

4-28-2016

State Dependent Price Setting Rules Under Implicit Thresholds: An Experiment

Justin Deloy LeBlanc
University of Arkansas

Andrea Civelli
University of Arkansas

Cary A. Deck
Chapman University, deck@chapman.edu

Klajdi Bregu
University of Arkansas

Follow this and additional works at: http://digitalcommons.chapman.edu/esi_pubs

 Part of the [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

LeBlanc, J. D., Civelli, A., Deck, C., & Bregu, K. (2016). State dependent price setting rules under implicit thresholds: An experiment. *Journal of Economic Dynamics and Control* 68: 17-44. <http://doi.org/10.1016/j.jedc.2016.04.004>

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Publications by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

State Dependent Price Setting Rules Under Implicit Thresholds: An Experiment

Comments

NOTICE: this is the author's version of a work that was accepted for publication in *Journal of Economic Dynamics and Control*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in *Journal of Economic Dynamics and Control*, volume 68, in 2016. DOI: [10.1016/j.jedc.2016.04.004](https://doi.org/10.1016/j.jedc.2016.04.004)

The Creative Commons license below applies only to this version of the article.

Copyright

Elsevier

State dependent price setting rules under implicit thresholds: An experiment

Justin Deloy LeBlanc, Andrea Civelli, Cary Deck, Klajdi Bregu



PII: S0165-1889(16)30051-3
DOI: <http://dx.doi.org/10.1016/j.jedc.2016.04.004>
Reference: DYNCON3291

To appear in: *Journal of Economic Dynamics and Control*

Received date: 11 August 2015
Revised date: 6 January 2016
Accepted date: 19 April 2016

Cite this article as: Justin Deloy LeBlanc, Andrea Civelli, Cary Deck and Klajdi Bregu, State dependent price setting rules under implicit thresholds: An experiment, *Journal of Economic Dynamics and Control* <http://dx.doi.org/10.1016/j.jedc.2016.04.004>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting galley proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

State Dependent Price Setting Rules Under Implicit Thresholds: An Experiment

Justin Deloy LeBlanc*
University of Arkansas

Andrea Civelli
University of Arkansas

Cary Deck
University of Arkansas
Economic Science Institute, Chapman University
University of Alaska Anchorage

Klajdi Bregu
Univeristy of Arkansas

Abstract

How firms make their pricing decisions is a fundamental question of macroeconomics. We use a laboratory experiment to examine individual choices in a price updating task that provide insight into how well state dependent models reflect behavior. We find that in general subjects behave as if they recognize the importance of a state dependent pricing strategy, but they are unable to ascertain this threshold with precision and they also exhibit a substantial degree of time dependence. As a result, they update prices too frequently, and perform statistically significantly fewer real effort profit-generating tasks than theoretically optimal under full state dependence, which results in statistically significantly lower profits as well.

Keywords: Laboratory Experiments, Price Rigidity, State Dependent Models, Dynamic Pricing

JEL Classification: C91, D03, D21, E30.

*University of Arkansas, Walton College of Business, Department of Economics, Business Building 402, Fayetteville, AR 72701. We would like to thank Gabriele Camera, Aspen Gorry, and Dongya Koh for very useful comments and suggestions.

1 Introduction

How firms make their pricing decisions is a fundamental question of macroeconomics. The degree of price rigidity is commonly considered the main factor affecting the transmission of monetary policy to the real economy and, not surprisingly, an extensive literature has explored this matter over the years. A growing body of empirical work has tried to quantify the degree of flexibility of prices. The most recent empirical evidence seems to unanimously support the conclusion that price adjustments occur at higher frequencies than initially believed (see, for instance, Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2007; Bils and Klenow, 2004, who find durations between four and eight months). At the same time, many competing theories have been proposed to provide a micro-foundation of price rigidity and embed it in macroeconomic models, from the more simple staggered price model of Calvo (1983) to the menu costs of Mankiw (1985) and the state dependent models (also known as “Ss”) introduced by Caplin and Leahy (1991). With so many alternative models still actively used in the theoretical work, economists seek to understand the intrinsic validity of the assumptions required by each, and controlled laboratory experiments provide a useful tool for doing so. In this paper, we investigate price rigidities relying on a fully experimental context. The purpose of our work is demonstrating to what extent a state dependent optimal pricing model has a legitimate *behavioral* basis.

