American Economic Journal: Macroeconomics 2016, 8(1): 285–310
http://dx.doi.org/10.1257/mac.20130267

Time and State Dependence in an Ss Decision Experiment†

By Jacopo Magnani, Aspen Gorry, and Ryan Oprea*

Flow earnings in a laboratory experiment decline the further a Brownian state variable, z, evolves from its optimal level, z^*. Optimal state dependent models predict subjects will pay a fixed cost to return z to z^* only when z strays outside a critical inaction region around the optimum. On average, subjects adjust at states remarkably close to optimal threshold levels but, as in the field, do not establish true "state dependent" inaction regions, suggesting significant "time dependent" components in adjustment rules. Structural estimates of subjective observation cost qualitatively account for variation in time dependence observed across treatments. (*JEL* C91, D21, D80)

When the world changes, firms must change with it to continue optimizing. But when it is costly to make changes, it pays to be judicious about timing. Indeed, Ss models show that when the future is uncertain and adjustment is costly, it is best to ignore minor deviations from optimality and delay reoptimizing until failing to do is sufficiently costly. Such decision problems are canonical, arising throughout economics. Indeed, Ss models have become a workhorse in the last three decades, helping economists understand sticky prices, inventory dynamics, cash flows, hiring and firing decisions, and investments.[1] Ss models point out that agents should choose when to reoptimize soley based on the state of the world ("state dependent" adjustment) and not on any direct function of time itself ("time dependent" adjustment).

* Magnani: Krannert School of Management, Purdue University, 403 W. State Street, West Lafayette, IN 47907 (e-mail: jacopo.magnani@googlemail.com); Gorry: Department of Economics and Finance, Utah State University, 3565 Old Main Hill, Logan, UT 84322 (e-mail: aspen.gorry@gmail.com); Oprea: Department of Economics, University of California, Santa Barbara, 3014 North Hall, Santa Barbara, CA 93106 (e-mail: roprea@gmail.com). We are grateful to Andrew Caplin, James Costain, Daniel Friedman, Charles Holt, and audiences at session 7 of the 2013 Stanford Institute for Theoretical Economics, the 2013 North American Economic Science Association meetings, and the WZB Berlin Social Science Center for helpful comments and James Pettit for developing the experimental software.

† Go to http://dx.doi.org/10.1257/mac.20130267 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] As summarized in Caplin and Leahy (2010), the development of Ss models grew out of an interest in understanding how to manage inventories. Arrow, Harris, and Marschak (1951) showed that Ss policies are cost minimizing in infinite horizon inventory problems while Scarf (1960) demonstrated that optimal decisions in a large class of such inventory problems take an Ss form. Given the pervasiveness of both uncertainty and fixed costs of adjustment in economics, Ss models have since been deployed in a wide range of applications in economics including the interaction of portfolio choice and housing demand (Grossman and Laroque 1990); the demand for durables (Bertola and Caballero 1990); portfolio choice and stock ownership (Vissing-Jørgensen 2002); investment (Pindyck 1988, and Abel and Eberly 1994); the decision of countries to exit the Eurozone (Alvarez and Dixit 2013); and price adjustment (Barro 1972 and Sheshinski and Weiss 1977).

Do people use optimal state dependent rules when making adjustments? Among the best studied Ss-type problems empirically is firms' timing of price changes in the face of "menu costs." A central ambition of the empirical literature on price adjustment is to discern whether firms tend to employ the sorts of state dependent price changes predicted in Ss models or whether they instead employ "time dependent" rules of the sort described by Taylor (1980), a question of critical importance to monetary policy.[2] Although state dependent models are empirically successful in several dimensions, much of the evidence so far suggests that they fail to describe the data in important ways: firms also show some evidence of time dependence in their actual decision making.[3] To illustrate, Figure 1 plots log price adjustments using AC Nielsen data (the data were constructed in Midrigan (2008) and used as an empirical benchmark in Costain and Nakov (2011a, 2011b)). As in state dependent models, adjustments are bimodal with modes above and below zero. However, as in time dependent models the sizes of price adjustments are diffuse and not tightly clustered around modes; adjustment data provides little evidence that firms establish the inaction regions predicted by Ss models.[4]

There are two broad classes of explanations available for the predictive failures of Ss models. One class is institutional: the environments studied in the field may include institutional features that rationalize time dependence but are not included in the structure of simple Ss models.[5] Another is cognitive: the models have the description of the decision problem right but limits in cognitive capacity or attention prevent firms from fully assessing or responding to the current state.[6] Such boundedly rational agents' adjustment rules include properties of time dependent models that can generate patterns like those in Figure 1.

[2] While monetary policy has sizable real effects when agents behave in a time dependent manner, monetary policy can be largely, or even entirely, ineffective with state dependent agents, a fact first highlighted by Caplin and Spulber (1987).

[3] Indeed, the recent empirical literature has produced mixed support for both pure state and time dependent models. While Midrigan (2010) finds significant evidence of selection effects in firms' price changing decisions, consistent with state dependence, Nakamura and Steinsson (2008) find evidence of seasonality in price changes and no evidence for the upward sloping hazard functions of price adjustment that are predicted by standard state dependent models. Klenow and Kryvtsov (2008) find that purely state dependent models are unable to generate as many small and large price changes as are observed in the data. On the other hand, they also find that the correlation between the time between adjustments and the size of adjustments is weaker than would be expected from a purely time dependent model.

[4] Of course underlying parameters of the economy and institutional features of the decision task are unobservable in naturally occurring data such as this, making it difficult to make direct and easily interpretable comparisons to theoretical models. This observability problem is a central motive for conducting direct tests using controlled and sterile experiments as a complement to conventional empirics.

[5] In order to explain the empirical distribution of price changes, "second-generation" state dependent pricing models have been developed that rationalize the empirical evidence by including additional institutional details in the model. Some of these competing explanations are: variable menu costs as in Dotsey, King, and Wolman (1999); infrequent shocks under small menu costs as in Gertler and Leahy (2008); and economies of scope for changing the prices of multiple products as in Midrigan (2011) and Alvarez and Lippi (2012). For a survey of the literature, see Klenow and Malin (2011).

[6] Alvarez, Lippi, and Paciello (2011) show that costly information processing (of a sort modeled in Caballero 1989; Reis 2006; and Alvarez, Lippi, and Paciello 2012) in a menu cost Ss model can also generate an optimal blend of time and state dependent behavior. Alternately, Woodford (2008) builds a model based on the concept of rational inattention (costly information gathering) developed by Sims (1998, 2003, 2006) where the standard Ss model is generated as a limiting case of no information costs. More recently (Costain and Nakov, 2011a, 2011b) develop a similar theory of smoothly state dependent pricing that assumes agents face cognitive costs in processing and using information to make optimal decisions.

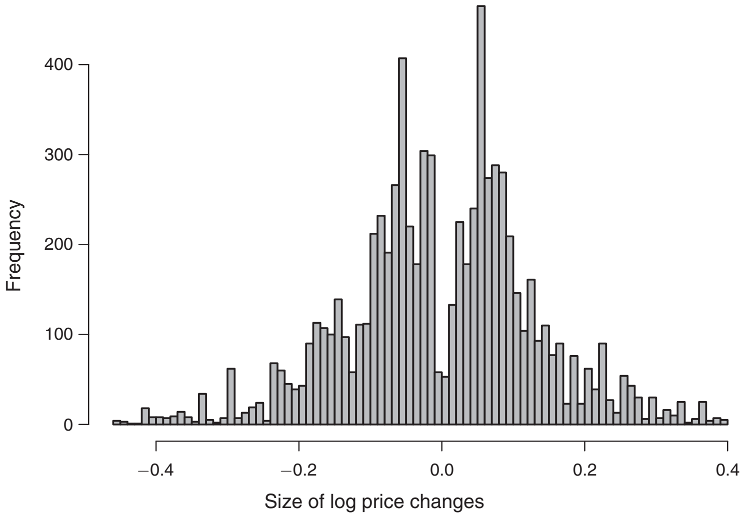

Figure 1. Histogram of AC Nielsen Data of log Price Adjustment Data Were Constructed by Midrigan (2008) and Used as an Empirical Benchmark in Costain and Nakov (2011a, 2011b)

It is difficult to directly judge between these two classes of explanations using conventional empirics because the underlying decision problem faced by firms is only imperfectly observed in naturally occurring data. In this paper we use laboratory experiments to help shed some new light on this question. Experiments have, for our purposes, one important advantage over the field: they give the analyst perfect information and perfect control over the exact decision problem faced by subjects. If we precisely implement the simple environment of a standard Ss pricing model in the controlled and sterile setting of the lab and observe deviations in behavior matching the distinctive ones observed in the field, the cause can only be bounded rationality;[7] alternative explanations are unavailable in our laboratory environment.

When designing the experiment, we made an ex ante choice to compare standard Ss predictions to the alternative predictions of a specific model of bounded rationality. We chose Alvarez, Lippi, and Paciello (2011),[8] which adds to a standard Ss model subjective costs incurred by agents when they track the state or (equivalently) assess its relationship to the optimal state.[9] Optimizing against these costs generates

[7] Here and elsewhere we use "bounded rationality" to describe settings in which deviations from orthodox predictions are driven by agents optimizing in the face of subjective costs (such as costs of cognition, attention, or perception) that are usually not included in standard models. In these models, agents are rational in the sense that they make constrained optimal decisions given these costs.

[8] Among the large menu of bounded rationality models to emerge in recent years, we chose Alvarez, Lippi, and Paciello (2011) because the model conveniently has limiting cases of purely "state dependent" and purely "time dependent" behavior depending on the magnitude of the observation and menu costs; and it predicts a distinctive effect of volatility on behavior that does not arise in standard Ss models and therefore allows us to conduct a simple (and therefore quite credible) treatment level test of the model in the lab.

