}$), reducing the number of free parameters by one, the BIC statistic further increases by 7 log points. Hence, we conclude that the adjustment decision is state-dependent — with the Bayesian gap providing at least an imperfect proxy for the relevant state — but that the Caballero-Engel hypothesis of a continuous hazard function fits the data better than a model with sharp thresholds.

4.2 Optimizing Models with a Fixed Cost of Adjustment

We next consider the extent to which our data are consistent with models in which the decision whether to adjust the slider is optimal (expected-payoff maximizing) on each trial, subject to a fixed cost of adjustment (so that it need not be optimal to move the slider every

period, as Bayesian optimality would otherwise require). Note that while there is no cost in terms of points deducted from the subject's score when the slider is moved, one might suppose that there is nonetheless a utility cost of having to move the mouse to adjust the slider position, a utility cost of the additional delay before proceeding to the next trial,[12] or a cognitive cost of having to decide where to move the slider.

Consider first the joint hypothesis that there is a constant fixed cost of moving the slider at all times, and that the slider will be adjusted to its optimal location in the case that the fixed cost is born, as in classic "menu cost" models of price adjustment such as the model of Sheshinski and Weiss (1977, 1983). Then the slider should be adjusted on trial t if and only if

$$V_t(\hat{p}_{t-1}) \;<\; \max_p V_t(p) - \kappa, \tag{5}$$

where for any possible slider position $0 \leq p \leq 1$, $V_t(p)$ is the expected cumulative payoff from all remaining trials net of adjustment costs on trials after t, if the slider is set at p at the end of trial t, and $\kappa > 0$ is the fixed cost of adjustment. Note that the value function $V_t(p)$ will depend not only on the number of the trial t, but also on the (optimal Bayesian) posterior π_t at that time (i.e., the posterior taking into account the rings drawn on trial t and earlier); the argument π_t has been suppressed to simplify the notation. Note that although the pattern of behavior documented in the previous section rules out models with a fixed cost of adjustment, we can still determine how well such a model can predict the decision of whether to adjust the slider.

We can alternatively assume, as in the generalized model of state-dependent pricing proposed by Dotsey, King and Wolman (1999), that the fixed cost of adjustment is a random variable, drawn independently on each trial from some distribution with CDF $G(\kappa)$. If we again assume that the slider will be moved to its optimal location conditional on adjustment, then the probability of adjustment on trial t should be given by a non-decreasing hazard function

$$\Lambda(\Delta_t) \;=\; G(\Delta_t), \tag{6}$$

[12] We note that there is a time cost of slider adjustment for our subjects: on average, subjects take only 1.4 seconds to request the next ring on trials where they do not move the slider, but 4.4 seconds on trials where they move the slider before requesting the next ring.

where the "value gap" Δ_t is defined as

$$\Delta_t \equiv \max_p V_t(p) - V_t(\hat{p}_{t-1}). \tag{7}$$

If the distribution of fixed costs is continuous (contains no atoms), the hazard function $\Lambda(\Delta)$ must be a continuous, non-decreasing function with $\Lambda(0) = 0$. It is possible, however, to have $\Lambda(0) > 0$ if we suppose that there is a positive probability of a zero fixed cost, as in the model of Calvo (1983).[13]

The empirical content of either version of the model depends on evaluating the value function $V_t(p)$. In any non-terminal period t, this should satisfy a Bellman recursion

$$V_t(p) = p_t^* \left[r(p;1) + W_{t+1}(p;\hat{\pi}(\pi_t,1))\right] + (1-p_t^*)) \left[r(p;0) + W_{t+1}(p;\hat{\pi}(\pi_t,0))\right], \tag{8}$$

where $r(p;s)$ is the reward function for a single trial (given by (1) above) if the slider position is p and ring drawn is s (equal to 1 for a green ring and 0 for a red ring);

$$\pi_{t+1} = \hat{\pi}(\pi_t, s)$$

is the updated posterior for trial $t+1$ implied by Bayes' Rule if the ring drawn on trial $t+1$ is s; and

$$W_{t+1}(p;\pi_{t+1}) = \Lambda(\Delta_{t+1}) \max_{p'} V_{t+1}(p') + (1 - \Lambda(\Delta_{t+1}) V_{t+1}(p) \tag{9}$$

is the expected continuation value, where Δ_{t+1} is again the value gap defined by (7).[14]

This definition makes the value function $V_t(p)$, and the hence the predictive content of conditions (5) or (6) depend on what is assumed about adjustment of the slider on later trials. But we would like to compare alternative models of the adjustment decision on a single trial using a measure of the value gap that is (to the extent possible) independent of

[13] The model of Calvo (1983) also assumes that the cost of adjustment is effectively infinite, except at those times when a zero adjustment cost is drawn; but one can combine the assumption of a positive probability of drawing a zero adjustment cost with an assumption that the adjustment cost is finite at all other times, so that the hazard approaches 1 for all large enough values of the "value gap," contrary to what is assumed in the Calvo model.

[14] This expression assumes that $t+1$ is also a non-terminal trial. If instead $t+1 = T$, the final trial of the session, we define $W_T(p;\pi_T) = 0$, regardless of the slider position and posterior.

assumptions regarding behavior in later trials.

Our method involves estimating an *atheoretical* statistical model of subjects' behavior after any trial t conditional on the slider setting $\hat{p}_t$ chosen on that trial and the posterior π_t at that point, and using this atheoretical model to compute an implied value function $V_t(p)$. We then use this *same* value function to evaluate the empirical fit of *all* of the different optimizing models of the adjustment decision that we wish to compare in this section. Note that if we are able to compute a value function $V_t(p)$ that reflects subjects' actual behavior, then verifying that the optimality condition (5) or (6) holds on each individual trial will constitute a demonstration that all of the predictions of the optimizing model are satisfied, even if a theory of optimizing behavior was not used in computing $V_t(p)$.

The continuation value can be broken into two parts,

$$V_t(p) \;=\; R_t(p) \;-\; K_t(p), \tag{10}$$

where $R_t(p)$ denotes the expected value of the remaining rewards $\sum_{\tau=t}^{T} r(\hat{p}_\tau; s_\tau)$ if the slider position $\hat{p}_t = p$; and $K_t(p)$ denotes the expected value of the cumulative adjustment costs on trials $t+1 \leq \tau \leq T$ on which the slider is adjusted. Using (1) we can further write

$$R_t(p) \;=\; -\sum_{\tau=t}^{T} \mathrm{E}[(\hat{p}_\tau - p^*_\tau)^2 \,|\hat{p}_t = p, \pi_t], \tag{11}$$

neglecting an additive term that is independent of the value of p, and thus irrelevant to our calculation of the continuation value $V_t(p)$.[15] We can then use a purely empirical model of the joint dynamics of the slider position and of the Bayesian posterior mean to compute a numerical estimate of the function $R_t(p)$ for any hypothetical slider position p that may be chosen on trial t, given the posterior π_t at that time.

We can similarly write

$$K_t(p) \;=\; \sum_{\tau=t+1}^{T} \mathrm{E}[\Gamma(\Lambda_\tau)|\hat{p}_t = p, \pi_t], \tag{12}$$

[15]Note that it is only the *relative* continuation values associated with different possible slider positions that matter for optimality conditions such as (5) or (6). Thus adding or subtracting a term that is independent of p to the definition of $V_t(p)$ makes no difference to the predictions of the model.

where Λ_τ is the probability of adjustment on trial τ (given the slider position $\hat{p}_{\tau-1}$ and the posterior π_τ on that trial), and

$$\Gamma(\Lambda) \equiv \Lambda \, \mathrm{E}[\kappa \,|G(\kappa) \leq \Lambda] \tag{13}$$

multiplies the probability of adjustment by the mean adjustment cost conditional on κ being below the threshold that is required for adjustment on a trial with that particular probability of adjustment.

Given an estimate of the function $\Gamma(\Lambda)$, we can use a purely empirical model of the adjustment hazard (more precisely, of the joint dynamics of the slider position, the Bayesian posterior, and the probability of adjustment) to obtain a numerical estimate of the function $K_t(p)$ for each trial. However, the function $\Gamma(\Lambda)$ depends on the distribution of fixed costs $G(\cdot)$, which in turn can only be inferred using a measure of the continuation value we are trying to estimate. Hence we must jointly estimate the parameters of the distribution of fixed costs needed to evaluate the function $\Gamma(\Lambda)$ *and* determine the best-fitting model of the adjustment decision, using the measure of the value gap Δ_t that is implied by this parameterization of $\Gamma(\Lambda)$.

In the numerical results presented below, we examine how well the adjustment decision can be explained by variation in the value gap, when we use a particular measure of the value based on the function $\Gamma(\Lambda)$ implied by our best-fitting value-gap-based model.[16] The contour lines associated with equal magnitudes for the value gap are plotted in Figure 11. We see that the value gap is related to the Bayesian gap (which corresponds, in the figure, to the distance of points from the diagonal). If the value gap were purely a monotonic function of the absolute value of the Bayesian gap, the contour lines would be the family of parallel straight lines with a slope exactly equal to one. In this case, the predictions of the optimizing model would be identical to those of the gap-based model with a symmetric hazard function. This is not quite true, since in the optimizing model, the value of adjusting also takes into account future expected adjustment costs, in addition to the current gap; nonetheless, the

[16] This is the value gap measure derived using the function $\Gamma(\Lambda)$ implied by the distribution of fixed costs $G(\kappa)$ corresponding to the best-fitting logistic-polynomial hazard function, displayed in the middle panel of Table II.