Ss models hold at their core that agents should choose when to re-optimize solely based on the state of the world (“state-dependent” adjustment) and not on time *per se* (i.e., “time-dependent” adjustment). In our experiment, we follow the approach of Magnani et al. (2015) to setup an experimental environment in which the subjects must adopt a state dependent pricing rule in order to maximize the profits of their production activity. The profits earned by the subjects are deleteriously affected by the deviation of the price charged at a point in time from an optimal reference price, which depends on the underlying economic fundamentals. Prices can be reset to the optimal price at any time, but the subjects face an adjustment cost that determines the standard optimal threshold for price adjustments of the Ss models. We depart from Magnani et al. (2015)’s design in how we model the adjustment cost, which entails two fundamental differences with respect to their basic framework. First, in our experiment subjects must complete a real effort task in order to produce a unit of output and make profits. Second, the subjects must complete a time-consuming task in order to adjust the price they charge, thereby foregoing profits to do so. As a result of the introduction of these two tasks, the cost of updating in our model reflects a real trade-off between production and price setting, and it is not exogenously given to the subjects as in

Magnani et al. (2015). The trade-off relies on the perception of time cost of the subjects, which is an individual specific cost. By internalizing the updating cost and examining how this affects subjects' behavior, we are able to highlight a new dimension of the inability of price setters to accurately solve the problem facing them, which does not seem strictly related to economic uncertainty *per se*. This new dimension of the analysis provides additional insights on the role of attention and cognitive load in the price setting mechanism.

The main results of the experiment are essentially twofold. First, we find that subjects tend to respond in a way that is akin to state dependence. In each of the situations we examine, the average subject exhibits behavior consistent with the employment of a noisy decision rule for a threshold. They recognize the importance of an inaction region across treatments, although the precision of these regions is lower the higher the degree of uncertainty in the economy, and they exhibit some elements of time dependence in their pricing decisions. This first result is in line with the evidence documented by Magnani et al. (2015), who explain the large deviations of the subjects' decisions from the optimal thresholds with a model of bounded rationality. In their model, the limitations to rationality increase as a function of the cognitive load fundamentally caused by the volatility of the economic environment in which subjects find themselves operating. Second, subjects in our experiment make an additional systematic error in the identification of the optimal thresholds due to our implicit adjustment cost. Comparing the experimental results to the optimal thresholds obtained by simulation of the profit-maximizing decision rules that a fully rational agent should have followed, we find that subjects adjust their prices too "early" with respect to the state, in the sense that they reset prices when the profit from a unit of output sold is still higher than the optimal threshold point. Since this type of error does not seem to depend on the uncertainty of the environment, we link it to the form of the adjustment cost used in our experiment.

The results obtained from the experiment suggest that a state dependent model is insufficient to fully capture the price setting behavior. Rather, an hybrid model that mixes, at the same time, state and time dependent elements would be necessary to adequately fit the data from the laboratory. Similar conclusions were suggested by the analysis of field data as well; for instance, Nakamura and Steinsson (2007) and Klenow and Kryvtsov (2008) show that the shape of the hazard functions of price adjustments in field microdata is not directly attributable to simple menu-cost models. Also Magnani et al. (2015) conclude that bounded rationality is the most plausible explanation of this hybrid dynamics of pricing decisions given their experimental setup. Subjects experience cognitive costs in assessing the state

of the economy and formulating their optimal response, and higher cognitive loads decrease subjects ability to fully adopt state dependent rules by increasing these costs. While there are some exceptions, most studies have found negative effects of cognitive load on basic economic behavior (see Deck and Jahedi, 2015, for a recent survey).

We further explore this type of explanation with the introduction of the trade-off between production and price setting tasks. The purpose of this feature of our design is to have a simple and direct way to tamper with the cognitive capacity of the subjects; if the cognitive load explanation is valid, this additional burden should negatively affect the performance of subject across treatments. The implicit cost of adjustment achieves this goal in two ways: the first is by forcing subjects to infer their cost of resetting prices in terms of forgone profits while they accomplish the updating task; the second is by adding a real source of distraction for the subjects that diverts their attention from screening the state of the economy to executing the production task. While it is hard to precisely disentangle the two channels, they both operate in the same direction by increasing the likelihood of optimization errors. This implicit trade-off provides a simple micro-foundation of the attention allocation mechanism in the laboratory, but it can be considered fairly realistic for a large group of firms in which one or only very few workers are employed. The Small Business Administration (SBA), for instance, reports that over 75% of businesses in the US have no employees. Hence for most companies, pricing, production, and other tasks likely fall to a single person. Whilst experimentally adequate in general, this trade-off would be clearly much less realistic for larger and more sophisticated enterprises in which separate divisions are in charge of production and pricing decisions. However, other distracting factors would likely impact the pricing divisions of larger firms and hence our experimental set-up likely has relevance for them as well.