[9] In Alvarez, Lippi, and Paciello (2011) subjects must pay a cost each time they want to "review" information about the state. While this model of paying a fixed observation cost to observe the state is possibly narrow, the model accurately captures any cognitive cost that the agent faces in observing or processing information about the state as long as the agent must pay such costs before gaining any information about the current state. This is broad

time dependence in decision making that can resemble in several key respects patterns observed in naturally occurring data. An important component of the structure of our study is that the specific Ss model we took to the lab, our experimental treatments, and our structural estimation were all designed with Alvarez, Lippi, and Paciello (2011) in mind prior to running the experiment. We thus tied our hands on the source of our alternative hypotheses, our structural model, and our comparative statics (see below) prior to collecting the data reported here.

In our experiment, subjects observe a graphical visualization of their static price and are also shown the evolution of the optimal relative price, which fluctuates according to a geometric Brownian motion. The further the optimal price strays from the current price (either above or below), the lower their flow earnings. A subject decides only when to pay a fixed monetary cost to adjust her price to the current optimum. Ss theory gives us clear predictions: subjects should set thresholds above and below the optimum, exercising their option to adjust only when the relative price passes one of these thresholds. Because we have access to the exact parameters of the game, we can precisely evaluate the theory's predictive power and compare it to time dependent alternatives (specifically, Alvarez, Lippi, and Paciello 2011)—something that cannot be done using naturally occurring data.

Our first contribution is to show that subjects' adjustment patterns in the experiment match those observed in the field. As in the field, adjustments in our data are bimodal, conforming with the two-sided adjustment rules predicted by the theory. Moreover, because we have access to the induced parameters of the economy we are able to show, for the first time, that these adjustment points are, on average, nearly identical to those predicted in Ss models and obey Ss predicted comparative statics to a remarkable degree. However, as in the field, adjustments are extremely diffuse: subjects often make price adjustments far from optimal threshold points and do not establish true inaction regions. Subjects therefore show strong evidence of the same sort of time dependence measured in the field, though here only bounded rationality is a plausible explanation.

Our second contribution is to examine comparative static predictions both from the orthodox Ss model and from a bounded rationality extension as in Alvarez, Lippi, and Paciello (2011). We focus attention on volatility, a variable of recent interest in the Ss empirical literature (e.g., Bloom, Bond, and Van Reenen 2007; and Vavra 2013). In Ss models, an increase to volatility has a distinct predicted effect: it will push out adjustment thresholds, increasing the size of average price changes. However, Alvarez, Lippi, and Paciello (2011) predict a second effect of volatility not predicted in orthodox models: it directly decreases state dependence, leading to greater dispersion in price changes.[10] By varying volatility in our design we therefore generate both a test for state dependent comparative statics (by looking for a change in the average size of adjustments) and a test for the presence of the type of

enough to capture ideas of costly state observation, costly attention, costly acquisition of information, and costly information processing.

[10] This effect arises because the decision maker establishes time-dependent inattention intervals between observations that directly produce greater variability when the state is more volatile. Though this variability is costly, an optimizing decision maker only partially compensates by paying closer attention to the state in Alvarez, Lippi, and Paciello (2011) leaving an effect of volatility on variability.

bounded rationality in Alvarez, Lippi, and Paciello (2011) (by looking for a change in the variability of adjustments). We find strong evidence of both. Price changes in our high volatility (VOL) treatment are significantly larger than changes in our baseline (BASE) treatment as the Ss model would predict. Moreover, we find that VOL adjustments are far more diffuse and thus exhibit weaker state dependence than in the BASE treatment, as Alvarez, Lippi, and Paciello (2011) predict.

Our third contribution is to dig deeper into the interpretation of the subjective costs measured in the experiment. One possibility is that these costs are essentially driven by difficulties in computing (or learning) the optimal adjustment rule. An alternative is that they are driven by difficulties in paying attention to or exerting control over the timing of adjustment. In order to distinguish between these two classes of explanation, we ran the GUIDE treatment. In GUIDE we implemented an exact replication of the BASE treatment except that we told subjects that a threshold rule was profit maximizing; and showed them the thresholds at which to adjust, *on their screens and throughout the experiment*. Remarkably, decision making in GUIDE is no more optimal and no more state dependent than decision making in BASE. Typical deviations from the optimum and measures of diffusion were of the same magnitude and earnings were nearly identical. These results strongly suggest that the bounded rationality observed in our experiment are due to costs of attention and control rather than cognitive or learning constraints.

Our final contribution is to use the unique precision of laboratory data to assess the predictions of the model in Alvarez, Lippi, and Paciello (2011) by structurally estimating subjective observation costs at the individual level. Laboratory data give us access to all of the relevant parameters of the game and therefore allow us to precisely estimate the level of observation costs that are consistent with each subject's observed behavior in the lab. Because of this, we can collect unusually credible structural estimates of the subjective costs of implementing state dependent decision rules for every subject in our sample.[11] By comparing these subject-level cost estimates across treatments we are able to assess the degree to which forces described in Alvarez, Guiso, and Lippi (2012) organize our results. For example, we find that distributions of estimates are identical in BASE and GUIDE subjects, confirming that the bounded rationality in our data does not arise from difficulties in learning optimal behavior. On the other hand, these costs are estimated to be larger in the VOL treatment than in the BASE treatment—a finding that suggests that volatility has additional effects on the cost of implementing state dependent strategies beyond those modeled in Alvarez, Lippi, and Paciello (2011).

Although our experiment is framed as a price adjustment problem, our findings have a broader scope: costly adjustment against uncertain fundamentals arises throughout economics. As a result, our findings shed light on an important general component of economic decision making. Indeed, given the importance of Ss models in economics, it is surprising that (to the best of our knowledge) no previous experimental studies have examined behavior in these common settings. The closest

[11] Alvarez, Guiso, and Lippi (2012) comes closest to our paper on this by directly estimating costs of observation of the state variable in an investment problem using a novel household survey dataset. While the mechanism they study is related to ours, the estimate is from a very different setting than firms' pricing decisions.

study to ours is perhaps Oprea, Friedman, and Anderson (2009) who study how subjects learn to value real options in a related dynamic-stochastic environment, though the economic forces studied in that paper are quite different.[12] More distantly related are several experiments that study pricing frictions in macroeconomic-inspired settings. Davis and Korenok (2011) study a multiplayer market in which agents are directly induced to use time-dependent pricing rules and examine how these timing frictions effect the speed of price convergence over a number of trading periods. Wilson (1998) studies the behavior of a monopolist facing nominal price shocks and shows that even when faced with menu costs, prices converge quickly to monopolistic levels in dynamic settings. Finally, Noussair, Pfajfar, and Zsiros (2013) implement a multiagent dynamic stochastic general equilibrium (DSGE) economy and show that menu cost treatments reduce the variability of inflation relative to other treatments. None of these papers, however, study a Ss-type setting and none study the contrast between state and time dependence that animates our environment. Duffy (2008) provides a comprehensive overview of the experimental macroeconomics literature.

Our paper also relates to a broader literature on bounded rationality and cognitive costs, usefully surveyed in Conlisk (1996), who notes that "deliberation about an economic decision is a costly activity, and good economics requires that we entertain all costs." A recent example by Caplin, Dean, and Martin (2011) develops and experimentally tests a model of satisficing where cognitive fixed costs shape sequential search over a choice set. They find that these costs increase with complexity, a result that seems related to our finding that volatility increases such costs in Ss decisions.

The rest of the paper is organized as follows. In Section I, we describe a simple Ss model inspired by the menu cost literature and compare its predictions with a model that includes both menu and observation costs that nests both state and time dependent models as special cases. In Section II, we discuss our strategy for implementing the model in the lab and motivate our experimental design. In Section III, we present our results. Finally, we conclude with a discussion in Section IV. Theoretical details and instructions to subjects are collected in online Appendices.

I. Theoretical Model, Ss Predictions, and Bounded Rationality

In this section we discuss the simple Ss pricing model that we take to the lab, and compare its optimal state dependent predictions to purely time dependent and boundedly rational alternatives. Section IA describes optimal state dependent predictions for a a standard, continuous time Ss pricing model with menu costs as described by Stokey (2008) (these environments were first studied by Barro (1972) and Sheshinski and Weiss (1977)). Full details on the motivating economic problem

[12] In Oprea, Friedman, and Anderson (2009), subjects choose when to make a one-time fixed investment in a project whose value evolves according to a geometric Brownian motion with positive drift. The reward function, action space and predictions are very different from those in the Ss models studied here. Moreover, while Oprea, Friedman, and Anderson (2009) are focused on understanding learning rules driving aggregate long-run time series, ours is focused on understanding cross sectional bounded rationality and time dependent departures from threshold adjustment rules.

and technical details on the solution are provided in the online Appendix. Section IB compares optimal state dependent adjustment to time dependent adjustment (specifically the Taylor 1980 adjustment rule). We then describe a model with both menu costs and bounded rationality—modeled as costly observation as in Alvarez, Lippi, and Paciello (2011)—that features purely state and time dependent adjustment behaviors as special cases.