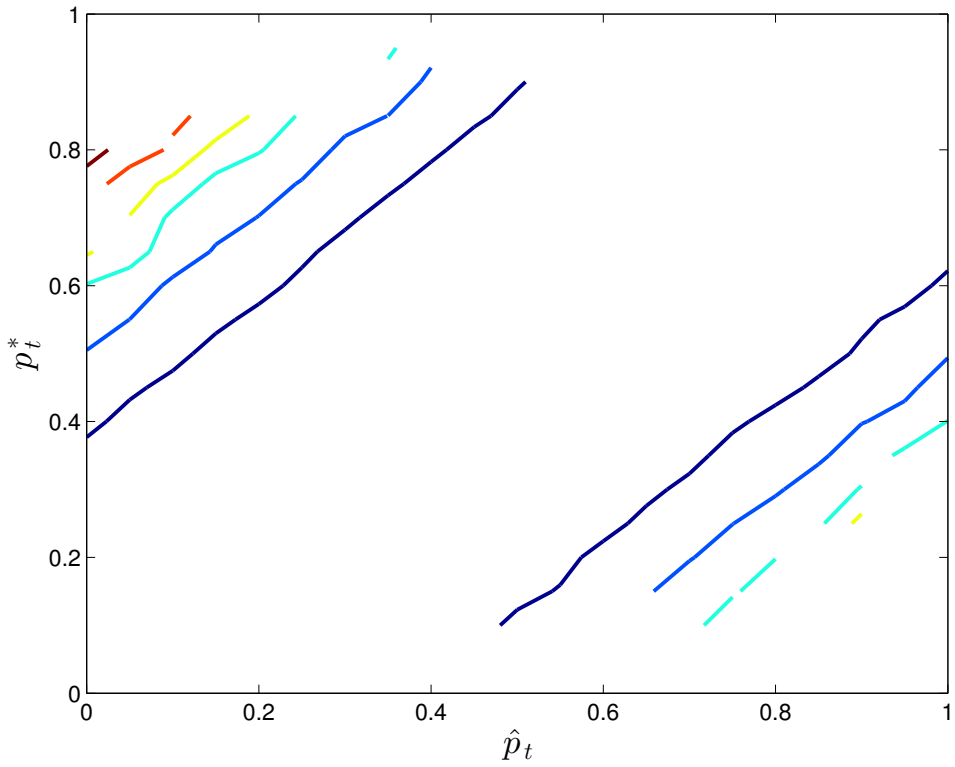


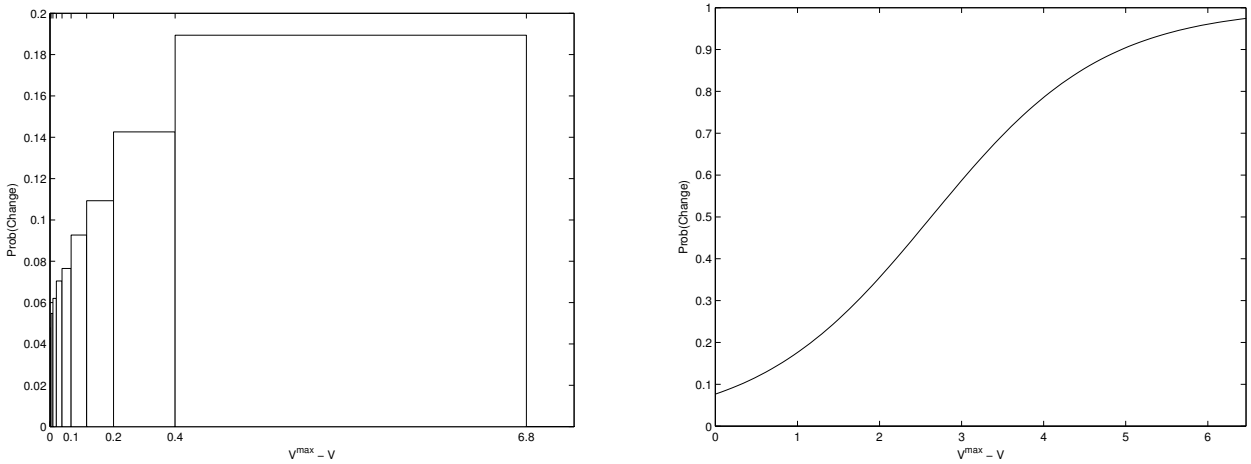
Figure 11: A contour plot of the average value gap Δ_t associated with alternative pairs of values for the slider position $\hat{p}_t$ [the horizontal axis] and the Bayesian posterior mean p_t^* [the vertical axis], when the continuation value function $V_t(p)$ is numerically approximated on the basis of an empirical model of the subsequent adjustment dynamics.

lines are close to those predicted by the Bayesian gap model.

Figure 12 plots the fraction of trials in which the slider is adjusted, grouping the trials in bins with an equal number of observations as in Figure 10(b), but now using the value gap, rather than the Bayesian gap, as the basis for classification.[17] The probability of adjustment is monotonically increasing in the value gap, as implied by (6). Moreover, adjustment only occurs with positive probability if the value gap is positive. On the other hand, there does not appear to be a sharp threshold of the kind required by equation (5), corresponding to an optimizing *Ss* model of adjustment. Nor is there any indication that the adjustment hazard falls to zero as the value gap approaches zero, as (6) would require in the absence of an atom at zero in the distribution of fixed costs.

As in the case of the models based on the Bayesian gap, we can estimate a continuous

[17]The value gap is measured, in the horizontal axis of this figure, in terms of the number of additional points (net of any increase in expected subsequent adjustment costs) that the subject can expect to obtain, on average, by adjusting the slider, assuming that it is adjusted to the optimal position.



(a) Equal-weight bins (b) Polynomial

Figure 12: Adjustment probability as a function of the value gap, computed under an assumption that if the slider is adjusted, it will be moved to the currently optimal position. (a) Nonparametric empirical hazard. As in Figure **??**(b), boundaries are chosen so that each bin contains an equal number of observations. (b) The best-fitting logistic-polynomial hazard function as a function of the value gap. (Coefficients given in Table II.)

hazard function using the BIC to penalize excessive flexibility in the family of functions considered. If we again consider the class of logistic-polynomial hazard functions defined by (3)–(4), we find that the best-fitting model makes the log odds of adjustment a linear function of the value gap, with a positive intercept and positive slope, as reported in the middle panel of Table II. The implied best-fitting value-based hazard function is shown in the second panel of Figure 12. This corresponds to a random-fixed-cost model (as in Caballero and Engel, 1999, or Dotsey, King and Wolman, 1999) in which the distribution of fixed costs has a CDF also given by the curve in that figure. The implied distribution has an atom at zero fixed cost (about 8 percent of the total probability), and a continuous distribution otherwise, with a positive probability density at zero and at all positive levels; the fixed cost can be unboundedly large, but is finite with probability one.

We again use the BIC to measure the degree to which the data favor this model. Again, we can clearly reject the hypothesis of no state-dependence: the best-fitting linear function $\lambda(\Delta)$ achieves a statistic that is lower by 1,156 log points than in the case of the best-fitting constant hazard. We also reject the hypothesis of a constant fixed cost (the optimizing *Ss* model). Even if we assume (as in our discussion above) that it is not possible to reduce the probability of adjustment below some $\epsilon_1 > 0$ or to raise it above $1 - \epsilon_1$ (with ϵ_1 a free

parameter that can be chosen to fit the data), the lowest BIC statistic obtainable with such a model remains 1,116 log points higher than that obtained with our preferred random-menu-cost model, as is also shown in the table.

Finally, we can also consider the performance of an intermediate, step-function variant of the flat hazard model, in which we suppose that for value gaps below a certain threshold, the probability of adjustment is $\epsilon_1 > 0$, while above that threshold, the probability rises to $\epsilon_2 > 0$ (though not necessarily constrained to be equal to 1-ϵ_1). Allowing for this flexibility significantly improves the fit over the flat hazard model, closing more than 40% of the difference between the BIC of the flat hazard model and that of the strictly increasing hazard model. Nevertheless, the BIC remains 413 log points higher. Thus, the data strongly supports a probability of adjustment that increases continuously with the value gap.

While our value gap measure does have some ability to predict the trials on which subjects are more likely to adjust the slider, it is a less successful predictor than the Bayesian gap; the best-fitting Bayesian gap-based model (the one with a quadratic polynomial) achieves a lower BIC, by 163 log points, than the best value gap-based model. This failure of the optimizing model to do better, even when we allow for a very flexible possible distribution of fixed costs, may indicate that in fact our subjects use a simple rule of thumb rather than a fully optimal strategy, or perhaps that they are myopic optimizers. Alternatively, it may simply indicate that our numerical approximation to the value function $V_t(p)$ is not accurate enough; we should expect our numerical estimates of the Bayesian gap to be much more accurate, since they are not limited by the use of any approximations, and perhaps this makes our measure of the Bayesian gap a better predictor even though subjects are not as myopic as would be required for the Bayesian gap to actually determine their decisions.

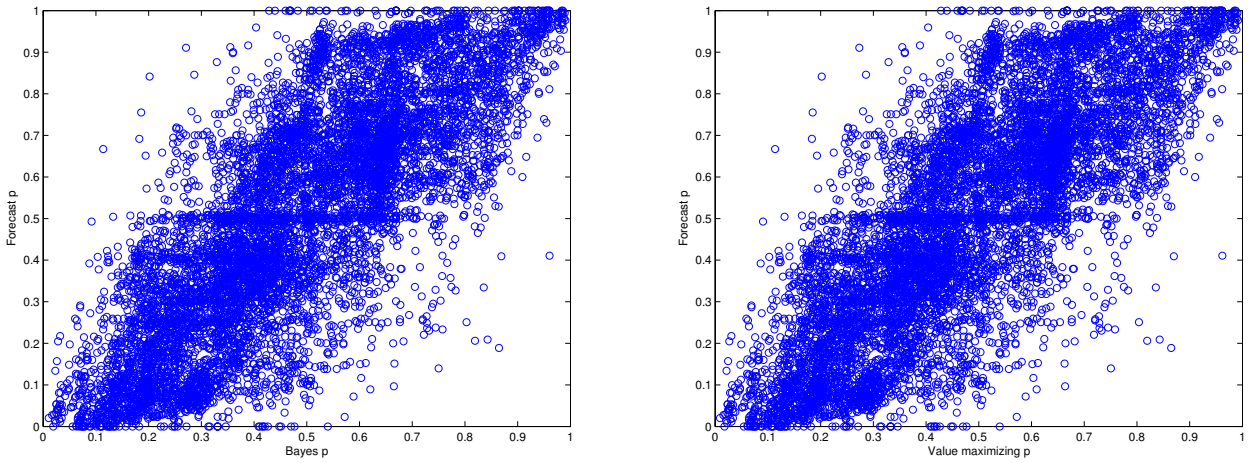
It also possible that our value gap-based model fails to describe subjects' behavior not because they do not decide optimally whether to adjust the slider, but because of the sub-optimality of their adjustment of the slider when they choose to move it. The value gap defined in (7) depends on the assumption that if the slider is adjusted, it will be set at the position that maximizes $V_t(p)$. But while this is a standard assumption in state-dependent pricing models (and many models of inattentive adjustment as well, such as the models of Calvo, 1983, or Woodford, 2009), we show below that it is not consistent with our evidence

— at least, not if our numerical approximation to the value function is at all accurate.

4.3 Models of the Slider Position Conditional on Adjustment

We also wish to consider alternative models for the decision of where to move the slider, conditional on adjustment. The simplest slider position model assumes that the slider position is set to the currently optimal Bayesian estimate whenever the slider is adjusted, as in the model of Caballero and Engel (1993). As shown in Figure 13, the relationship between the optimal Bayesian forecast and the position chosen by the subjects upon adjusting the slider is very noisy. Nonetheless, the Bayesian setting has predictive value: running a regression yields a slope that is very close to 1.

A more sophisticated class of models assume that upon adjustment, the decision variable moves to the value maximizing position conditional on the current state (e.g., Dotsey et al (1999), Woodford (2009), etc.). Panel (b) of Figure 13 shows a scatter plot of the slider positions chosen on those trials where subjects adjust the slider plotted against the position p that maximizes $V_t(p)$ on that trial. This plot also shows a large dispersion in the actual slider positions that are chosen, even conditioning on the currently optimal slider position (at least the one that would be implied by our estimate of the value function) — especially when the currently optimal position is in the middle of the interval of possible positions. This suggests a considerable degree of random error (presumably reflecting inattentiveness) when the slider is moved. Such errors must be taken into account in judging what an optimal decision should be as to whether the slider should be adjusted at all, as they should reduce the expected increase in payoff that can be achieved by adjustment. However, once we recognize that the decision as to where to move the slider is subject to random error, we should consider that the randomness in the decision to adjust the slider may reflect inattention as well, rather than random variation in the cost of adjustment. We take up explicit models of inattentive decisionmaking in both parts of our decision problem in the next section.



(a) Bayesian estimate

(b) Value maximizing estimate

Figure 13: The distribution of slider positions that are chosen, conditional on adjustment. (a) Scatter plot of the slider position [vertical axis] against the Bayesian optimal setting; (b) scatter plot of the slider position [vertical axis] against the one that would maximize the continuation value function.