The effect of the endogenous cost of adjusting prices is a uniform shift of thresholds across decision environments. Sellers, pressured to actively engage in production, seem to systematically underestimate the actual cost they must bare to update prices, and as a consequence they end up updating relatively more often than optimal. This behavioral pattern can be explained in two ways. First, the literature on cognitive load (Deck and Jahedi, 2015) has demonstrated that higher cognitive load leads people to be more impatient; subjects update more often in the short period in the experiment. Second, this pattern is also suggestive of models, as for instance rational inattention, in which information processing limitations trigger decisions about attention allocation. Intuitively, subjects focus more attention on production which, given a limited processing capacity, leads them to acquire a worse esti-

mate (signal) of their implicit costs. Finally, our experimental results offer some guidance also in interpreting some of the empirical evidence on price changes. Specifically, two observations are worth mentioning. First, the flat hazard rates found for the experimental updating decisions suggest that models with constant probability of updating, such as the Calvo pricing, might still have some practical utility in approximating the observed price changes in spite of the oversimplified mechanism of price setting they embed. Klenow and Kryvtsov (2008) find that flat hazard rates are typical with field data as well. The second observation is that the large differences in price change frequencies across sectors documented by both Nakamura and Steinsson (2007) and Klenow and Kryvtsov (2008) could find at least a partial explanation in our price setting trade-off mechanism. Price durations in the field could reflect the interaction between the price setting activity and other phases of the production and commercialization of a good.

The rest of the paper is laid out as follows: Section 2 elaborates on the relation between this paper and the literature in detail. Section 3 describes a simple theoretical model that provides us with some qualitative predictions on the relative movements of observations from treatment to treatment. In Section 4, we discuss the experimental design and implementation, and in Section 5 we present and discuss the results. Finally, in Section 6 we offer some discussion of our findings and some ideas for potential future work.

2 Related Literature

Given the complexity of bringing macroeconomic models into a laboratory setting, it is not especially common to use experiments to examine these models. With regard to price rigidity specifically, there is a particular gap in the literature. To the best of our knowledge, the only paper currently examining these issues is Magnani et al. (2015), who devise an experiment to test a state dependent pricing model. They find that adjustment distributions in the laboratory closely resembled field data, and that mean switching behavior obeyed comparative static predictions, on average. In their paper, subjects face only a simple price adjustment task in which they must incur a fixed adjustment cost in order to reset the current price to the observable “optimal price.” We extend their framework by including real production and price updating tasks, which allows us to define an adjustment cost that reflects the implicit trade-off between the two tasks. This new mechanism negatively affects the decision process of the price setters by exposing the subjects to an increased cognitive load; in response to that, subjects systematically over-update.

The interpretation of our results is closely related to the literature on the negative effects of cognitive load on people's economic decisions. Kahneman (2002, 2011) develops a dual system model of behavior.¹ In this framework, people have both an impulsive system (System 1) and a deliberate reasoning system (System 2). When faced with a choice, System 1 quickly arrives at a choice, which System 2 can override if brought to bear. However, when a person is placed under cognitive load, the ability of System 2 to counteract System 1 is reduced. Psychologists have recognized that people have limited cognitive resources since at least Miller (1956), who found that people can only hold seven items effectively in short term memory. As a result, much of the work on the negative effects of cognitive load involves examining choices people make while memorizing a 7 digit number. Using a within-subject design, Deck and Jahedi (2015) demonstrate that those whose math performance is most impacted by cognitive load are also the people who become more impatient and more risk averse when under cognitive load. This pattern of behavior provides strong evidence in favor of the dual system model. Relatedly, studies have found that people with high cognitive ability are less risk averse and more patient (see, for example, Dohmen et al., 2010; Benjamin et al., 2013).² Thus, increasing cognitive load effectively serves to reduce the cognitive ability of the decision maker. In our setting, the heightened cognitive load could explain two of the behavioral patterns that we observe. First, because cognitive load leads people to be more impatient, they tend to update before they should. Second, because their cognitive ability is reduced, they have difficulty identifying a consistent threshold.