A. *Model and State Dependent Optima*

Consider a firm that maximizes its flow profits at each moment by minimizing the expected discounted value of a quadratic instantaneous loss function:

$$B(p - p^*)^2, \tag{1}$$

where B is an exogenous constant parameter, p is the current (log) price and p^* is the target price, i.e., the price that maximizes instantaneous profits.[13] The target price follows a Brownian motion:

$$dp^*(t) = \mu dt + \sigma dw(t) \tag{2}$$

with drift μ, volatility σ, where $w(t)$ is a standard Wiener process. For simplicity and to match the setting of our experiment, we restrict attention to the case of $\mu = 0$.[14] The problem can be described entirely in terms of a single state variable: the price gap, $z = p - p^*$. The firm's task is to decide whether to exert control over the state, z, by discontinuously adjusting it to the optimal level z^* at each moment in time. While typical Ss problems have the firm choose the optimal level z^*, we simplify the problem for subjects by assuming that z^* is known. The firm faces a fixed cost of m each time it makes an adjustment and thus faces a sequence of optimal stopping problems: to choose the optimal times, $\{T_i\}$, to adjust the state to the optimal level, z^*, to optimize its value:

$$v(z_0) = \min_{\{T_i\}} E_0\left\{ \int_0^{T_1} e^{-rt} B(z(t))^2 dt + \sum_i \left[-e^{-rt} m + \int_{T_i}^{T_{i+1}} e^{-rT_i} B(z(t))^2 dt \right] \right\}. \tag{3}$$

The optimal solution to this problem is to establish an "inactivity band" for $z(t)$, $(-\overline{z}, \overline{z})$: whenever $z(t)$ lies within the inactivity band the firm makes no adjustment and whenever $z(t)$ reaches or crosses these bounds, the firm pays a real fixed cost m to adjust $z(t)$ to z^*. Importantly, this optimal decision rule is unrelated to calendar

[13] Having a quadratic approximation to the profit function is a common assumption in both the menu cost and observation cost literature. For menu cost models see the seminal work by Dixit (1991) and Caplin and Leahy (1997). Stokey (2008) reviews the literature. Caballero (1989) and Moscarini (2004) make the same assumption in models with costly observation.

[14] In particular, the combination of quadratic loss and zero drift implies that the optimal Ss bands will be symmetric. This has the advantage of providing subjects with the simplest possible setting to play in the lab. Also, as shown in Proposition 1 of Alvarez, Lippi, and Paciello (2011) $\mu = 0$ implies that there will be no planned price changes between observations in the model with both menu and observation costs, which maps more closely to our experimental setting where subjects do not have the option to create price plans.

time: the timing of switches are entirely dictated by the current state of the world, $z(t)$. The optimal rule is thus purely "state dependent."

Inside the inaction region the value function is characterized by the Hamilton-Jacobi-Bellman equation, and outside the inaction region it is equal to the value at the optimal return point net of the menu cost. Thus, the value of the firm over the entire state space is defined by:

$$rv(z) = B(z(t))^2 - \mu v'(z) + \frac{1}{2}\sigma^2 v''(z), \qquad z \in (-\overline{z}, \overline{z}) \tag{4}$$

$$v(z) = v(z^*) + m, \qquad z \notin (-\overline{z}, \overline{z}). \tag{5}$$

Using these value functions, we calculate the optimal bands using standard methods that exploit value matching and smooth pasting conditions (full details are provided in the online Appendix).

Our experimental design, described below, will leverage the comparative static prediction of the model that an increase in volatility (an increase in σ) widens the optimal inaction band $(-\overline{z}, \overline{z})$. See Stokey (2008), for a discussion of the comparative statics of the model.

B. *State Dependence*, *Time Dependence, and Bounded Rationality*

What alternative decision rules might agents use in the setting of the model described in the previous section? Consider the following simple descriptive framework: agents set a time T from their previous adjustment at which to adjust the state and can condition this T on the current state z if they choose.

The optimal solution described in the previous section has a sharp prediction in this framework: agents will set $T = \inf\{t : |z(t)| = \overline{z}\}$, adjusting the first time the state variable hits the optimal thresholds. This is a strongly state dependent prediction as adjustment is entirely governed by (and is perfectly responsive to) the current value of $z(t)$. Panel A of Figure 2 plots a histogram of the distribution of switching points (i.e., sizes of price changes) from a simulation populated by purely state dependent agents. The simulation is parameterized by the vector $(\mu, \sigma, r, m, B) = (0, 0.24, 0.008, 25, 4.17)$ (the same parameters we use in our BASE treatment in the experiment). Solid vertical lines mark the boundaries $\overline{z}$ and $-\overline{z}$. The observable pattern generated by this decision rule is clear: we observe adjustments near $\overline{z}$ and $-\ \overline{z}$ but very little adjustment away from these boundaries.

By contrast, purely time dependent models predict that subjects will adjust based on calendar time rather than state. The Taylor (1980) model in particular predicts that agents will simply set $T = \tau$ (for some fixed τ), ignoring the state, z, altogether. Panel B of Figure 2 shows a histogram of price change sizes from a simulation using $\tau = 15$. Unlike the state dependent behavior in panel A, the time dependent distribution of price changes are strongly unimodal and quite diffuse, with adjustments occurring at a wide range of relative prices. In the setting described in the previous section, such time dependent rules constitute a substantial deviation from profit maximization.

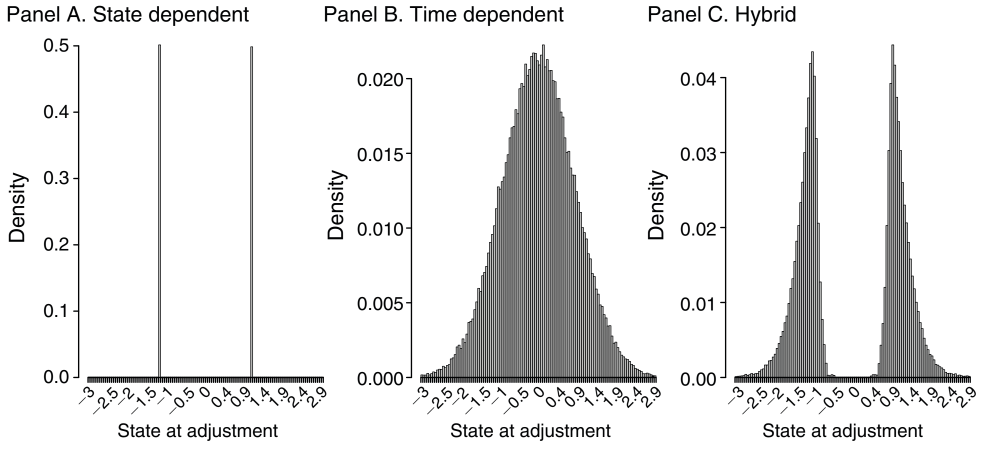

Figure 2. Histogram of Simulated Adjustments for Panel A: State Dependent Agents; Panel B: Time Dependent Agents; and Panel C: Costly Attention Agents

What could cause agents to make time dependent instead of optimal state dependent decisions? One hypothesis recently advanced in the literature is that subjects may experience private costs from implementing decision rules that require observation of or attention to the state. While a time dependent rule requires only limited attention to the passing of time, conditioning adjustment on the state can require considerable cognitive effort. If the additional cognitive effort necessary to condition adjustment on state is subjectively costly, agents may avoid using purely state dependent rules for economically sensible reasons.

Consider Alvarez, Lippi, and Paciello (2011),[15] who model a decision maker (DM) that experiences a fixed cost from observing (alternatively paying attention to, thinking about or noticing) the state z, denoted κ. In our setup, κ should be viewed as a mental cost paid by the subject associated with observing the state in contrast to m, which is a real cost set in the treatment and subtracted from profits. The DM's choice is then not only when to adjust the price, but also when to observe the state to gauge the adequacy of the current price. Upon observing the state and recording its value z, the DM decides whether to adjust the state and the time until the next observation $T(z)$.[16] Define $U(z)$ as the value of observing the state and choosing not to pay the additional cost to adjust the state at z, W as the value of observing the state and paying the menu cost to adjust the state to z^*, and $V(z)$ as the optimized value function between $U(z)$ and W that minimizes the individual's expected loss.

[15] Out of all of the many models of bounded rationality in Ss style problems we selected Alvarez, Lippi, and Paciello (2011) for its simplicity, tractability, clean comparative statics, fit to a laboratory setting, and the fact that it features purely state and time dependent models as special cases.

[16] This is the optimal policy in the case where there is no drift in the target price process ($\mu = 0$) that we implement in the lab; see proposition 1 in Alvarez, Lippi, and Paciello (2011). Under this solution there is no price adjustments between observation of the state. Such intermediate price adjustments may be optimal in high inflation environments.

Using this notation, we rewrite the $\mu = 0$ case of the model of Alvarez, Lippi, and Paciello (2011) as follows

$$(6) \qquad V(z) = \min\{U(z), W\}$$

$$(7) \qquad U(z) = \kappa + \min_T B\int_0^T e^{-rt}\left(z^2 + \sigma^2 t\right) dt + e^{-rT}\int_{-\infty}^{\infty} V\left(z - s\sigma\sqrt{T}\right) dN(s)$$

$$(8) \qquad W = m + \kappa + \min_T B\int_0^T e^{-rt}\left(z^{*2} + \sigma^2 t\right) dt + e^{-rT}\int_{-\infty}^{\infty} V\left(z^* - s\sigma\sqrt{T}\right) dN(s).$$

Whenever the DM observes the state z, she optimally chooses if she should adjust the state to z^* and also optimally sets the time until the next observation $T(z)$. An optimizing DM always adjusts z to $z^* = 0$ (because $\mu = 0$) and plans shorter inattention spells the further z is from z^* when she does not adjust (because $T(z)$ is maximized at $z = 0$). Whenever the DM is paying attention (i.e., is not in an *inattention spell*), she chooses to adjust if $W \leq U(z)$, generating two symmetric adjustment thresholds, $\tilde{z}$ and $-\tilde{z}$ (at which $W = U(z)$). As κ rises from zero, this inaction region narrows relative to the standard Ss region (i.e., $\tilde{z}$ falls below $\overline{z}$). Intuitively, when the state variable is sufficiently close to the Ss threshold, the option value of waiting to obtain more precise information on the state is outweighed by the observation cost: the DM adjusts at states both within and outside the original Ss inaction band because the DM does not continuously pay attention to the state.[17]

This model includes purely state and time dependent adjustment rules as special cases. When κ is zero, and observation of the state is costless, the DM continuously pays attention ($T \rightarrow 0$), and the optimal inaction region coincides with the Ss band in the basic menu cost model: the DM is capable of precisely adjusting as a function of state as visualized in Figure 2, panel A. On the other hand, in the limit when the ratio of the observation cost to the menu cost κ/m is infinite, the agent ignores the state, setting a fixed time interval between adjustments and implementing the Taylor rule, generating the sort of outcomes pictured in Figure 2, panel B.