5 A Model of Inattentive Adjustment

We now propose a model of discrete adjustment in which both the decision when to adjust the slider and the decision where to move it conditional on adjustment are assumed to be optimal, but in terms of an objective that includes costs of paying closer attention (to one or both decisions). The particular theory of attention costs that we consider is based on the theory of "rational inattention" proposed by Sims (2003, 2011) and applied to the discrete adjustment of firms' pricing policies by Woodford (2009) and Stevens (2015).[18] We generalize the models of inattentive adjustment used in the latter two papers, by relaxing the assumption that agents learn the true state upon adjustment, and by introducing the possibility of intrinsic preference for particular slider positions.

[18]This approach is used to model randomness in experimentally observed decisions by Cheremukhin *et al.* (2011). Their application, however, is to an experiment involving a sequence of independent static decision problems, and so does not require them to address the issues arising in the case of dynamic decision problems that we take up here.

5.1 Modeling Attention Costs

The theory of rational inattention assumes that it is not costless for a decisionmaker to vary her actions depending on the situation in which she finds herself; instead, there is a cost of executing strategies that require more precise information about the decision situation. Following Sims (2003), we suppose that the relevant measure of the degree of precision of information required by a given strategy is measured by the Shannon *mutual information* between the decisionmaker's action and the external state that defines the decision situation.[19] Moreover, also as in Sims (2003), we suppose that a subject's actions when trial t is reached can depend on any aspect of the complete history h_t of states (visible to the subject in principle) that have been realized up until that time, subject only to a cost that is proportional to the mutual information between those actions and the "state" specified by the history h_t.

In our case, the history h_t consists of the complete history of ring draws through trial $t-1$, together with the complete history of slider positions chosen through trial $t-1$. It does *not* include the history of evolution of the underlying state $\{p_t\}$, as this is not visible to the subject, no matter how closely he may pay attention.[20] In assuming that the entire history is equally available to the decisionmaker, in the sense that it is no more costly for her to condition her action upon a ring draw 100 trials earlier than to condition it upon the most recent draw, we effectively assume that the relevant cognitive constraint relates not to the precision with which events are observed when they occur, or stored in memory, but *only* to the precision with which current actions may depend on information retrieved from memory.

In this spirit, we assume that while a subject's action may depend on what he recalls about his previous actions (and not simply the ring draws that have been observed), it is costly for

[19] The mutual information is a property of the joint distribution of the state and the action that measures how much information about the state can be obtained from observing the action, and vice versa (Cover and Thomas, 2006).

[20] While the theory of rational inattention may be viewed as endogenizing the partial information that a decisionmaker possesses about her situation — and even treating it as if it is a matter of choice on her part — this does not mean that we suppose that there is any choice about the degree to which the available evidence (the sequence of ring draws) will be informative about the underlying state. The choice only relates to the degree of attention that the subject pays to the sequence of ring draws and to his own past actions.

the subject's behavioral strategy to condition on such memories, to the same extent as it is costly for it to condition on any other information contained in memory. There is assumed to be no "free" information (upon which actions may be conditioned without incurring any attention costs), not even about previous contents of one's own mind.[21]

Let a subject's behavioral rule be specified by a sequence of functions $\Lambda_t(h_t)$, specifying the probability Λ_t of adjusting the slider on trial t for each possible history h_t prior to that trial's ring draw, and a sequence of functions $\mu_t(h_t)$, specifying a probability measure μ_t over possible new locations for the slider (if it is adjusted on trial t) for each possible prior history h_t. (Both of these functions are specified for all $1 \leq t \leq T$.) The theory of *rational inattention* postulates that these functions maximize the objective

$$\mathrm{E}\left\{\sum_{t=1}^{T}\left[r(\hat{p}_t; s_t) \; - \psi_1 I_1 \; - \psi_2 \Lambda_t I_2\right]\right\} \tag{14}$$

where I_1 is a measure of the amount of information used in deciding on each trial whether to adjust the slider; I_2 is a corresponding measure of the amount of information used in deciding where to set the slider, on those trials where it is adjusted; $\psi_1, \psi_2 > 0$ are attention cost parameters for the two types of information; and $\mathrm{E}\{\cdot\}$ indicates an expectation over the possible sequences of realizations of ring draws, the possible outcomes of the random decision whether to adjust the slider each period, and the possible outcomes of the random decision as to where to set the slider when it is adjusted. The final term inside the brackets is multiplied by Λ_t because the cost $\psi_2 I_2$ is paid only in those states (i.e., realizations of the random decision whether to adjust the slider) in which the slider is adjusted, which states occur with conditional probability $\Lambda_t(h_t)$ given a history h_t. Hence, as in Alvarez et al (2015) and Stevens (2015), we distinguish between the cost of conducting reviews and the cost of changing the decision variable.

[21] In this respect our assumption about information costs differs from that of Sims (2003). Sims assumes that a decisionmaker arranges to receive a signal σ_t each period, that can depend on both the history of external states and the previous history of signals, and that the action in any period t can depend on the entire history of signals $(\sigma_t, \sigma_{t-1}, \ldots)$; the information cost each period depends only on the informativeness of that period's new signal σ_t, measured by the mutual information between the distributions of σ_t and h_t *conditional upon the prior history of signals* $\{\sigma_{t-1}, \sigma_{t-2}, \ldots\}$. Thus the information revealed by prior signals is assumed to remain available "for free" (with no attention cost) in all subsequent periods. Our alternative assumption instead follows the information-cost formulations used by Woodford (2009) and Steiner, Stewart and Matejka (2016).

There are two possible interpretations of the objective (14): under one, any amount of information can be used in each decision, subject to a positive per-unit cost of the information used, and (14) represents the subject's expected monetary payoff net of attention costs (which have been expressed in terms of monetary equivalents, even though they represent a disutility of paying more attention, rather than any monetary charge). Under the other, there is a finite upper bound on the information that can be used in either decision, and (14) is the Lagrangian for the constrained optimization problem. The first-order conditions for optimality of the subject's behavioral rule presented below are the same in either case. Under the second interpretation, there is no reason why the two cost parameters ψ_1, ψ_2 need be equal in value (assuming that the attention constraints are two independent constraints); but even under the first interpretation, they need not be equal (the sources from which additional attention can be mobilized may be different for the two types of decisions, and so have different costs), and we treat these as two independent parameters of our model.

Note also that both I_1 and I_2 are numbers that are independent of the trial t on which the decision is made, or the history prior to that trial. This is because we assume that the subject's attentional strategy (producing subjective states with certain probabilities, conditional on the history h_t) is chosen once for the entire session: it cannot be contingent upon any aspects of the subject's situation at a given point in the session, because the subject is assumed not to have any access to such facts (including the number of rings that have already been drawn), other than *through the attentional strategy* itself.[22]

If the measure of information that determines the attention costs is Shannon mutual information, we have

$$I_1 \equiv \frac{1}{T} \mathrm{E} \left\{ \sum_{t=1}^{T} D(\Lambda_t(h_t) || \bar{\Lambda}) \right\}, \tag{15}$$

where

$$\bar{\Lambda} = \frac{1}{T} \mathrm{E} \left\{ \sum_{t=1}^{T} \Lambda_t(h_t) \right\}, \tag{16}$$

[22] Again, the model of Sims (2003) is different. Sims assumes costless access to the history of prior signals (including the passage of time), so that a separate attentional strategy can be chosen for the case that one finds oneself in period t following a particular signal history $\{\sigma_{t-1}, \sigma_{t-2}, \ldots\}$. Hence there is assumed to be a separate cost for the amount of information used in each such contingency.

and

$$I_2 \equiv \frac{\mathrm{E}\left\{\sum_{t=1}^{T} \Lambda_t(h_t)\, D(\mu_t(h_t)||\bar{\mu})\right\}}{\mathrm{E}\left\{\sum_{t=1}^{T} \Lambda_t(h_t)\right\}}, \tag{17}$$

where

$$\bar{\mu} = \frac{\mathrm{E}\left\{\sum_{t=1}^{T} \Lambda_t(h_t)\, \mu_t(h_t)\right\}}{\mathrm{E}\left\{\sum_{t=1}^{T} \Lambda_t(h_t)\right\}}, \tag{18}$$

in which expressions $D(\mu||\lambda)$ denote the Kullback-Leibler divergence[23] between two probability measures μ and λ.[24] That is, the mutual information is the expected value (averaging over all of the possible states in which a given decision may be taken) of the Kullback-Leibler divergence between the conditional probability measure describing the decisionmaker's probabilistic action in that state and a reference probability measure (the same in all states) which is equal to the unconditional probability measure over the same set of possible actions.

We can alternatively write the mutual information (15) as

$$I_1 = \min_{\bar{\Lambda}} \frac{1}{T}\, \mathrm{E}\left\{\sum_{t=1}^{T} D(\Lambda_t(h_t)||\bar{\Lambda})\right\},$$

where the constraint (16) is no longer imposed, since one can show that (16) is the measure that minimizes the mean Kullback-Leibler divergence. We can similarly rewrite (17). This allows us to write the optimization problem that defines rational inattention as the choice of functions $\{\Lambda_t(h_t)\}$ and $\{\mu_t(h_t)\}$ *and* reference measures $\bar{\Lambda}$ and $\bar{\mu}$ so as to maximize

$$\mathrm{E}\left\{\sum_{t=1}^{T} \left[r(\hat{p}_t; s_t) - \psi_1 D(\Lambda_t||\bar{\Lambda}) - \psi_2 \Lambda_t D(\mu_t||\bar{\mu})\right]\right\}. \tag{19}$$

This has the advantage of being time- and state-separable, so that we can use standard recursive methods to characterize an optimal behavioral rule.

[23] The Kullback-Leibler divergence is defined as

$$D(\mu||\lambda) = \int \log \frac{d\mu}{d\lambda} d\mu.$$

[24] In a small abuse of notation, we have used the notation Λ_t to refer to the probability measure that assigns probability Λ_t to one of two possible states, and probability $1 - \Lambda_t$ to the other; and similarly, $\bar{\Lambda}$ for the probability measure that assigns probabilities $\bar{\Lambda}$ and $1 - \bar{\Lambda}$ to the same two possible states.

5.2 The Complete Rational Inattention Problem

Thus far, we have assumed that the objective is simply to maximize the expected cumulative monetary payoff, net of information costs. We can generalize the theory by assuming that there may also be intrinsic costs or convenience associated with particular actions. For example, there may be an effort cost associated with moving the mouse in order to adjust the slider position. Or there may be a preference for choosing certain slider positions, independent of the monetary payoff expected to result from choosing those positions; our subjects' apparent preference for positions corresponding to round numbers suggests that this is the case.

Let us suppose that there is a cost $c(p)$ associated with choice of slider position p (born only at times when the slider is adjusted), where $c(p)$ is some measurable function (not necessarily non-negative in value, as there might be a positive psychic reward from some position choices). Similarly, let us suppose that there is a cost c^{adj} associated with adjustment of the slider, and a cost c^{non} of non-adjustment (which quantities are also not necessarily non-negative). We can then generalize the objective of the rational inattention problem to obtain

$$\mathrm{E}\left\{\sum_{t=1}^{T}\Big[r(\hat{p}_t;s_t)\ -\psi_1 D(\Lambda_t||\bar{\Lambda})\ -\psi_2\Lambda_t D(\mu_t||\bar{\mu}) \right.$$
$$\left. -\Lambda_t c^{adj}-(1-\Lambda_t)c^{non}-\Lambda_t\int c(p)d\mu_t(p)\Big]\right\}. \tag{20}$$

The problem is then to choose a reference adjustment probability $\bar{\Lambda}$, a reference measure over slider positions $\bar{\mu}$, and sequences of functions $\Lambda_t(h_t)$ and $\mu_t(h_t)$, so as to maximize (20).

We begin by assuming that values for $\bar{\Lambda}$ and $\bar{\mu}$ are specified, and optimizing over the other functions. We then have a problem that can be treated recursively. Let us consider first the problem of choosing a probability measure μ_t over new slider locations, conditional on having decided to adjust the slider on trial t, in the case of some history h_t. The optimal measure μ_t must maximize

$$\int V_t(p)d\mu_t(p)\ -\int c(p)d\mu_t(p)\ -\psi_2 D(\mu_t||\bar{\mu}), \tag{21}$$

where $V_t(p)$ is again a continuation value if slider position p is chosen on trial t; but this is now defined as the expected cumulative payoff from all remaining trials, net of both information costs and the intrinsic costs of actions on those remaining trials.

The definition of the function $V_t(p)$, which depends only on behavior anticipated on subsequent trials, is independent of the choice of $\hat{p}_t$. Hence the solution to this maximization problem is given by

$$\mu_t(p) \;=\; \frac{\exp\{\psi_2^{-1}(V_t(p) - c(p))\}\bar{\mu}(p)}{\int \exp\{\psi_2^{-1}(V_t(p') - c(p'))\}d\bar{\mu}(p')} \tag{22}$$

for all p in the support of the reference distribution $\bar{\mu}$. (The distribution μ_t furthermore has the same support as $\bar{\mu}$.) In the limit as $\psi_2 \to 0$, the optimal μ_t approaches a measure that assigns probability one to the p that maximizes $V_t(p) - c(p)$, and this model of slider movement reduces to the one assumed in section 4.2;[25] but for any $\psi_2 > 0$, the slider location decision is instead random, as suggested by Figure **??**.

Because only the *relative* preference for different slider locations matters for predicted behavior,[26] we can choose a normalization for the cost function (adding or subtracting some constant from $c(p)$ for all p) such that

$$\int \exp\{-\psi_2^{-1}c(p)\}d\bar{\mu}(p) \;=\; 1. \tag{23}$$

We can then define an alternative reference measure

$$\tilde{\mu}(p) \equiv \exp\{-\psi_2^{-1}c(p)\}\bar{\mu}(p),$$

in terms of which the prediction (22) takes the simpler form

$$\mu_t(p) \;=\; \frac{\exp\{\psi_2^{-1}V_t(p)\}\tilde{\mu}(p)}{\int \exp\{\psi_2^{-1}V_t(p')\}d\tilde{\mu}(p')}. \tag{24}$$

This model of the location decision implies a value of adjustment V_t^{adj} (given the history

[25] In section 4.2, we assumed no intrinsic preference for particular slider positions, but the theory could easily have been extended to allow for such preferences.

[26] Note that anything that we might wish to assume about the overall cost of adjustment as opposed to non-adjustment can be achieved through the specification of $c^{adj} - c^{non}$.

h_t) equal to the maximized value of (21). The solution (22) implies that this will equal

$$V_t^{adj} \;=\; \psi_2 \,\log \int \exp\{\psi_2^{-1} V_t(p)\} d\tilde{\mu}(p). \tag{25}$$

The optimal adjustment hazard Λ_t, given the history h_t, is then the value that maximizes

$$\Lambda_t(V_t^{adj} - c^{adj}) \,+ (1-\Lambda_t)(V_t(\hat{p}_{t-1}) - c^{non}) \,- \psi_1 D(\Lambda_t||\bar{\Lambda}). \tag{26}$$

This optimal value is given by

$$\Lambda_t \;=\; \frac{\bar{\Lambda}\,\exp\{\psi_1^{-1}(V_t^{adj} - c^{adj})\}}{(1-\bar{\Lambda})\,\exp\{\psi_1^{-1}(V_t(\hat{p}_{t-1}) - c^{non})\} \,+\, \bar{\Lambda}\,\exp\{\psi_1^{-1}(V_t^{adj} - c^{adj})\}}. \tag{27}$$

Note that (27) implies that the probability of adjusting is again given by a continuous hazard function, $\Lambda_t = \Lambda(\Delta_t)$, where now the relevant "value gap" is defined by

$$\Delta_t \;=\; V_t^{adj} \;-\; V_t(\hat{p}_{t-1}) \tag{28}$$

instead of (7). An important difference between (28) and (7) is that the new definition no longer implies that the gap is necessarily non-negative; hence the hazard function must also be defined for negative gap values. The hazard function is predicted to be a continuous, monotonically increasing function of the gap, positive but finite for all values of the gap. In particular, the hazard is necessarily positive in the case of a zero value gap, unlike the prediction of a random-fixed-cost model in which the distribution of fixed costs is assumed to contain no atoms.

More specifically, (27) implies that

$$\log \frac{\Lambda_t}{1-\Lambda_t} \;=\; \lambda(\Delta_t) \;\equiv\; \log \frac{\tilde{\Lambda}}{1-\tilde{\Lambda}} \,+ \psi_1^{-1}\Delta_t, \tag{29}$$

where $\tilde{\Lambda}$ is an alternative reference probability,

$$\tilde{\Lambda} \;\equiv\; \frac{\bar{\Lambda}\exp\{-\psi_1^{-1}c^{adj}\}}{(1-\bar{\Lambda})\exp\{-\psi_1^{-1}c^{non}\} + \bar{\Lambda}\exp\{-\psi_1^{-1}c^{adj}\}}.$$

Once again, only the relative cost $c^{adj} - c^{non}$ has behavioral consequences; so we can adopt a normalization of these costs in which

$$(1 - \bar{\Lambda}) \exp\{-\psi_1^{-1} c^{non}\} + \bar{\Lambda} \exp\{-\psi_1^{-1} c^{adj}\} = 1; \tag{30}$$

in this case, we can simply write

$$\tilde{\Lambda} \equiv \bar{\Lambda} \exp\{-\psi_1^{-1} c^{adj}\}.$$

Equation (29) then implies that the hazard function must belong to a specific parametric family of functions, parameterized by the values of ψ_1 and $\tilde{\Lambda}$. The parameter ψ_1 determines the steepness of this function for low values of the gap; the value of $\tilde{\Lambda}$ instead shifts the probability of adjustment up or down for all values of the gap. We further note that the log odds of adjustment are predicted to be a linear function of Δ_t, as in the case of the best-fitting random-fixed-cost model discussed in section 4.2 (though here the definition of Δ_t is different). In the limit as $\psi_1 \to 0$, the hazard function approaches a step function, according to which the probability of adjustment is zero for all $\Delta_t < 0$, and one for all $\Delta_t > 0$.

This model of the optimal adjustment hazard implies that the optimized value of the objective (26) is equal to

$$W_t(\hat{p}_{t-1}) = \psi_1 \log \left[\tilde{\Lambda} \exp\{\psi_1^{-1} V_t^{adj}\} + (1 - \tilde{\Lambda}) \exp\{\psi_1^{-1} V_t(\hat{p}_{t-1})\} \right]. \tag{31}$$

The continuation value function $V_t(p)$ is then again given by (8), but with $W_{t+1}(p)$ defined by (31) in the case that t is a non-terminal trial, rather than by (9).[27]

Adjustment behavior is then completely described by equations (24) and (29), together with equations (8), (25) and (31) that recursively define the continuation values following any partial history, given the model parameters ψ_1, ψ_2, the reference adjustment probability $\tilde{\Lambda}$, and the reference measure $\tilde{\mu}$ over slider positions. It remains to be discussed how the reference probabilities are endogenously determined.

Recall that given the information cost parameters and the parameters specifying the

[27] Once again, in the case that $t = T$, we define $W_{T+1}(p) = 0$.

intrinsic costs of different actions, the reference probabilities $\bar{\Lambda}$ and $\bar{\mu}$ are chosen so as to maximize (20). The variational method discussed in Jung *et al.* (2015)[28] can be used to show that the first-order conditions for these choices to be optimal are that

$$J(p) \;\leq\; 1 \tag{32}$$

for each possible slider position p, with equality at all positions p in the support of $\bar{\mu}$, where we define

$$J(p) \;\equiv\; \mathrm{E}\left[\frac{\exp\{\psi_2^{-1}(V_t(p)-c(p))\}}{\int \exp\{\psi_2^{-1}(V_t(p')-c(p'))\}d\bar{\mu}(p')}\Big|\text{ adjust on trial } t\right]; \tag{33}$$

and similarly, that

$$J^{adj} \;\leq\; 1, \qquad J^{non} \;\leq\; 1,$$

with equality in the case of any action (adjustment or non-adjustment) that occurs with positive probability under the reference measure $\bar{\Lambda}$, where we define

$$J^{adj} \;\equiv\; \mathrm{E}\left[\frac{\exp\{\psi_1^{-1}(V_t^{adj}-c^{adj})\}}{(1-\bar{\Lambda})\,\exp\{\psi_1^{-1}(V_t(\hat{p}_{t-1})-c^{non})\} \;+\; \bar{\Lambda}\,\exp\{\psi_1^{-1}(V_t^{adj}-c^{adj})\}}\right],$$

$$J^{non} \;\equiv\; \mathrm{E}\left[\frac{\exp\{\psi_1^{-1}(V_t(\hat{p}_{t-1})-c^{non})\}}{(1-\bar{\Lambda})\,\exp\{\psi_1^{-1}(V_t(\hat{p}_{t-1})-c^{non})\} \;+\; \bar{\Lambda}\,\exp\{\psi_1^{-1}(V_t^{adj}-c^{adj})\}}\right].$$

In these expressions, $\mathrm{E}[\cdot]$ denotes an expected value averaging over all possible specifications of the date t and history h_t under which an adjustment decision may be made, weighting each in proportion to the ex ante probability of being in that circumstance (given the data-generating process for ring draws). In (32), we additionally condition on t being a trial on which the slider is adjusted.

[28] See also Woodford (2009) and Stevens (2015) for treatment of related problems.

5.3 Behavioral Implications of the Model

In terms of the alternative reference measure $\tilde{\mu}$ (and assuming the normalization (23) of costs), (32) can equivalently be written as

$$c(p) \;\geq\; \psi_2 \log \mathrm{E}\left[\frac{\exp\{\psi_2^{-1} V_t(p)\}}{\int \exp\{\psi_2^{-1} V_t(p')\} d\tilde{\mu}(p')} | \text{ adjust on trial } t\right]. \tag{34}$$

In fact, we see that (in the absence of any direct evidence, or a priori restriction, on the intrinsic costs $c(p)$) this places no restriction on the measure $\tilde{\mu}$. Given any measure $\tilde{\mu}$, we can choose a cost function $c(p)$ that satisfies (34) for all p, as well as the complementary slackness condition (that (34) must hold with equality, for p in the support of $\tilde{\mu}$), and any such cost function will also satisfy the normalization (23).