At intermediate values of κ, agents display a hybrid pattern between these two extremes. Panel C of Figure 2 plots switching points from a simulation using intermediate values of κ. The results contain ingredients of both the state dependent behavior in panel A and the time dependent behavior in panel B. As in panel A, behavior is bimodal and peaked near optimal switching points, $\underline{z}$ and $\overline{z}$. As in panel B, adjustments are diffuse, with frequent adjustments at states far from the optimum. Importantly, this pattern of adjustment bears a strong resemblance to patterns documented in field data, as illustrated in Figure 1.

[17] Alvarez, Lippi, and Paciello (2011) derive highly tractable expressions for $T(z)$ and $\tilde{p}$ as an approximation of the optimal solution when the menu cost is small relative to the attention cost. They find the approximation to be accurate when the menu cost is less than half of the observation cost. We do not make special assumptions on the size of attention costs and therefore must solve the general model numerically. Most of the comparative statics from the expressions developed in Alvarez, Lippi, and Paciello (2011) extend to the general solution.

The purpose of our experiment is to compare the organizing power of these models, and one particularly useful tool for making this comparison is to evaluate comparative static predictions that differ across models. Our treatment design is centered on varying Brownian volatility, σ, which has distinct predicted effects on behavior under state and time dependence. Under the hypothesis of state dependence, increasing σ will cause average changes to become larger by widening the inaction bands (i.e., increasing $\overline{z}$). By contrast, under the hypothesis of time dependence, increasing σ will instead cause the *dispersion* of price changes to rise (i.e., the variance in the histogram in Figure 2, panel B will increase). Finally, under Alvarez, Lippi, and Paciello (2011) we expect to see both effects: average price changes will become larger *and* price changes will become more diffuse.[18]

II. Experimental Implementation and Treatment Design

In this section we describe our strategy for taking the model described in Section IA to the laboratory. In Section IIA, we show how we convert the continuous time, infinite horizon model into a closely related near-continuous time, indefinite horizon game. In Section IIB, we describe the visual software we built to test the theory and provide details regarding the implementation of the experiment. Finally, in Section IIC, we describe and motivate our experimental treatments.

A. *Binomial Conversion*

There are two barriers to directly implementing the model described in Section IA in the laboratory. First, true Brownian motion evolves in infinitesimal time increments and cannot therefore be generated in simulation. Second, the model has an infinite horizon and experimental sessions are obviously of finite duration.

To cope with the first problem, we turn to a well-known binomial approximation standardly used to simulate Brownian motion. We divide the game into a number of ticks, each of length Δt. In each tick, the series $p^*(t)$ iterates up with probability p and down with probability $1 - p$. If the tick moves up, it rises by h (the step size parameter) and, if down, falls by h.

As Δt approaches zero, the parameters converge to the two key Brownian parameters, drift

$$\alpha = \lim_{\Delta t \to 0} \frac{(2p - 1)h}{\Delta t} \tag{9}$$

[18] More precisely, when volatility increases, there are two predicted effects on the dispersion of adjustments in the attention cost model. The first effect is that agents will optimally set shorter durations between observations of the state. This is because the cost of observation is the same, but the benefits of observing the state are higher, since the state moves away from its optimal value at a faster rate. On the other hand, the direct effect is that for a fixed time between observations the dispersion of the value of z when the DM attends to the state increases, as the standard deviation of $z(T)$ conditional on some prior value of z is given by $\sigma\sqrt{T}$. As often happens in this type of problem, the direct effect dominates the indirect effect on the decision rule, and hence the standard deviation of adjustments is increasing in σ. We thank an anonymous referee for helpful comments on this point.

and volatility

$$(10) \qquad \sigma^2 = \lim_{\Delta t \to 0} \frac{4p(1-p)h^2}{\Delta t}.$$

Thus, by setting a very low value of Δt (one-fifth of a second in our experiment), we can generate random series that are nearly continuous and closely mimic the Brownian motion described in the theory. See Dixit (1993) for details.

To deal with the second problem, we implement a per-tick probability q of the game ending in lieu of a discount rate. As shown in Oprea, Friedman, and Anderson (2009), this converts to the Brownian discount rate in the limit according to the following formula:

$$(11) \qquad r = \frac{-\ln(1-q)}{\Delta t}.$$

B. *Implementation*

We ran the experiment using a custom piece of software programmed in a new Javascript environment called Redwood. The subject display, shown in Figure 3, consists of three panels, each visualizing a different part of the decision problem.

First, on the top panel, we show subjects their current price, $p(t)$ and the optimal price $p^*(t)$ (labeled the "Ideal Price" on the screen). Subjects see $p(t)$ as a stable red line and $p^*(t)$ as a point fluctuating over time with previous values drawn as a trailing blue line.

The bottom two panels display detailed information on the real time earnings consequences of subjects' decisions. The middle panel shows subjects their flow profits.[19] Positive flows are shown as regions shaded in green, negative flows as regions shaded in red. The bottom line charts the total profits accumulated so far.

A subject's only decision is when to press a button labeled "Move Current Price." When this button (or the space bar) is pressed, the red line immediately jumps to join the current ideal price, maximizing instantaneous earnings, and a cost of m is incurred. The experiment is divided into 20 periods (each period is a complete play of the indefinitely lengthed decision problem described above), with the current period ending with a fixed probability, q, in each tick. When the period ends, $p(t)$ is automatically and costlessly returned to $p^*(t)$ and a new period begins.

Data was collected in the Experimental and Behavioral Economics Laboratory (EBEL) laboratory at the University of California, Santa Barbara in May of 2014. A total of 74 subjects were drawn from an undergraduate subject pool using students

[19] Though we frame the problem above as an expected discounted loss minimization problem, we present subjects as an isomorphic expected discounted profit-maximization problem. In particular, we provide subjects with a flow constant C from which the instantaneous loss function is subtracted (similarly the menu cost is subtracted from current profits whenever an adjustment occurs). The optimal policy in either case is the same. In our discretized version of the problem, flow profits accruing to subjects in each tick is thus:

$$\pi(t) = \Delta t C - \Delta t B[p(t) - p^*(t)]^2,$$

where $C > 0$.

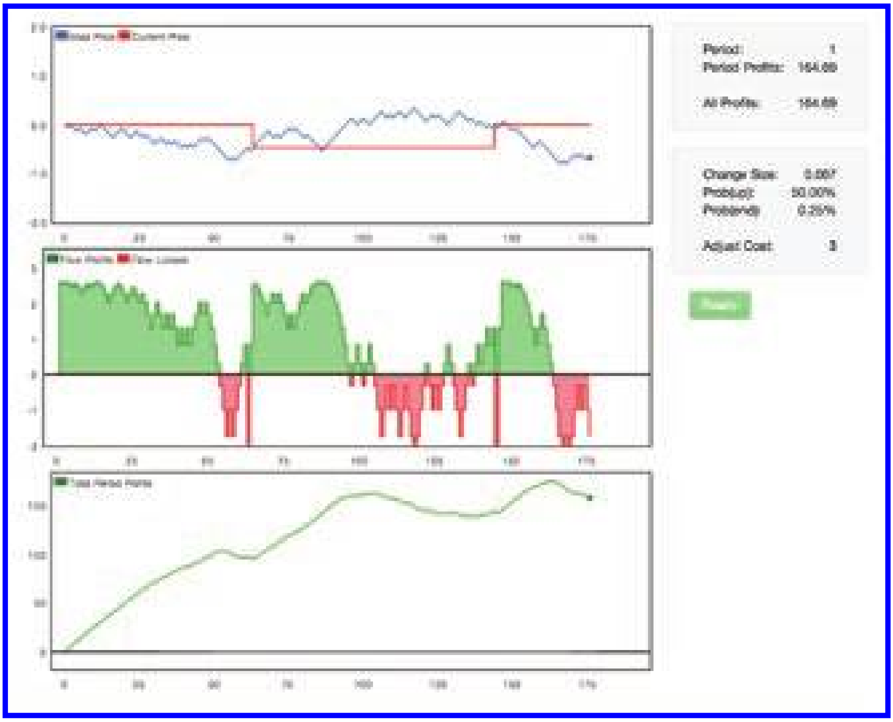

Figure 3. Screen Shot of Custom Experimental Software

from across the curriculum, recruited using the ORSEE software (Greiner 2004). Subjects were randomly assigned to visually isolated terminals and interacted with no other subjects during the session. Instructions, reproduced in the online Appendix, were read aloud prior to the beginning of the experiment. Subjects were shown the exact binomial parameters used in the experiment on their screen and were given a graph in their instructions showing how flow profits are impacted by deviations from z^*.

Each period's length was randomly predrawn prior to data collection such that all subjects experienced identical period lengths (though of course subjects did not know the duration of any given period until the period ended). Subjects were paid a $5 showup fee and 7.5 cents for each 100 points of profit earned. Sessions lasted roughly 1 hour and 30 minutes including instructions and average subject earnings were $19.43.