Similarly, in terms of the alternative reference adjustment probability $\tilde{\Lambda}$ (and assuming the normalization (30)), the conditions (5.2) can equivalently be written as

$$c^{adj} \;\geq\; \psi_1 \log \mathrm{E}\left[\frac{\exp\{\psi_1^{-1} V_t^{adj}\}}{(1-\tilde{\Lambda})\, \exp\{\psi_1^{-1} V_t(\hat{p}_{t-1})\} \;+\; \tilde{\Lambda}\, \exp\{\psi_1^{-1} V_t^{adj}\}}\right], \tag{35}$$

$$c^{non} \;\geq\; \psi_1 \log \mathrm{E}\left[\frac{\exp\{\psi_1^{-1} V_t(\hat{p}_{t-1})\}}{(1-\tilde{\Lambda})\, \exp\{\psi_1^{-1} V_t(\hat{p}_{t-1})\} \;+\; \tilde{\Lambda}\, \exp\{\psi_1^{-1} V_t^{adj}\}}\right]. \tag{36}$$

Once again, we note that these conditions place no restriction on the value of $\tilde{\Lambda}$. Given any value for $\tilde{\Lambda}$, we can find values for c^{adj} and c^{non} that satisfy (35)–(36) together with the complementary slackness conditions, and any such values will also satisfy the normalization (30).

Hence the only restrictions that the rational inattention model places on observable behavior can be summarized as follows: there must exist information-cost parameters $\psi_1, \psi_2 > 0$, a reference adjustment probability $\tilde{\Lambda}$, and a reference measure $\tilde{\mu}$ over possible slider positions, such that following any history of ring draws and previous slider adjustments h_t, the probability of adjusting the slider on trial t will be given by (29), and conditional upon adjustment, the new location of the slider will be drawn from the measure given by (24). In these formulas the value functions $V_t(p)$ and V_t^{adj} are the (uniquely determined) ones recursively defined by equations (8), (25), and (31), given the values of $\psi_1, \psi_2, \tilde{\Lambda}$, and $\tilde{\mu}$.

5.4 Interpretation in Terms of Control Costs

It is worth noting that our theory is observationally equivalent to one in which a subject's behavioral rule is chosen to maximize an objective of the form

$$\mathrm{E}\left\{\sum_{t=1}^{T}\left[r(\hat{p}_t;s_t)\ -\psi_1 D(\Lambda_t||\tilde{\Lambda})\ -\psi_2\Lambda_t D(\mu_t||\tilde{\mu})\right]\right\}, \tag{37}$$

with the reference probabilities $\tilde{\Lambda}$ and $\tilde{\mu}$ treated as (exogenously given) parameters of the cost functions, rather than as endogenous variables to be optimized, as in formulation (19) of the rational inattention model (without intrinsic preferences for particular actions). In this formulation, the model can be viewed not as one that limits the degree to which actions can be sensitive to (and thus require information about) alternative decision situations, but rather as one that limits the accuracy with which actions can be controlled.

In the alternative theory, behavioral rules are chosen in order to maximize the expected monetary payoff net of "control costs," which are proportional to the degree to which the probability of observing different possible actions on a given occasion differs from certain "default" probability distributions, which represent an "uncontrolled" pattern of behavior by the subject.[29] The objective (37) can be given an interpretation of this kind, in which the probabilities $\tilde{\Lambda}$ and $\tilde{\mu}$ represent the "default" pattern of behavior, and control costs are assumed to be proportional to the relative entropy between the behavioral distribution and the default distribution, as in the "logit model" of discrete adjustment dynamics proposed by Costain and Nakov (2014).[30]

While the entropy control-cost model and the rational inattention model are equivalent, if we consider the least restrictive form of each type of model — that is, if in the control-cost model, we are completely agnostic about what the default probabilities might be, and in the rational inattention model, we are completely agnostic about the intrinsic preferences for alternative actions — they are no longer equivalent if we adjoin more special hypotheses to these relatively flexible frameworks. For example, it is common in the literature on entropy

[29] For the general idea of modeling behavior subject to control costs, see, e.g., van Damme (1991, chap. 4).

[30] This particular specification of control costs has been proposed in static contexts by Stahl (1990) and Mattson and Weibull (2002). For applications to explaining experimental data, see Anderson *et al.* (2002).

control-cost models to assume for the default probabilities a uniform distribution over the set of possible actions, so that the control cost is simply proportional to the entropy of the behavioral probability distribution.[31] However, uniformity of the reference distribution (for example, $\tilde{\mu}(p)$ independent of p) would not correspond to uniformity of the intrinsic cost of action ($c(p)$ independent of p) in the rational inattention model. In fact, in a version of the rational inattention model where we assume on a priori grounds that $c(p)$ is independent of p,[32] the reference distribution $\tilde{\mu}$ satisfying condition (32) would assign non-zero probability to only a discrete set of slider positions, for reasons discussed by Jung *et al.* (2015).

Once we determine the reference probabilities $\tilde{\Lambda}$ and $\tilde{\mu}$ that best account for our data, we can interpret these as direct estimates of our subjects' default pattern of behavior, under the control-cost interpretation of our model. Alternatively, we can compute the implied intrinsic costs of alternative actions, under the rational inattention interpretation. These are obtained by assuming that (34) must hold with equality for each slider position in the support of $\tilde{\mu}$, and similarly that both (35) and (36) must hold with equality, since we must have $0 < \tilde{\Lambda} < 1$ (given that we observe both trials on which the slider is adjusted and trials on which it is not). Condition (34) only establishes a lower bound for the intrinsic cost associated with slider positions that are never observed.

5.5 Empirical Fit of the Model

We now consider the parameterization of this model that best fits our experimental data. We wish to determine (*i*) the values for the information-cost parameter ψ_1 and the reference adjustment probability $\tilde{\Lambda}$ that best fit our data on the *timing* of slider adjustments, and (*ii*) the values for the information-cost parameter ψ_2 and the reference measure $\tilde{\mu}$ over new slider positions that best fit our data on the *positions* to which the slider is moved when it is adjusted. In the first case, we are choosing parameters for prediction (29) for the adjustment hazard, and in the second we are choosing parameters for prediction (24) for the distribution

[31] See for example Stahl (1990) and Mattson and Weibull (2002). Costain and Nakov (2014) consider a model like ours, in which there is a decision each period whether to adjust a continuous decision variable (a firm's price), and then a decision where to set the variable if it is adjusted; they assume that the default distribution $\tilde{\mu}$ must be a uniform distribution over some interval of candidate prices, though they let $\tilde{\Lambda}$ be a free parameter, as we do.

[32] Note that this would require $c(p) = 0$ for all p, under the normalization (23).

from which the new slider position is drawn. These problems are not independent, however, since the expected value of adjusting depends on the choice of a distribution of new slider positions, and likewise, this latter choice depends on the probability of adjustment in future periods. Hence, we estimate these parameters jointly. Moreover, we report below estimates of the relative size of the *intrinsic* costs, in addition to the estimated information costs.

As in Section 4.2, we test the fit of our model while basing our assumed continuation values on an empirical characterization of subjects' behavior in subsequent periods. This approach has the advantage that it does not require us to assume that optimal behavior is anticipated on the remaining trials as well. Testing the predictions of equations (24) and (29) for the decisions on a single trial t then requires an empirical estimate of the continuation values $V_t(p)$ for alternative slider positions p. In our model with information costs, the continuation value can be written in the form

$$V_t(p) \;=\; R_t(p) \;-\; \psi_1 H_{1t}(p) \;-\; \psi_2 H_{2t}(p), \tag{38}$$

where $R_t(p)$ is defined as in (10), and $H_{1t}(p)$ and $H_{2t}(p)$ are the expected cumulative costs (summing both information costs and the intrinsic costs of different actions) of the two decisions on subsequent trials, if the slider is set at p on trial t:[33]

$$H_{1t}(p;\pi_t) \;\equiv\; \sum_{\tau=t+1}^{T} \mathrm{E}[D(\Lambda_\tau||\tilde{\Lambda})|\hat{p}_t = p, \pi_t],$$

$$H_{2t}(p;\pi_t) \;\equiv\; \sum_{\tau=t+1}^{T} \mathrm{E}[\Lambda_\tau D(\mu_\tau||\tilde{\mu})|\hat{p}_t = p, \pi_t].$$

As in Section 4.2, we can compute a numerical estimate of the function $R_t(p)$ for each trial, on the basis of an atheoretical empirical model of the joint dynamics of the slider position and the Bayesian posterior on trials subsequent to t. The numerical estimate that we use is the same as in our estimation of the model of optimization subject to a fixed cost of adjustment. We similarly compute numerical estimates of the functions $H_1(p)$ and $H_2(p)$, using empirical models of the joint dynamics of the Bayesian posterior, the slider

[33] In writing H_{1t} and H_{2t} in this form, we make use of the reformulation (37) of the objective.

position, the adjustment probability Λ_t, and the time-varying measure over possible new slider positions μ_t. Our atheoretical characterizations of the dynamics of the adjustment probability and the measure over slider positions are estimations based on the tabulations of the experimental data shown in Figures 10 and 13.

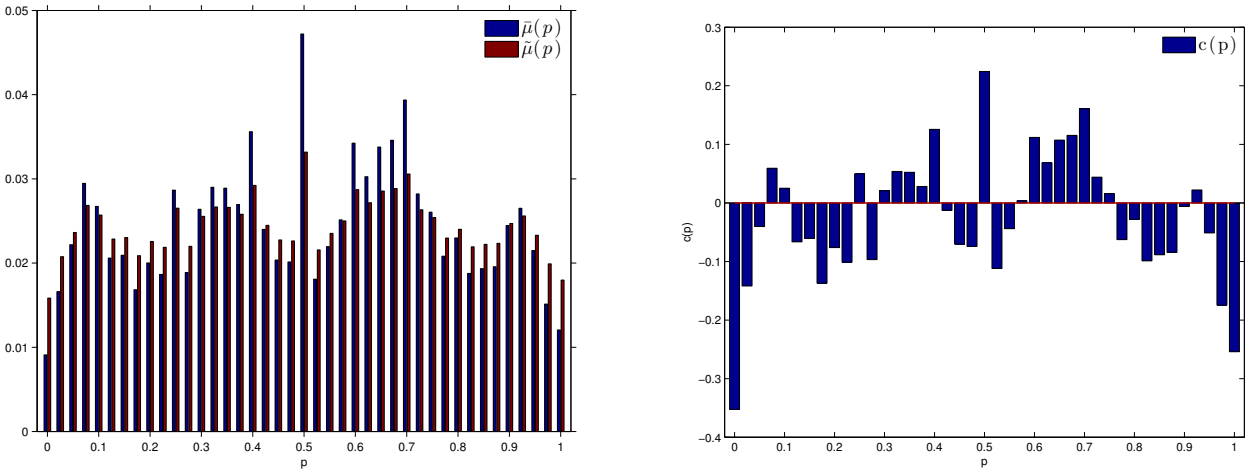
Computing the estimates of $H_1(p)$ and $H_2(p)$ requires values for $\tilde{\Lambda}$ and $\tilde{\mu}$ (in order to estimate empirical models of the dynamics of the quantities $D(\Lambda_t||\tilde{\Lambda})$ and $D(\mu_t||\tilde{\mu})$), while the estimate of $V_t(p)$ also requires values for ψ_1 and ψ_2. We start with a set of parameters and compute an estimate of the value function. Once we obtain a numerical estimate of the function $V_t(p)$, we can find the values of ψ_2 and $\tilde{\mu}$ that maximize the consistency of observed slider position choices with the model prediction (24). Using this estimated model of the slider position decision to estimate V_t^{adj}, we can then find the values of ψ_1 and $\tilde{\Lambda}$ that maximize the consistency of the observed timing of adjustments of the slider with the prediction (29). Finally, we check whether the estimated values for ψ_1, ψ_2, $\tilde{\Lambda}$, and $\tilde{\mu}$ coincide with the values assumed in computing the value function $V_t(p)$. The numerical estimates reported here represent a solution to this fixed-point problem. In order to compute the best-fitting distribution of slider choices, $\tilde{\mu}$, we assume the parametric form

$$\tilde{\mu}(p) = A\overline{\mu}(p)^{\gamma},$$

where A is a normalizing constant that ensures that $\tilde{\mu}$ is a probability distribution that integrates to 1. The parameter $\gamma \in [0, 1]$ determines how close the model-implied reference distribution is to the empirical unconditional distribution ($\gamma = 1$, which corresponds to the standard rational inattention model) versus the uninformative, uniform distribution ($\gamma = 0$, which corresponds to the alternative control cost model).