C. *Treatments and Motivation*

Table 1 summarizes our treatment design, displaying binomial parameters used,[20] other parameters used in the experiment, and both the thresholds of the optimal

[20] Our binomial parameters generate approximate Brownian parameters of $\alpha = 0$ (all treatments), $\sigma = 0.24$ (BASE/GUIDE), and $\sigma = 0.36$ (VOL). For all treatments, setting $q = 0.0016$ generates an approximate $r = 0008$.

Table 1—Parameter Values and Theoretical Predictions for Each Treatment

	Binomial			Other		Predictions	
Treatment	p	h	q	m	B	z^*	$\overline{z}$
BASE/GUIDE	0.5	0.107	0.0016	25	4.17	0	1.20
VOL	0.5	0.161	0.0016	25	4.17	0	1.47

Note: In all treatments, the time step in seconds is $\Delta t = 0.2$, and the profit-scaling constant is $C = 10$.

inactivity bands $(-\overline{z}, \overline{z})$ and the optimal price ratio, z^* for each treatment. Our design focuses on two main treatments (BASE and VOL) and one diagnostic treatment (GUIDE). Each subject was observed under only one treatment (we employed a completely between-subjects design). The treatments were parameterized to achieve several inferential goals.

First, in order to simplify the model predictions and the actual task we focused the design on a noninflationary setting. The probability that relative prices rise and fall were set to $p = 0.5$ generating a drift in the price index of $\alpha = 0$.

Second, we carefully selected our parameters to minimize flat payoffs around the optimum, a perennial concern in implementing dynamic stochastic problems in the lab. We considered dozens of parameterizations before selecting the ones we, in the end, took to the lab to ensure that subjects face substantial costs by deviating from optimal decisions. Payoff hills from the parameterizations we selected are shown in the online Appendix.

Third, as we discuss at the end of Section IB, we focused our treatment design on the effect of volatility: in our high volatility "VOL" treatment we nearly double volatility relative to our baseline "BASE" treatment. While a state dependent subject will respond to this parameter change by making larger price changes, a time dependent will make more diffuse price changes. An agent with both ingredients in her decision making, ala Alvarez, Lippi, and Paciello (2011), will make both larger and more diffuse price changes. We can thus use variation in volatility as a tool for evaluating time and state dependence in subject decision making.

Fourth, in order to better understand the nature of bounded rationality operating in our experiment, we replicated the BASE treatment but instead of having subjects glean optimal behavior themselves we explicitly tell them that they can maximize their expected earnings by implementing a threshold rule; and we show them on their screen the exact optimal threshold. In particular, we added "guidelines" around the current price (the stable line) on subjects' screen showing when $p^*(t)$ (the jagged line) deviates from $p(t)$ enough to optimally prescribe a price change. Doing this allows us to distinguish between two sources of bounded rationality. If deviations from state dependent predictions are substantially reduced in the GUIDE treatment relative to the BASE treatment, it is likely that cognitive difficulties in fully assessing the optimum are an important driver of time dependence. On the other hand, if BASE and GUIDE behaviors are similar, it is likely that such deviations are caused by attention or observation costs rather than costs associated with calculating the optimum.

III. Results

We present the results in three parts. Section IIIA presents an aggregate overview of the data. We show that while average decisions are meaningfully organized by and responsive to the underlying economic environment (suggesting state dependence), subjects do not choose sharp switching points or set true inaction regions (suggesting time dependence). Section IIIB reports reduced form estimates of the adjustment hazard function and confirms that there is significant time dependence in our data. The estimates also confirm that this time dependence is substantially higher in the VOL than in the BASE treatment but is completely unaffected by the information provided in the GUIDE treatment. Finally, Section IIIC digs deeper into the roots of time dependence by estimating a structural model of bounded rationality, discussed in Section IB. Analysis of our estimates suggests that volatility has a larger effect on dispersions in adjustment than is predicted by Alvarez, Lippi, and Paciello (2011).

A. *Aggregate Results*

Figure 4 plots histograms of the state, z, at the time of adjustment for each treatment.[21] For reference we plot vertical dashed gray lines at the bounds of the standard Ss inaction band, $-\overline{z}$ and $\overline{z}$, for each treatment.[22] We make five key observations about the data via patterns observable in these graphs.

First, it is clear that subjects fail in the aggregate to implement threshold adjustment rules along the lines prescribed by the purely state dependent model. States at adjustment (price changes) are instead highly diffuse.

Second, even though behavior is poorly described by a purely state dependent rule, it is also poorly described by a purely time dependent model. As in the field, decisions are, instead, bimodal, showing evidence consistent with both time and state dependence. We collect these observations as a first empirical result:

RESULT 1: *Adjustment distributions are bimodal but diffuse*, *with many switches occurring far from modal levels*, *closely resembling patterns documented in the field.*

A third pattern that is clear in Figure 4 is that modal adjustments in each treatment peak near optimal threshold levels in all treatments. In fact in no treatment can we reject the hypothesis that subjects' median adjustment times are equal to those predicted for their respective treatment. However, BASE/GUIDE adjustment states are both significantly lower than 1.47 (the VOL Ss threshold), while VOL adjustment

[21] These are weighted histograms with weights assigned to subjects so that each individual in the dataset has equal influence over densities. In particular, each decision contributes not 1 to the relevant bin in the histogram (as in a standard histogram) but instead $\frac{1}{N_i}$, where N_i is subject i's total number of adjustments over the course of the experiment. In this and other calculations reported in this subsection, we eliminate adjustments made within 0.2 seconds (one tick) of the previous adjustment because these can easily be generated accidentally by "double clicking" when using the software. Restoring these adjustments has no effect on the statistical results reported in this subsection.

[22] There is no evidence of learning in the data—first and second half behavior differ only in minor ways. Throughout the analysis we therefore use the full dataset in forming plots, estimates, and tests.

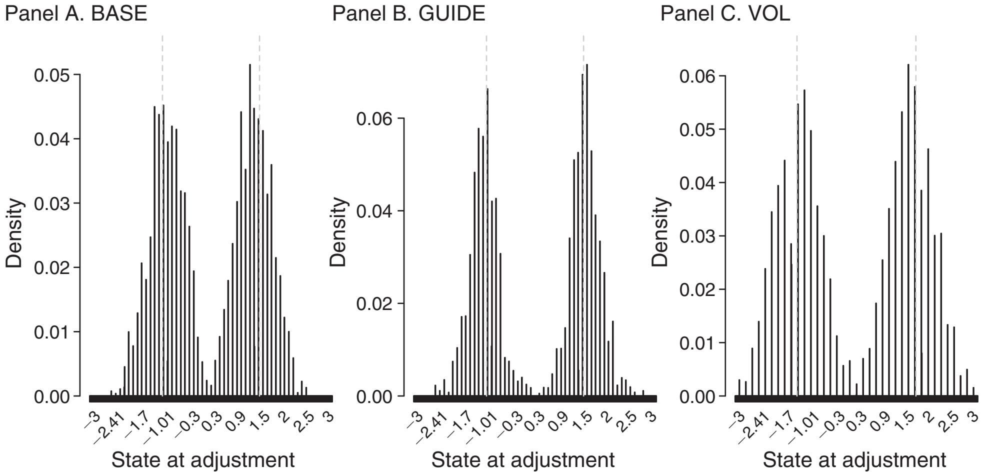

Figure 4. Weighted Histograms of Adjustment States for Each Treatment

states are significantly higher than 1.2 (the BASE/GUIDE Ss threshold).[23] Average behavior is thus surprisingly well organized by the adjustment predictions of the model and shows little sign of bias relative to those predictions.[24]

RESULT 2: *Optimal thresholds do a reasonably good job of organizing average observed switching states in all treatments.*

A fourth important observation is that explicitly showing subjects the optimal threshold in the GUIDE treatment doesn't seem to generate radically more state-dependent behavior than in the BASE treatment: BASE and GUIDE standard deviations of adjustment are virtually identical across treatments (0.29 versus 0.32) for the median subject; and the median subject earns virtually identical amounts on average across treatments (1.51 versus 1.539 points in the average tick). One implication of this striking similarity is that state dependence is likely an outgrowth of attention or concentration costs rather than calculative optimization costs.

A fifth and final pattern evident from Figure 4 is that increasing volatility, σ, as we do in the VOL treatment substantially increases dispersion in adjustment states relative to the BASE or GUIDE treatment. Indeed the standard deviation of adjustment states for the median subject is 0.46 compared to 0.29 in BASE and 0.32 in GUIDE. This increase is a key comparative static predicted by Alvarez, Lippi, and Paciello (2011) as discussed in Section IB.

[23] These tests are conducted using Wilcoxon tests applied to distributions of subject-wise median (absolute) adjustment states. All failures to reject are at $p > 0.2$, while all null rejections are significant at $p < 0.01$.

[24] Subjects' adjustment points are centered around the optimal thresholds from the very early rounds of each session. Comparing the first half of the session and the second half, we found no systematic difference in median adjustments. Adjustments in late rounds tend to have slightly lower dispersion (at the individual level), but this effect is small.

We collect these final two observations as a third result:

RESULT 3: *Increasing the volatility of the Brownian process* (*the VOL treatment*) *leads to a substantial increase in the dispersion of adjustment. However*, *giving subjects precise information on the optimal threshold rule* (*the GUIDE treatment*) *has no effect on behavior.*

B. *State Dependence*

Figure 4 provides clear evidence of aggregate dispersion in all treatments and larger dispersion in the VOL treatment relative to the others. However, the figure does not allow us to distinguish between two possible sources of dispersion: time dependence within subject or heterogeneity across subjects. Our motivating hypothesis is that dispersion is due to subjects failing to use consistent switching points (and failing to establish inaction regions) because of time dependent elements in their adjustment rules. However another possibility is that the result is a simple effect of aggregating subjects with heterogeneous inaction regions: subjects employing different switching points can generate the pattern observed in Figure 4 in the aggregate even if individual subjects employ perfectly state dependent rules.