We obtain a reference distribution of slider positions $\tilde{\mu}$ shown in Figure 14. Associated with this estimate is an information cost parameter $\psi_2 = 0.64$ for the use of more information in choosing the slider position. The estimated reference distribution is similar to the unconditional empirical frequency of slider position choices, also reproduced in the figure. However, the reference distribution is estimated to be significantly flatter than the empirical distribution: we obtain a coefficient $\gamma = 0.45$. This estimate points to the presence of sig-

nificant intrinsic gains and costs in the decision of where to position the slider conditional on adjustment, since in the absence of such costs the two distributions would be identical. In particular, the estimated value of gamma implies that our subjects like extreme slider positions (close to 0 or close to 1) more than would be predicted by a rational inattention model without any intrinsic costs or gains from choosing specific slider positions.



(a) Reference and empirical distributions

(b) Cost function

Figure 14: Intrinsic costs. (a) The model-implied reference distribution $\tilde{\mu}$ of slider positions conditional on adjustment. (b) The model-implied intrinsic relative cost function, $c(p)$ associated with different slider positions.

We can go further and use the estimated reference distribution to infer the intrinsic relative cost $c(p)$ associated with movement to each possible slider position, from the requirement that (34) must hold with equality at each slider position in the support of the reference distribution (which is to say, that is predicted to be chosen with positive probability). The implied intrinsic relative cost function is also shown in Figure 14. In the figure, negative values of the intrinsic cost function are associated with preferred slider positions (which are chosen by subjects more than would be predicted by the rational inattention model without intrinsic costs), while positive cost values are associated with disliked slider positions. In addition to our subjects preferring extreme values, we also note that subjects are more likely to choose certain interior slider positions than others as well.

The degree to which this model of the slider location decision matches the experimental data is shown by Figure 15. Here we show again the scatter plot of slider positions chosen on trials where the slider is adjusted, plotted against that Bayesian posterior mean for

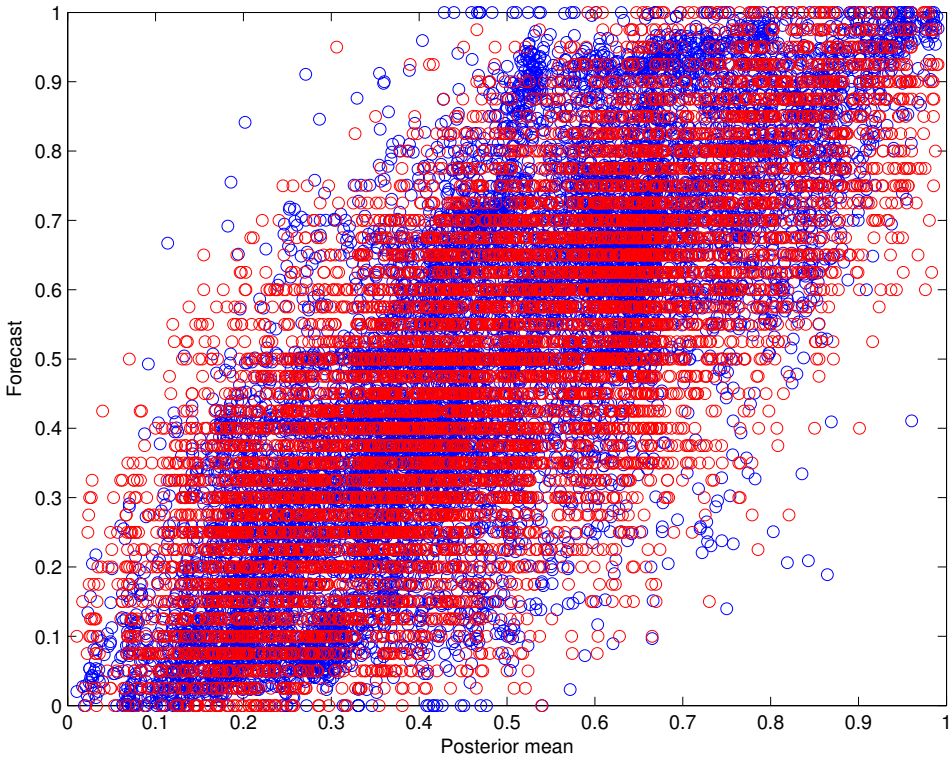


Figure 15: The model-implied joint distribution of slider settings conditional on adjustment and the Bayesian posterior mean.

the underlying state on that trial (from panel (a) of Figure 13), with the corresponding scatter plot from a simulation of the rational inattention model of the overlaid on it. In this simulation, we assume the same sequences of ring draws (and hence the same evolution of the Bayesian posterior) as in the experimental sessions, and we suppose that the slider is adjusted on exactly the same trials as in the experimental sessions (so that the set of Bayesian posteriors associated with trials on which the slider is adjusted is the same as in the experimental data); but we suppose that when the slider is adjusted, the new position is drawn from the distribution μ_t given by equation (24).

We see that the predicted joint distribution of the Bayesian posterior mean and the new slider position is reasonably similar to the one observed in our data.[34]

We turn next to the ability of the RI model to account for the timing of adjustments of

[34]We do not provide a quantitative measure of the goodness of fit, as there is no alternative stochastic model of the slider location decision with which we wish to compare this model. Because of the substantial stochastic variation in the chosen slider position, the rational inattention model is clearly much more successful in accounting for the data than the deterministic hypotheses considered above (that when the slider is adjusted, it is set to the current Bayesian posterior mean, or that it is set to the position that maximizes the continuation value).

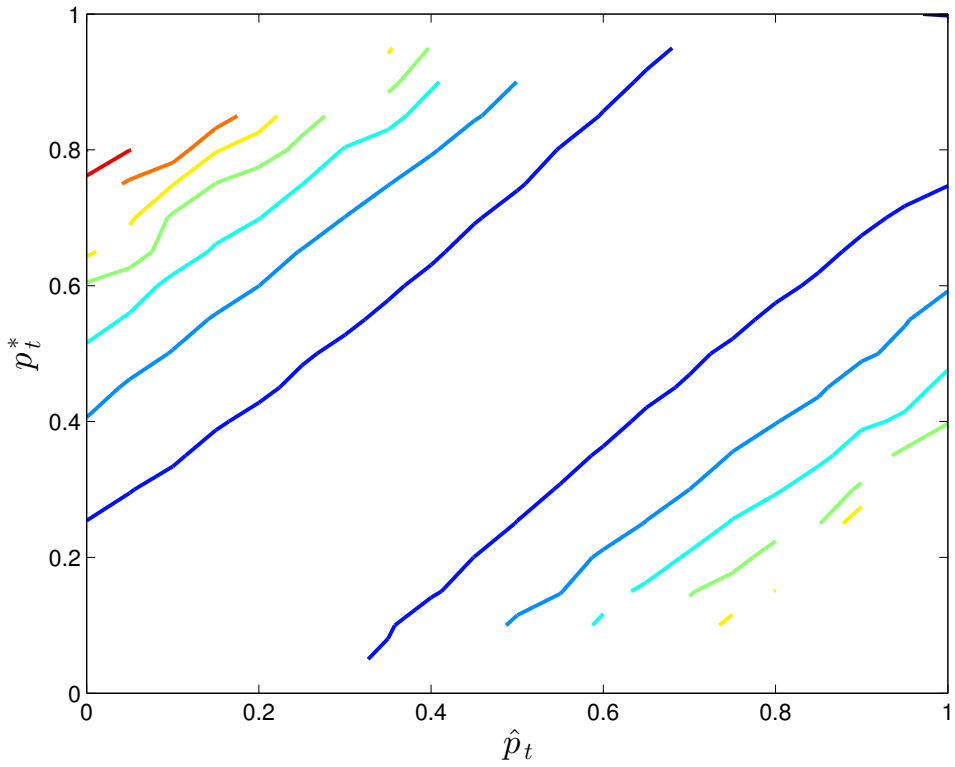
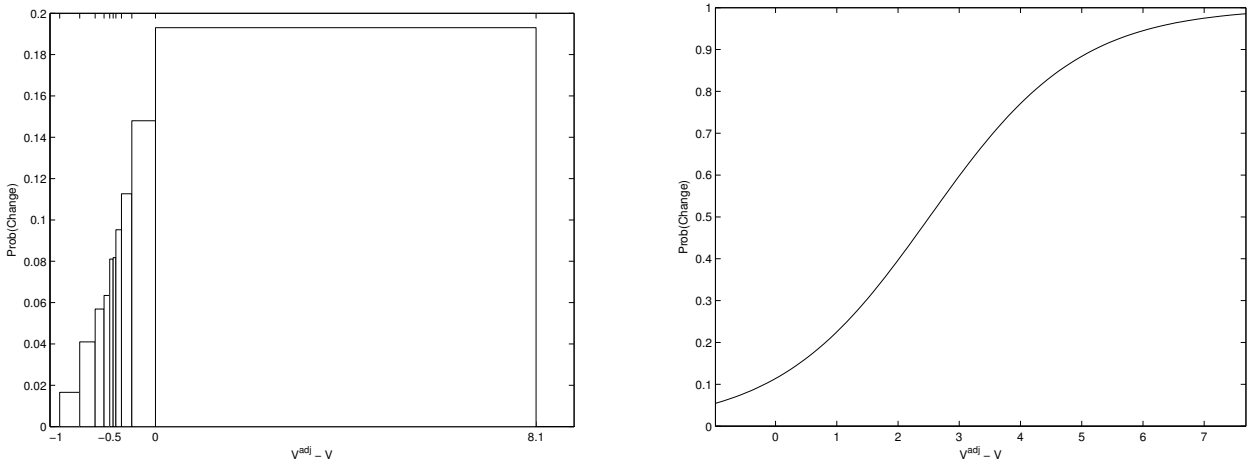


Figure 16: A contour plot of the average RI value gap Δ_t associated with alternative pairs of values for the slider position $\hat{p}_t$ [the horizontal axis] and the Bayesian posterior mean p_t^* [the vertical axis].

the slider. Like the models of the adjustment decision considered in the previous sections, the RI model implies that the adjustment hazard should be a function of a "gap" variable; but now the relevant gap Δ_t is given by (28). This differs from the value gap defined by (7) for two reasons: the continuation value function is now net of the expected information and intrinsic costs (as defined by (38)); and the expected value of adjustment is an average of the continuation values associated with different slider positions (defined by (25), rather than solely the maximum continuation value. The value gap implied by the RI model, under our estimated parameter values, is graphed in Figure 16, using the same format as Figure 11.

As a nonparametric illustration of the degree to which this gap variable is able to explain the timing of adjustments, Figure 17 plots the fraction of trials in which the slider is adjusted, grouping the trials in bins with an equal number of observations as in Figure 12, but now using the value gap (28) as the basis for classification. We see that the adjustment probability is monotonically increasing with this gap measure, as predicted by the model.