In order to investigate whether our results are in fact consistent with individual level failures of state dependence, we estimate the parameters of the following simple, reduced-form hazard function for each subject in our sample:

$$h(\tilde{z};\alpha,\beta) = \frac{1}{1+\alpha\exp\left(-\frac{\tilde{z}}{\beta}\right)}, \tag{12}$$

where $\tilde{z} \equiv \left|\ln\frac{z}{z^*}\right|$. Note that because we directly observe the state, we are able to estimate the hazard as a function of the state (hazard functions are often estimated with respect to time in the literature). Here, α is a simple shifter of the hazard function (essentially an intercept term), while β measures the slope of the hazard function, a simple measure of responsiveness to state. Under the hypothesis of perfect state dependence, this parameter will be close to zero while a large β suggests insensitivity to state, suggesting time dependence. We estimate equation (12) by maximum likelihood for each subject independently and plot cumulative distribution functions (CDFs) of the critical parameter β for each treatment in panel A of Figure 5.

We make three observations. First, values of β are far from zero in each treatment for most subjects, suggesting that hazard functions are flat relative to a state dependent threshold rule for individual subjects. Most individual subjects, therefore, exhibit considerable time dependence in their decision making: the diffuse adjustment pattern observed in Figure 4 is not a mere artifact of aggregation of heterogeneous subject decisions. Second, there is absolutely no evidence of a change in the degree of time dependence when we add guidelines in the GUIDE treatment relative to the parametrically identical BASE treatment: the CDFs of estimates are virtually identical in each case. Finally, values of β tend to be far higher in the VOL treatment

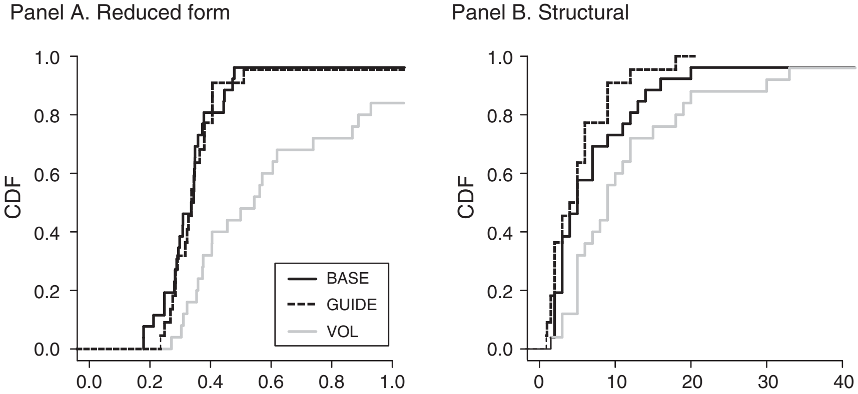

Figure 5. Estimates of Individual Level: Panel A: Switching Hazard Slopes (*Estimated with respect to state*); and Panel B: Structural Cognitive Costs

than in the BASE or GUIDE treatments. Indeed, the median VOL subject's β is over 50 percent larger than that of the median BASE subject and this difference is highly significant ($p < 0.001$, Wilcoxon test). Estimates thus confirm the pattern observed in Figure 4: volatility significantly increases time dependence in subjects' decision making.

RESULT 4: *Hazard functions of adjustment* (*estimated with respect to state*) *are gradually sloped across the sample in both treatments*, *suggesting significant time dependence in most subjects' decision making. This time dependence is unaffected by information on the optimal adjustment thresholds but increases substantially when we increase volatility in the underlying economy.*

C. *Structural Estimation*

To what degree can bounded rationality—subjective costs of observing or attending to state—account for the time dependence we measure in our data and the comparative static effects observed in our main treatment variation (an increase in Brownian volatility in the VOL treatment)? The distribution of switches in our data certainly resembles those generated by Alvarez, Lippi, and Paciello (2011). In order to better interpret our results, we structurally estimate their model for each subject in our dataset, focusing attention on the subjective cost parameter, κ.

We estimate the model by simulated method of moments. For each subject we choose the value of κ that minimizes the distance between two moments of the empirical distribution of the adjustment points, the median and standard deviation denoted by $\mathcal{M}$ and $\mathcal{S}$, and their theoretical counterparts obtained by simulation of the costly observation model.[25]

[25] Since Alvarez, Lippi, and Paciello (2011) does not have a closed form solution unless menu costs are substantially less than observation costs, we had to rely on a simulation-based estimation strategy. Among the other simulation-based methods, SMM is particularly useful for us because estimation is directly based on the same features of the data that we have discussed in previous sessions, namely the average size and the standard deviation of adjustments.

For a given observation cost κ, we solve the model numerically. The optimal policy of the costly observation model is described by the number $\tilde{z}$ defining the inaction region of the model and the function $T(z)$ that gives the planned inattention spell as a function of the state. We compute $\{\tilde{z}, T(z)\}$ by solving equations (6), (7), and (8) numerically on a tight grid of points in the state space.

With the policy $\{\tilde{z}, T(z)\}$ we simulate the model in the following way. When the DM observes the state he records the value z. There is an adjustment if $z \in [-\tilde{z}, \tilde{z}]$, in which case z is reset to 0. Whether there is an adjustment or not, the length of the next inattention spell is set to $T(z)$. The next value of the state is drawn from a normal distribution with mean z and standard deviation $\sigma\sqrt{T(z)}$. We repeat these steps starting at $z = 0$. Note that in the experiment some inaction spells are censored by the exogenous end of the round. Rounds are drawn from the discrete counterpart of an exponential distribution with parameter q. Similarly, the simulation is divided into rounds with a random length, drawn from an exponential distribution with parameter q. When a round ends, z is reset to 0, the observation is censored and a new round begins. We simulate 10,000 (uncensored) adjustment observations.

For each value of the observation cost κ, we compute the predicted median $M(\kappa)$ and standard deviation $S(\kappa)$ of the absolute value of the state at adjustment. Let $\mathbf{M}_S(\kappa)$ denote the 2×1 vector that collects the two simulated moments. Similarly let $\mathbf{M}$ be the vector of empirical moments. Then we compute the distance between empirical and simulated moments as follows:

$$J(\kappa) = \big(\mathbf{M} - \mathbf{M}_S(\kappa)\big)'\mathbf{Q}\big(\mathbf{M} - \mathbf{M}_S(\kappa)\big). \tag{13}$$

The matrix $\mathbf{Q}$ is a 2×2 diagonal matrix, whose entries are equal to the inverse of N times the variances of the sample moments, where N is the number of observations in the actual sample. We estimate the variance of the two sample moments by bootstrapping. The estimated κ is the value that minimizes $J(\kappa)$.

Figure 5, panel B plots distributions of the individual observation cost estimates by treatment. The median value of κ across treatments and subjects is five—one-fifth of the menu cost m used in our experiment. This cost is moderate but it is large enough to generate strong deviations from purely state dependent behavior. Indeed, Table 2 compares observed medians and standard deviations to those simulated using the median κ estimated in each treatment and shows how well the model organizes behavior in each treatment. Median adjustment points in the BASE and GUIDE treatments are each only about one step in the stochastic process away from the simulated median adjustment and the VOL simulated median is even closer. Most importantly, however, the simulations almost perfectly capture dispersion in each treatment and especially the large change in dispersion between the BASE and VOL treatments.

Figure 5 also allows us to compare the estimated subjective cost, κ, across treatments. BASE and GUIDE parameters are identical, further supporting the conclusion that GUIDE has no real effect on behavior. κ estimates, however, nearly double in the VOL treatment from a median of five in the BASE/GUIDE treatments to a median of nine.

Table 2—Empirical and Simulated Moments

Treatment	M	S	Simulated M	Simulated S
BASE	1.07	0.30	1.16	0.32
VOL	1.45	0.46	1.40	0.47
GUIDE	1.28	0.32	1.16	0.31

Is the difference in κ generated in VOL large? How important is this difference in determining the model's ability to explain the effect of volatility we observe in the data? In order to find out, we consider to what degree the model's accuracy in predicting the effect of volatility on time dependence (in the S simulations reported in Table 2)[26] depends on the fact that κ estimates are allowed to differ across treatment in that exercise. If we simulate allowing κ to vary by treatment (i.e., using the median κ parameter from each treatment) as we do in Table 2, we find that about 88 percent of the observed change in dispersion is predicted by the simulation. If we constrain κ to instead be seven (the median estimate across BASE and VOL) in each simulation, the simulation captures only 52 percent of the observed effect of volatility. (By contrast, conducting the same exercises on M reveals no difference in explanatory power using constrained versus unconstrained κ estimates.)

We conclude that though the model predicts the sign of the effect of volatility on time dependence in our data, it somewhat understates the size of this effect. Indeed, in order to explain the magnitude of the effect using the model, we must allow the parameter κ itself to be influenced by volatility (as we do by allowing κ to vary by treatment in the simulations in Table 2). One straightforward interpretation is that volatile environments have unmodeled direct effects on optimization costs. For instance, the marginal cost of either observing or processing information may depend in a direct way on the volatility of the environment – an idea with some intuitive psychological plausibility. We consider this and alternative reasons for the model's underestimate of the effect of volatility in Section IV.