We can test the model's more precise predictions by estimating a polynomial log odds



(a) Equal-weight bins (b) Polynomial

Figure 17: Adjustment probability as a function of the value gap implied by the rational inattention model. (a) Nonparametric empirical hazard. As in Figure 12(a), boundaries are chosen so that each bin contains an equal number of observations. (b) The best-fitting logistic-polynomial hazard function as a function of the RI value gap. (Coefficients given in Table II.)

model for the hazard function, again assuming a flexible relationship of the form (3)–(4). The RI model makes a very specific prediction, that $\lambda(\Delta)$ should be an increasing linear function (equation (29)). When we consider polynomials of arbitrary order, but use the BIC criterion to select the best model family, we find that in fact our data are best fit by a model according to which the log odds are an increasing linear function of the gap, as predicted by the theory. The best fitting parameters and the associated BIC value are shown in the bottom panel of Table II. The best-fitting parametric hazard function is plotted in the second panel of Figure 17.

Equation (29) allows us to interpret the coefficients (θ_0, θ_1) of the logistic-linear hazard function in terms of the information cost ψ_1 and the reference frequency of adjustment $\tilde{\Lambda}$ of the rational inattention model; the implied values of these two parameters are also given in Table II. Interestingly, the information cost for the decision of whether to adjust is estimated to be roughly twice as large as the information cost for the decision of where to position the slider conditional on adjustment. This finding is consistent with the results from the rational inattention price setting model considered by Stevens (2015), who finds that the parameterization that best fit the micro price data implies that the decision of when to review the firm's price plan is much more costly than the decision of what price to charge in

each period.

The estimated reference frequency of adjustment is $\tilde{\Lambda} = 11.4\%$, which is larger than the empirical frequency of adjustment, $\bar{\Lambda} = 8.9\%$. These two numbers allow us to compute an implied intrinsic cost of adjustment relative to non-adjustment, using (35)–(36). The value implied by our estimates, again using the numerical estimates for the continuation value function, is $c^{adj} - c^{non} = -0.34$, implying intrinsic *gains* from adjustment, so that subjects adjust more frequently than would be optimal given the precision of their signals. In fact, as shown by the nonparametric empirical hazard function graphed in the first panel of Figure 17, subjects often adjust when the model-implied value gap is negative.

Finally, the BIC allows us to compare the fit of the RI model of the timing of slider adjustments to the models considered in the previous section and to other versions of the noisy adjustment model (also shown in the table). We see that the RI model with intrinsic costs achieves the lowest BIC of any of the models considered. Relative to the best-fitting random-fixed-cost model, which also involves a logistic-linear hazard function, the BIC of the RI model is lower by 651 points.The RI model also fits better than the best-fitting model based on the Bayesian gap, the case of a logistic-polynomial hazard. The BIC of the RI model is lower than that of the best-fitting gap-based model by 489 points.

Allowing for intrinsic costs in both the adjustment decision and the position choice substantially improves the fit of the model. The standard RI model without any psychic costs would yield estimates of the two information costs that are approximately twice as large, and a BIC that is 331 log points worse than that generated by the benchmark model with intrinsic costs.

Finally, our model also outperforms the alternative stochastic choice model of control costs (Costain and Nakov, 2014), which assumes that the reference distribution for slider position choices is fixed at the uniform distribution. The control cost model yields a BIC statistic that is 74 log points higher than that of the best fitting RI model.

Hence the rational inattention model augmented with intrinsic costs succeeds in capturing the main features of our data, at least as far as the pooled data, describing the average behavior of our population of subjects, are concerned. It predicts both the timing of adjustments (that these are stochastic, with an adjustment hazard that is a continuously increasing

function of the expected value of adjusting) and the position to which the slider is moved if adjusted (this is also stochastic, with a probability of movement to each position that is a continuously increasing function of the expected value associated with that position). Moreover, the specific functional form (29) that is predicted for the adjustment hazard function is also consistent with our data. And the model allows (though it does not require) the kind of "round-number bias" observed on the part of our subjects.

6 Discussion

We have studied how successful our subjects are in tracking a changing environmental state, and adjusting their behavior so as to continue to maximize expected reward, in a controlled laboratory setting, where we know exactly what subjects are able to observe and what they know (or should know) about the data-generating process that gives rise to those observations, so that the rational-expectations (or perfect Bayesian) benchmark is unambiguously defined. While our experimental setting is very stylized, the pattern of discrete adjustments that we observe shares some notable features with data on the adjustment of individual goods prices by retailers, as illustrated by Figure 7. This suggests that our findings about adjustment behavior in our experiment may well be relevant to understanding discrete adjustment in economically relevant settings as well, such as price-setting behavior.

Moreover, we observe a pattern of discrete adjustment reminiscent of that observed in micro-level data on prices appears even in a setting where many factors that might be thought relevant to price-setting are clearly not present: here, for example, there are clearly no costs of communication of a new policy to customers, no reasons to doubt whether customers will recognize and respond to a price change, nor any reasons to fear the emotional reaction of customers to a price change that they notice. The discreteness thus must result from some kind of cognitive constraint on subjects' ability or willingness to more closely track the reward-maximizing slider setting. And while we might think of this as reflecting decisions made on the basis of imperfect information, it is not imperfect information that results from the structure of the physical environment, as in the model of Lucas (1972); instead, behavior seems to make use of less precise information about the changing environment than has been presented to the subjects, reflecting some form of inattention, limited memory, or other limit

in information processing capacity.

Our experiment allows us to reject many familiar hypotheses about adjustment to a changing environment, at least as an explanation of the behavior of our subjects. These include, most obviously, the benchmark of rational expectations (or perfect Bayesian inference from the available data), to which we compare our experimental data in Figures 3, 4, 6, and 7. But our data also clearly reject any model that implies continuous adjustment each time a subject has a reason to adjust their beliefs,[35] such as "adaptive learning" models, the "sticky information" models of Mankiw and Reis (2002) or Reis (2006), the "noisy information" model of Woodford (2003), or the rational inattention model of Mackowiak and Wiederholt (2009), given the evidence in Figure 7 for discrete adjustment. At the same time, we can also reject time-dependent models of discrete adjustment, such as the Taylor (1980) model of staggered adjustment, since Figure 7 shows that adjustments often occur very soon after a previous adjustment of the slider (and there is also no evidence of a preferred time length between adjustments).

Our experimental data on the timing of discrete adjustments are inconsistent with the assumption of the Calvo (1983) model of price adjustment, according to which an adjustment is equally likely to occur over any time interval, regardless of the degree of appropriateness of the existing position of the decision variable to current conditions. We have demonstrated in Figures 10, 12, and 17 that it is possible to define measures of the degree of inappropriateness of the existing slider position which have some ability to predict the timing of slider adjustments; and hazard functions that allow some degree of dependence on these "gap" measures fit the data better than one that assumes no state-dependence of the adjustment decision, to an extent that results in a lower BIC criterion, despite the penalty for additional free parameters in the case of the state-dependent models.

Yet the kind of state-dependence indicated by our data is not the kind assumed in the common type of state-dependent models of discrete adjustment (such as the literature on state-dependent pricing), namely an "*Ss*" model for the timing of adjustments. In the case

[35] Note that the sticky information models of Mankiw and Reis imply that there is no continuous updating of information, but that subjects' decision variables should nonetheless be constantly adjusted, if their estimate of their optimal action changes despite an unchanged information set. In our experiment, this would be the case, because of the possibility of a regime change after each ring is drawn.

of none of the "gap" measures that we consider do we find evidence of sharp thresholds such that adjustment occurs if and only if a threshold is crossed; nor do the data support a generalization of "*Ss*" models in which the probability of adjustment changes discontinuously at the threshold from some positive value to some higher probability that is still less than 1. It might be thought that our failure to find a sharp threshold reflects the fact that our gap measures are inaccurate proxies for the incentive for adjustment that actually determines the timing of adjustments; but the fact that we observe slider adjustments of many sizes, including a large number of very small adjustments but also a surprising number of extremely large adjustments (Figure 7) is also difficult to reconcile with an "*Ss*" model, regardless of the nature of the correct gap measure. Under an "*Ss*" model, if adjustments do not occur constantly (implying that there must be a substantial distance between the threshold that one must reach to justify adjustment and the optimal setting that is chosen when adjustment occurs), then one should never observe small adjustments. And one should also see few really large adjustments, since the threshold required for adjustment should typically be crossed before the size of the gap between the current setting and the optimal one can grow that large. (Figure 7) shows that large changes in the optimal slider position from one trial to the next are very infrequent.) The distribution of adjustment sizes that we observe is inconsistent with an "*Ss*" model on both counts.

The kind of models that are instead consistent with at least the broad features of our data are models of discrete adjustment in which the timing of adjustment is stochastic, though the probability of adjustment is a continuous, increasing function of some measure of the inappropriateness of the current slider setting. Models of this kind include the "generalized *Ss* model" proposed by Caballero and Engel (1993), the optimizing models with a random fixed cost of adjustment proposed by Caballero and Engel (1999) and Dotsey, King and Wolman (1999), rational inattention models of discrete adjustment like the one proposed by Woodford (2009), and the "logit model" of discrete adjustment proposed by Costain and Nakov (2014). Among the several models of this type that we fit to the timing of slider adjustments in our data, we find that the best-fitting model (under a BIC criterion) is a generalization of the rational inattention model of Woodford (2009), which can also be viewed as a generalization of the logit model of Costain and Nakov (2014). In this type of

model, the randomness of decision whether to adjust the slider at a given point in time is interpreted as resulting from inattention, or more generally from imprecision in the subject's subjective awareness of the precise situation that has been revealed by the sequence of ring draws observed to that point (due, perhaps, to a memory limitation). Such models are also able to explain why we observe a leptokurtic distribution of adjustment sizes, as shown in Woodford (2009) and Costain and Nakov (2014), though we do not analyze this issue quantitatively here.

Our data are nonetheless inconsistent with the model of Woodford (2009) — and similarly with the "generalized *Ss* model" of Caballero and Engel (2003), or optimizing models with a random fixed cost — in one important respect. These models (like standard models of optimization subject to a fixed cost, the Calvo model, and most other models of discrete adjustment in the economics literature) assume that when adjustment *does* occur, the decision variable is immediately moved to the currently optimal setting. This implies in particular that the setting chosen conditional upon adjustment is a deterministic function of the state (or more precisely, the information available to the decisionmaker) at the time of adjustment. Instead, Figure 13 shows that our data are inconsistent with the hypothesis of immediate adjustment to the optimal slider setting when the slider is moved, whether the current "optimal setting" is interpreted myopically (as the current best estimate of the hidden state, p_t^*) or in a way dictated by dynamic programming (as the position p that would maximize the continuation value $V_t(p)$). It seems instead that the best account of our data is provided by a model that assumes that there is an important random element in the slider location decision, while at the same time not assuming that the location chosen is entirely arbitrary. A model that fits our pooled data reasonably well is a rational inattention model of this decision as well, or equivalently a generalization of the logit model proposed by Costain and Nakov (2014).[36]

The fact that our data are best fit by a model of inattentive choice of the new slider position, conditional on the slider being moved, makes it plausible that inattention, rather than a random cost of adjustment, is responsible for the randomness in the timing of adjustments

[36]The kind of model of the location decision that best fits our data differs from the kind proposed by Costain and Nakov, in that they assume *a priori* that the reference measure $\tilde{\mu}$ should be a uniform distribution, while we find a better fit with a non-uniform distribution

as well. Thus the model of the adjustment dynamics most supported by our data is one that attributes departures from the benchmark of fully optimal behavior to inattention on the subject's part — but that proposes that behavior is nonetheless close to optimal, given the amount of information about the sequence of ring draws that is used in adjusting the slider.

The similarity of the pattern of adjustment of the slider in our experiment to retailers' adjustments of the prices that they charge raises the possibility that the failure of such prices to more perfectly track the retailer's currently optimal (profit-maximizing) price should be similarly attributed to rational inattention of the kind that we model. Such a conclusion would have important consequences for monetary economics, as discussed in Woodford (2009) and Stevens (2015). We believe that experimental studies such as this one can play an important role in advancing our understanding, not simply of the general importance of inattention in adjustment dynamics, but of the specific models of inattentive decisionmaking that are most consistent with what we know about human capabilities — and by doing so, can contribute to the construction of more empirically realistic macroeconomic models.

References

Alvarez, Fernando E., Francesco Lippi & Luigi Paciello (2011), "Optimal Price Setting With Observation and Menu Costs," *The Quarterly Journal of Economics* 126(4): 1909–1960.

Anderson, Simon, Jacob Goeree & Charles Holt (2002), "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," *Southern Economic Journal* 69(1): 21–47.

Burnham, Kenneth P. & David R. Anderson (2002), *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, Springer, New York, 2d ed.

Caballero, Ricardo J. & Eduardo M.R.A. Engel (1993), "Microeconomic Adjustment Hazards and Aggregate Dynamics," *The Quarterly Journal of Economics* 108(2): 359–383.

Caballero, Ricardo J. & Eduardo M.R.A. Engel (1999), "Explaining Investment Dynamics in U.S. Manufacturing: A Dynamic *(S,s)* Approach," *Econometrica* 67(4): 783–826.

Caballero, Ricardo J. & Eduardo M.R.A. Engel (2007), "Price Stickiness in *Ss* Models: New Interpretations of Old Results," *Journal of Monetary Economics* 54(S): 100–121.

Calvo, Guillermo (1983), "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12: 383–398.

Cheremukhin, Anton, Anna Popova & Antonella Tutino (2011), "Experimental Evidence on Rational Inattention," *Federal Reserve Bank of Dallas Working Paper* 1112.

Costain, James & Anton Nakov (2014), "Logit Price Dynamics," *European Central Bank Working Paper* 1693.

Cover, Thomas M. & Joy A. Thomas (2006), *Elements of Information Theory*, Wiley-Interscience, New York, 2d ed.

Dotsey, Michael & Robert G. King (2005), "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics* 52(S): 213–242.

Dotsey, Michael, Robert G. King & Alexander L. Wolman (1999), "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics* 114: 655–690.

Gallistel, Charles R., Monika Krishan, Ye Liu, Reilly Miller & Peter E. Latham (2014), "The Perception of Probability," *Psychological Review* 121(1): 96–123.

Hansen, Mark H. & Bin Yu (2001), "Model Selection and the Principle of Minimum Description Length," *Journal of the American Statistical Association* 96: 746–774.

Jung, Junehyuk, Jeong-Ho Kim, Filip Matejka & Christopher A. Sims (2015), "Discrete Actions in Information-Constrained Decision Problems," *unpublished* .

Kaplan, Greg & Guido Menzio (2015), "The morphology of price dispersion," *International Economic Review* 56(4): 1165–1206.

Levy, Daniel, Dongwon Lee, Haipeng Chen, Robert J. Kauffman & Mark Bergen (2011), "Price points and price rigidity," *Review of Economics and Statistics* 93(4): 1417–1431.

Lucas, Jr., Robert E. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4: 103–124.

Maćkowiak, Bartosz A. & Mirko Wiederholt (2009), "Optimal Sticky Prices under Rational Inattention," *The American Economic Review* 99(3): 769–803.

Magnani, Jacopo, Aspen Gorry & Ryan Oprea (2016), "Time and state dependence in an Ss decision experiment," *American Economic Journal: Macroeconomics* 8(1): 285–310.

Mankiw, N. Gregory & Ricardo Reis (2002), "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics* 117(4): 1295–1328.

Mattson, Lars-Goram & Jorgen Weibull (2002), "Probabilistic Choice and Procedurally Bounded Rationality," *Games and Economic Behavior* 41: 61–78.

Peterson, Cameron R. & Lee R. Beach (1967), "Man as an Intuitive Statistician," *Psychological Bulletin* 68: 29–46.

Reis, Ricardo (2006), "Inattentive Producers," *The Review of Economic Studies* 73(3): 793–821.

Ricci, Matthew & Randy Gallistel (2016), "Accuracy and Structure Detection in the Perception of Probability," *unpublished* .

Robinson, Gordon H. (1964), "Continuous Estimation of a Time-Varying Probability," *Ergonomics* 7: 7–21.

Sheshinski, Eytan & Yoram Weiss (1977), "Inflation and Costs of Price Adjustment," *Review of Economic Studies* 44: 287–303.

Sheshinski, Eytan & Yoram Weiss (1983), "Optimum Pricing Policy Under Stochastic Inflation," *Review of Economic Studies* 50: 513–529.

Sims, Christopher A. (2003), "Implications of Rational Inattention," *Journal of Monetary Economics* 50(3): 665–690.

Sims, Christopher A. (2011), "Rational Inattention and Monetary Economics," in *Handbook of Monetary Economics, volume 3A*, B.M. Friedman & M. Woodford, eds., Elsevier, Amsterdam.

Stahl, Dale (1990), "Entropy Control Costs and Entropic Equilibrium," *International Journal of Game Theory* 19: 129–138.

Steiner, Jakub, Colin Stewart & Filip Matejka (2015), "Rational Inattention Dynamics: Inertia and Delay in Decisionmaking," *unpublished* .

Stevens, Luminita (2015), "Coarse Pricing Policies," *unpublished* .

Taylor, John B. (1980), "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88: 1–23.

Van Damme, Eric (1991), *Stability and Perfection of Nash Equilibrium*, Springer, New York, 2d ed.

Woodford, Michael (2003), "Imperfect Common Knowledge and The Effects of Monetary Policy," in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, J. Stiglitz P. Aghion, R. Frydman & M. Woodford, eds., pp. 25–58, Princeton University Press, Princeton, NJ.

Woodford, Michael (2009), "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics* 56(S): 100–124.

A Appendix: Optimal Bayesian Inference

Given a sample of T observations, we wish to determine the posterior distribution over (n, p), where p is the most recent probability of drawing a 1 (namely a green ring) and n is the number of periods for which the current regime has lasted so far. A model of the data is specified by a probability p of drawing a 1 in the most recent regime and a partition $\pi = \{n_i\}$ of the sample into successive regimes, where n_i is the length of regime i. Let τ_i denote the last observation of regime i. The likelihood of the most recent n observations if the regime has been p over that time is

$$L(n, p) = p^{k_n}(1-p)^{n-k_n}, \tag{A.1}$$

where k_n is the number of successes in the n most recent observations. Let

$$L(n) \equiv \int L(n, p) f(p) dp. \tag{A.2}$$

and let $L_\tau(n)$ denote the average likelihood computed using the n observations ending with observation τ. The ex-ante joint probability of the model (π, p) being correct and the data being a particular observed sequence is given by

$$\mu(\pi)(\prod_{i=1}^{N(\pi)-1} L_{\tau_i}(n_i)) f(p) L(n, p),$$

where $N(\pi)$ is the number of regimes under partition π and μ_π is the ex-ante probability of partition π occurring in a sample of length T,

$$\mu(\pi) = (1-\delta)^{T-N(\pi)+1}(\delta)^{N(\pi)-1}. \tag{A.3}$$

Summing over the set $\Pi(n)$ of all possible partitions for which the final regime is of length n, we define

$$Q(n) \equiv \sum_{\pi \in \Pi(n)} \mu(\pi) \prod_{i=1}^{N(\pi)-1} L_{\tau_i}(n_i). \tag{A.4}$$

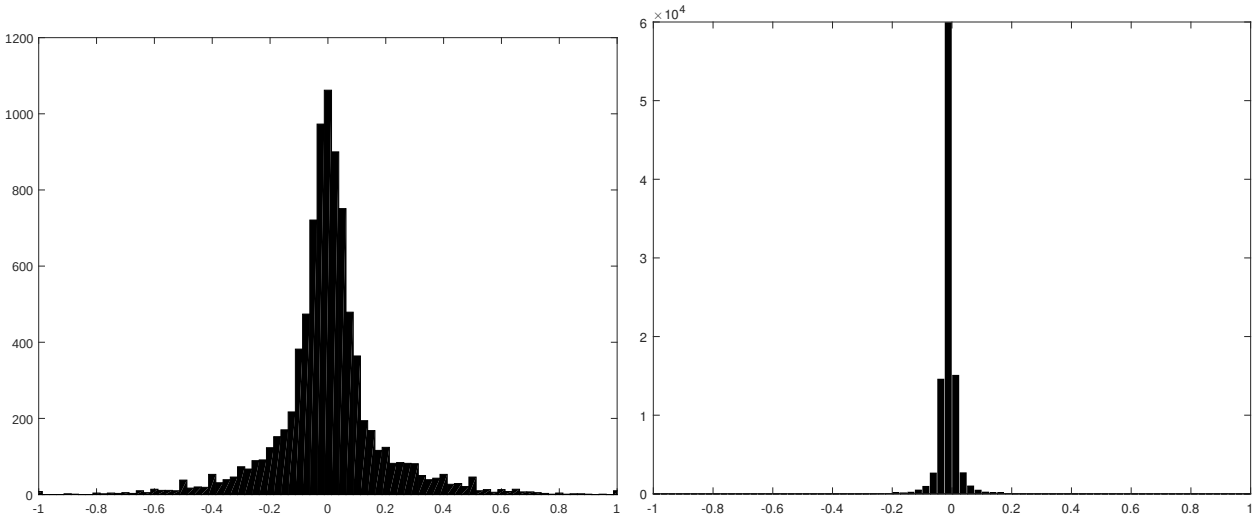
The posterior probability of (n, p) is

$$P(n, p) = \frac{Q(n) f(p) L(n, p)}{\Sigma_{n \geq 1} Q(n) L(n)}. \tag{A.5}$$

The expected value of p sums over all n and integrates over p using the measure $P(n, p)$. The Bayesian estimate for the probability of drawing a 1 on the next observation takes into account the fact that the regime might change on the next draw, which occurs with probability δ, and in which case, the estimate of the probability is 0.5:

$$B = (1-\delta) \int \sum_{n \geq 1} p P(n, p) dp + \frac{\delta}{2}. \tag{A.6}$$

Figure A.1 plots the distribution of the size of slider adjustments in our data (left panel) and in the Bayesian benchmark. There is considerable variation in the size of the slider adjustments when they occur, whereas according to the Bayesian benchmark, the slider should be adjusted on every trial, and the adjustments should seldom be very large, as shown by the right panel of the figure.



(a) Sample (b) Bayes

Figure A.1: Distribution of the size of changes in the slider position $\Delta\hat{p}_t$. Left panel: our data, counting only trials on which the slider is moved (all subjects and all sessions). Right panel: prediction under the hypothesis of Bayesian rationality.