RESULT 5: *Structural estimates suggest that subjects suffer modest optimization costs that can explain the dispersion observed in our data. The results further suggest that these costs are influenced by volatility.*

IV. Discussion

As Caplin and Leahy (2010) point out, Ss models have become a "touchstone" of economics, used to explain phenomena ranging from investment dynamics to sticky pricing, hiring decisions to inventory problems. However the predictive and explanatory power of these models is difficult to fully evaluate using conventional empirics. Perhaps the best developed empirical Ss literature so far focuses on price

[26] We ran a group of subjects through another robustness treatment, VOL', with very different optimal inaction bands but identical volatility parameters. We observed virtually identical adjustment dispersion in this setting, suggesting that the large effect of volatility measured in our data may be stable across environments.

adjustment in the face of menu costs (e.g., Klenow and Kryvtsov 2008, Nakamura and Steinsson 2008, and Midrigan 2010) and findings so far provide a mixed picture of the descriptive accuracy of simple Ss state dependent predictions. On the one hand, price adjustments tend to be bimodal and responsive to shocks as state dependent models predict. On the other, adjustments are not tightly clustered around modal adjustment points, with many changes too small or large to fit the stark predictions of the basic Ss model. Field data instead resemble a hybrid of state dependent adjustment of the sort described by Ss models and time dependent adjustment of the sort described by the famous Taylor (1980) rule.

We report the results of a laboratory experiment designed to not only help interpret findings from this active branch of the empirical literature but also to dig deeper into the fundamental relationship between behavior and Ss theory than has been possible using naturally occurring data. In our experiment, subjects repeatedly make adjustment decisions in a precise implementation of a standard menu cost pricing problem (Stokey 2008). By controlling the parameters and form of the experimental task, we can directly evaluate Ss models and significantly narrow the set of available explanations for observed deviations from their predictions.

Price adjustments in our data are highly bimodal and diffuse, strongly resembling patterns documented in the field; we are able to "grow" data in the lab with the distinctive characteristics of naturally occurring data. Because we know the precise optima of the laboratory task ex ante, we can go further, testing for the first time whether average switching points below and above optimal prices are biased relative to Ss threshold predictions. We show that in fact adjustments are statistically indistinguishable from optimal Ss thresholds *on average* and obey comparative static predictions: exogenously increasing volatility across treatments widens the range of average adjustments below and above the optimum to almost exactly the degree prescribed by Ss models.

Though it is clear that Ss-like forces powerfully shape average outcomes, it is equally clear that subjects do not use stable threshold rules and do not establish true inaction regions. Indeed, subjects make numerous and often sizable errors, frequently adjusting far from average (and optimal) adjustment points. As in the field, this pattern suggests that decision rules are not tightly tied to the current state, contra Ss models. Unlike in the field, the set of explanations capable of accounting for this "time dependence" is quite small in the lab.

One avenue pursued in the literature for explaining these patterns in the field is to replace the setup of simple Ss models with more elaborate "second generation" models. However, our results suggest that such elaborations may not be necessary: strong evidence of time dependence emerges in our experiment even though we *know* subjects are acting in a simple, first generation Ss environment. An alternative, recently pursued in the literature (Alvarez, Lippi, and Paciello 2011; Woodford 2008; Costain and Nakov 2011a, 2011b; and Moscarini 2004), is to posit that agents experience cognitive or attention costs that can also generate time dependent patterns of adjustment. Such bounded rationality explanations seem to be the only remaining explanation for time dependence emerging in the sterile, controlled setting of our experiment. Though our results cannot, of course, rule out that second generation Ss models apply in some field environments, they do suggest that optimization costs are very real and are at least a plausible contributing factor to time dependent patterns of adjustment documented in the world.

To guide our design (and to tie our hands to some degree), we chose a specific model of bounded rationality—Alvarez, Lippi, and Paciello (2011)—ex ante to anchor our investigation prior to collecting this data. In this model, agents experience costs from observing (or equivalently evaluating) the state. Purely state dependent or time dependent decision rules are special cases of the model when observation costs are zero or infinite (respectively), but at intermediate costs, optimizing agents choose hybrid decision rules that have both state and time dependent features. For our purposes, the model has the virtue of providing a tractable framework that includes the purely state and time dependent benchmarks as special cases; and allows us to structurally estimate the magnitude of subjective costs that subjects experience while making choices in our experiment.

Results from our investigation provide support for the Alvarez, Lippi, and Paciello (2011) model on some key dimensions. First, a central, distinctive comparative static prediction of the model is borne out in the data: when we increase Brownian volatility in the VOL treatment, we observe an increase in the standard deviation of price adjustments and the distribution of adjustment hazard rates relative to the BASE treatment. This is a key implication of the model with both observation and menu costs that is absent in standard Ss models. Second, structural SMM estimates of the subjective costs that animate Alvarez, Lippi, and Paciello (2011) reveal moderate costs that nonetheless are large enough to generate data that looks much like ours. In particular, simulations based on these parameters fit key moments in the distribution of adjustments rather well at the treatment level.

On the other hand, the Alvarez, Lippi, and Paciello (2011) model does not fully explain the deviations observed in subject behavior. One discrepancy between its predictions and our results is that Alvarez, Lippi, and Paciello (2011) predicts fewer very small price changes than we observe given the magnitude of the estimated costs in our data. A more important challenge to the model is that individual level structural estimates suggest a significantly different distribution of subjective costs in the high volatility VOL treatment than in the low volatility BASE treatment. This difference in costs is important for the model to fully account for the size of the effect of volatility on time dependence: were we to constrain estimates to the average across treatments, we would predict somewhat smaller effects of volatility on time dependence than we observe.

What explanations are available for this pattern in the data? Perhaps the most straightforward candidate is that volatility has an unmodeled direct effect on the sorts of costs discussed in Alvarez, Lippi, and Paciello (2011) (i.e., on κ). For instance, high volatility state variables may simply be more difficult to pay attention to or think about than lower volatility variables. The idea that features of the environment—like volatility—might generate additional optimization costs by increasing the difficulty of information processing was recently raised by Moscarini (2004), who pointed out "a fundamental disconnection between the difficulty of the decision problem and the cost of observations" in the stylized fixed cost of attention models. Intriguingly for our purposes, Moscarini (2004) shows that in a standard information theoretic framework, à la Shannon (1948), the cost of processing information should be directly increasing with volatility. This is an appealing explanation for our results and one that should be further investigated in future work.[27]

Though caution must be taken in projecting results from stylized laboratory experiments onto the field, there are several reasons to suspect our main results are useful for interpreting naturally occurring behavior. First, the remarkable fact that subjects in the GUIDE treatment—in which subjects are explicitly shown the optimal thresholds throughout the session—do no better than subjects in the baseline treatment strongly suggests that lack of sophistication is not the main determinant of the time dependence we observe in our data. Our subjects are sophisticated enough that they cannot improve their decision making by being explicitly told how to optimize. Moreover, there is evidence that subject behavior is strongly shaped by and responsive to relevant theoretical forces. Second, the deviations from theoretical predictions we do observe—time dependent diffusion of adjustment points—are very similar to those observed in naturally occurring data. This match between field and experiment suggests that we are likely measuring phenomena that are not artifacts of the lab. Finally, our ability to interpret these deviations is uniquely enabled by our laboratory environment: precise information on parameters and structure in the decision making environment are generally unobservable in the field. Though laboratory experiments are unlikely to (and shouldn't) supplant traditional empirics, observability of key variables makes laboratory evidence a powerful complement to the ongoing conversation between theory and empirics in this literature.

Still, it is important to bear in mind that the magnitudes of observation costs faced by subjects in our experiment are likely to be different from those experienced by the average decision maker in the field. After all, our experiment was designed to minimize the costs of making optimal choices: we do not require subjects to calculate optimal price ratios or current values of the price index and all information is neatly presented to subjects using a highly intuitive graphical interface. That we observe strong evidence of bounded rationality even here is striking. In the field, the problem of choosing optimal adjustment times is likely much more difficult, as agents must gather and interpret noisy information from multiple sources and estimate optimal prices themselves. As a result, our estimates probably substantially understate the magnitude of the observation costs most decision makers actually face. On the other hand, incentives to optimize are considerably larger for decision makers in the field and firms have access to sophisticated pricing models, counterbalancing this cost difference to a great degree. Unless subjective costs and optimization incentives scale identically in the move from lab to field, the degree of time dependence generated by bounded rationality may be larger (or smaller) than that observed in our experiment. However, to the degree that bounded rationality is a key driver of time dependence in the field (as it certainly is in the lab), the considerable

[27] Models with varying costs may also explain the large number of very small price changes we observe relative to Alvarez, Lippi, and Paciello (2011). In many such information-theoretic models of inattention, the decision maker only obtains an imperfect signal of the state when updating occurs, whereas the decision maker in the Alvarez, Lippi, and Paciello (2011) model obtains full information once the attention cost is paid. Imperfect information about the state in turns allows for more potential mistakes in price adjustments and can result in more small price changes than in fixed attention cost models (see Woodford 2008, for example). An alternate approach is to assume that the fixed cost of adjustment is random as in Alvarez, Le Bihan, and Lippi (2013). In contrast, we find in our experimental setting that small adjustments still occur even though the adjustment cost is not small or random.

time dependence documented in recent work suggests that subjective costs are large enough relative to incentives to influence pricing behavior.

Though the precise magnitudes of our subjective cost estimates cannot be readily exported to the field, the economic forces shaping them likely can. Perhaps our most practically relevant finding is that the degree of state dependence evident in subjects' behavior is directly eroded by volatility. This result is consistent with subjects suffering a cost from paying attention to the state as in Alvarez, Lippi, and Paciello (2011), but our effect sizes suggest more complex effects of volatility may be at work. Regardless, our results suggest that the descriptive validity of the orthodox Ss models can depend on details of the economic environment. The idea that optimization can be strengthened or weakened by fundamentals of an economy has potentially important implications for the interpretation of theory—and especially comparative static results – in a variety of policy settings. To pick an example close to the setting of the model motivating our experiment, Vavra (2013) argues that increasing volatility should reduce the effectiveness of monetary policy. Results like ours suggest that increasing volatility may also decrease the degree of state dependence in decision making, thereby generating a countervailing improvement in the effectiveness of monetary policy. Similar effects of the economic environment on the subjective costs of optimizing may arise in other settings in which the efficacy of policy depends on the degree to which agents reliably optimize. Investigating the relationship between optimization and environment in such settings may be a fertile area for future research.

REFERENCES

Abel, Andrew B., and Janice C. Eberly. 1994. "A Unified Model of Investment under Uncertainty." *American Economic Review* 84 (5): 1369–84.

Alvarez, Fernando E., and Avinash Dixit. 2013. "A Real Options Perspective on the Future of the Euro." https://www.princeton.edu/~dixitak/home/EuroCollapse.pdf.

Alvarez, Fernando, Luigi Guiso, and Francesco Lippi. 2012. "Durable Consumption and Asset Management with Transaction and Observation Costs." *American Economic Review* 102 (5): 2272–2300.

Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi. 2013. "Small and Large Price Changes and the Propagation of Monetary Shocks." Einaudi Institute for Economics and Finance (EIEF) Working Paper Series 1318.

Alvarez, Fernando E., and Francesco Lippi. 2012. "Price Setting with Menu Cost for Multi-Product Firms." National Bureau of Economic Research (NBER) Working Paper 17923.

Alvarez, Fernando E., Francesco Lippi, and Luigi Paciello. 2011. "Optimal Price Setting with Observation and Menu Costs." *Quarterly Journal of Economics* 126 (4): 1909–60.

Alvarez, Fernando, Francesco Lippi, and Luigi Paciello. 2012. "Monetary Shocks in a Model with Inattentive Producers." http://home.uchicago.edu/~falvare/ALP_Inattentive_producers_nov.pdf.

Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. 1951. "Optimal Inventory Policy." *Econometrica* 19 (3): 205–72.

Barro, Robert J. 1972. "A Theory of Monopolistic Price Adjustment." *Review of Economic Studies* 39 (1): 17–26.

Bertola, Giuseppe, and Ricardo J. Caballero. 1990. "Kinked Adjustment Costs and Aggregate Dynamics." In *NBER Macroeconomics Annual 1990*, Vol. 5, edited by Olivier Jean Blanchard and Stanley Fischer, 237–96. Cambridge: MIT Press.

Bloom, Nick, Stephen Bond, and John Van Reenen. 2007. "Uncertainty and Investment Dynamics." *Review of Economic Studies* 74 (2): 391–415.

Caballero, Ricardo J. 1989. "Time Dependent Rules, Aggregate Stickiness and Information Externalities." Columbia University Department of Economics Discussion Paper 428.

Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. "Search and Satisficing." *American Economic Review* 101 (7): 2899–2922.

Caplin, Andrew, and John Leahy. 1997. "Aggregation and Optimization with State-Dependent Pricing." *Econometrica* 65 (3): 601–26.

Caplin, Andrew, and John Leahy. 2010. "Economic Theory and the World of Practice: A Celebration of the (*S, s*) Model." *Journal of Economic Perspectives* 24 (1): 183–202.

Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102 (4): 703–25.

Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature* 34 (2): 669–700.

Costain, James, and Anton Nakov. 2011a. "Distributional dynamics under smoothly state-dependent pricing." *Journal of Monetary Economics* 58 (6–8): 646–65.

Costain, James, and Anton Nakov. 2011b. "Price Adjustments in a General Model of State-Dependent Pricing." *Journal of Money, Credit and Banking* 43 (2–3): 385–406.

Davis, Douglas, and Oleg Korenok. 2011. "Nominal shocks in monopolistically competitive markets: An experiment." *Journal of Monetary Economics* 58 (6–8): 578–89.

Dixit, Avinash. 1991. "A simplified treatment of the theory of optimal regulation of Brownian motion." *Journal of Economic Dynamics and Control* 15 (4): 657–73.

Dixit, Avinash. 1993. *The Art of Smooth Pasting.* London: Routledge.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics* 114 (2): 655–90.

Duffy, John. 2008. "Macroeconomics: A Survey of Laboratory Research." http://www2.econ.iastate.edu/tesfatsi/LaboratoryMacroeconomics.JDuffy2008.pdf.

Gertler, Mark, and John Leahy. 2008. "A Phillips Curve with an Ss Foundation." *Journal of Political Economy* 116 (3): 533–72.

Greiner, B. 2004. "The Online Recruitment System ORSEE 2.0: A Guide for the Organization of Experiments in Economics." University of Cologne Department of Economics Working Paper 10.

Grossman, Sanford J., and Guy Laroque. 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica* 58 (1): 25–51.

Klenow, Peter J., and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (3): 863–904.

Klenow, Peter J., and Benjamin A. Malin. 2011. "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics*, Vol. 3A, edited by Benjamin M. Friedman and Michael Woodford, 231–84. Amsterdam: North-Holland.

Magnani, Jacopo, Aspen Gorry, and Ryan Oprea. 2016. "Time and State Dependence in an Ss Decision Experiment: Dataset." *American Economic Journal: Macroeconomics*. http://dx.doi.org/10.1257/mac.20130267.

Midrigan, Virgiliu. 2008. "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations." Unpublished.

Midrigan, Virgiliu. 2010. "Is Firm Pricing State or Time Dependent? Evidence from U.S. Manufacturing." *Review of Economics and Statistics* 92 (3): 643–56.

Midrigan, Virgiliu. 2011. "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations." *Econometrica* 79 (4): 1139–80.

Moscarini, Giuseppe. 2004. "Limited information capacity as a source of inertia." *Journal of Economic Dynamics and Control* 28 (10): 2003–35.

Nakamura, Emi, and Jon Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.

Noussair, Charles N., Damjan Pfajfar, and Janos Zsiros. 2013. "Pricing decisions in an experimental dynamic stochastic general equilibrium economy." Unpublished.

Oprea, Ryan, Daniel Friedman, and Steven T. Anderson. 2009. "Learning to Wait: A Laboratory Investigation." *Review of Economic Studies* 76 (3): 1103–24.

Pindyck, Robert S. 1988. "Irreversible Investment, Capacity Choice, and the Value of the Firm." *American Economic Review* 78 (5): 969–85.

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

Scarf, Herbert E. 1960. "The Optimality of (S, s) policies in the dynamic inventory problem." In *Mathematical Methods in the Social Sciences*, edited by K. Arrow, S. Karlin, and P. Suppes, 196–202. Stanford: Stanford University Press.

Shannon, C. E. 1948. "A Mathematical Theory of Communication." *Bell System Technology Journal* 27: 379–423, 623–56.

Sheshinski, Eytan, and Yoram Weiss. 1977. "Inflation and the Costs of Price Adjustment." *Review of Economic Studies* 44 (2): 287–303.

Sims, Christopher A. 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy* 49 (1): 317–56.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Sims, Christopher A. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review* 96 (2): 158–63.

Stokey, Nancy L. 2008. *The Economics of Inaction: Stochastic Control Models with Fixed Costs.* Princeton: Princeton University Press.

Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88 (1): 109–32.

Vavra, Joseph. 2013. "Inflation Dynamics and Time-Varying Uncertainty: New Evidence and an Ss Interpretation." National Bureau of Economic Research (NBER) Working Paper 19148.

Vissing-Jørgensen, Annette. 2002. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *Journal of Political Economy* 110 (4): 825–53.

Wilson, Bart J. 1998. "Menu costs and nominal price friction: An experimental examination." *Journal of Economic Behavior and Organization* 35 (3): 371–88.

Woodford, Michael. 2008. "Information-Constrained State-Dependent Pricing." National Bureau of Economic Research (NBER) Working Paper 14620.

This article has been cited by:

1. Matthew Embrey, Christian Seel, J. Philipp Reiss. 2024. Gambling in risk-taking contests: Experimental evidence. *Journal of Economic Behavior & Organization* **221**, 570-585. [Crossref]
2. MARCUS GIAMATTEI. 2022. Can Cold Turkey Reduce Inflation Inertia? Evidence on Disinflation and Level-k Thinking from a Laboratory Experiment. *Journal of Money, Credit and Banking* **54**:8, 2477-2517. [Crossref]
3. Jinrui Pan, Jason Shachat, Sijia Wei. 2022. Cognitive Stress and Learning Economic Order Quantity Inventory Management: An Experimental Investigation. *Decision Analysis* **19**:3, 229-254. [Crossref]
4. Curtis Kephart, David Munro. 2021. Market concentration and the dynamics of prices and markups. *SSRN Electronic Journal* **29**. . [Crossref]
5. Samuel N. Cohen, Timo Henckel, Gordon Douglas Menzies, Johannes Muhle-Karbe, Daniel John Zizzo. 2018. Switching Cost Models as Hypothesis Tests. *SSRN Electronic Journal* . [Crossref]
6. Fernando Alvarez, Francesco Lippi, Juan Passadore. 2017. Are State- and Time-Dependent Models Really Different?. *NBER Macroeconomics Annual* **31**:1, 379-457. [Crossref]
7. Timo Henckel, Peter G. Moffatt. 2017. Sticky Belief Adjustment: A Double Hurdle Model and Experimental Evidence. *SSRN Electronic Journal* . [Crossref]
8. Justin D. LeBlanc, Andrea Civelli, Cary Deck, Klajdi Bregu. 2016. State dependent price setting rules under implicit thresholds: An experiment. *Journal of Economic Dynamics and Control* **68**, 17-44. [Crossref]