State dependent models are only one of the competing models of price rigidity currently in use. The evolution of these models dates back to the simplest mechanism of fixed contract durations, as proposed by Fischer (1977) and Taylor (1980). Both these papers develop rational expectation models wherein firms are locked into staggered contracts. Applying the model, Taylor finds that contracts of three or four quarters would be enough to explain unemployment rigidity. Staggered prices are introduced by Calvo (1983), which is considered the fountainhead of the New Keynesian approach to modeling price rigidity. Even though relatively more basic and stylized than a Ss model, this approach is still largely used in modern macroeconomic models due to the convenient tractability it offers. In his model, only a fixed fraction of firms are exogenously chosen to update their prices in each period, while the others are forced to keep the older price level. In a different approach, Mankiw (1985) suggests that prices are sticky because of “menu” costs. That is — real, explicit costs

¹Similar models have also been developed by Fudenberg and Levine (2006) and Mukherjee (2010).

²While cognitive load can be viewed as a state, cognitive ability is typically viewed as a trait.

incurred by producers to update their prices. A firm observes the deviation of its current price from the optimal price level, and it decides to update its price when the possible increase in revenues from adjusting is higher than the explicit updating cost. State dependent models can be considered a logical progression of these last two models. Our paper, similarly to Magnani et al. (2015), shows that price setting fundamentally follows a state dependent rule, even though it also exhibits traits attributable to time dependence and elements of price rigidities typical of the Calvo pricing.

A parallel strand of the same literature unfolded around the relation between prices and information in the economy. Hayek (1945) argues that prices are the means by which economic changes are communicated across the entire economy to its various participants through a market mechanism. As a consequence, more price flexibility would reflect improvements in the market communication system. Lucas (1973) proffers a model that derives a short-run Phillips Curve from the similar idea that prices convey noisy information about the stance of the economy and producers demand. More recently, stickiness in information propagation has been used by Mankiw and Reis (2002) to obtain sticky prices and inflation persistence in a way alternative to the New Keynesian Phillips Curve. Finally, Sims (2003) proposes a sophisticated model of price rigidities based on a limited information-processing capacity of the price-setting agents. This mechanism, known as rational inattention, improves on menu costs and Calvo/Ss models in that it can simultaneously generate the observed micro flexibility together with sufficient rigidity of prices on the macroeconomic level. The essence of the rational inattention model was laid out in Matejka (2010). In his model, firms face a large number of noisy signals about factors that affect pricing decisions. However, they face a constraint on the total amount of attention that they may allocate over these signals, so they must decide the signals on which they will focus. Even though we do not explicitly explore the information channels in this paper, we identify an additional set of errors in the pricing decisions which, as discussed above, can be interpreted in terms of bounded rationality and the effects of the reduction of subjects' attention on the optimal pricing due to the production activity.

Finally, our paper is related to the empirical literature that studies the characteristics of price adjustments from a more micro perspective at firm level. The main conclusion of this strand of literature is that prices are actually quite flexible and price adjustments occur relatively often, with a median frequency largely below one year. Blinder et al. (1998) surveys firm managers and concludes that prices change approximately only once per year; similar evidence is provided by Carlton (1986), Cecchetti (1986), Kashyap (1995), Levy et al.

(1997), MacDonald and Aaronson (2001), and Kackmeister (2001). This was a commonly held belief in the 1990's until Bils and Klenow (2004) studied frequency of price changes in 350 categories of goods and services representing approximately 70% of consumer spending and found that half of them had median frequency of price adjustments of approximately 4.3 months, substantially shorter than that reported by Blinder et al. (1998). Also Nakamura and Steinsson (2007) find that the typical price duration is below one year, around 7 – 8 months for median price adjustments, and that there were not gigantic differences across regular prices from the time periods of 1988-1997 as compared to 1988-2005. For these same periods, Klenow and Kryvtsov (2008) find qualitatively similar results for both posted prices (i.e., raw BLS data) and regular prices.³ Our experiment confirms that price setters understand and follow the right incentives to adjust prices relatively often, even in presence of hard costs of re-adjusting due to foregone production and profits. In this respect, the evidence we provide is broadly consistent with the conclusion that prices are indeed quite flexible as suggested by the most recent empirical literature.

3 Theoretical Framework

We first present a stylized theoretical model to help provide a structure for understanding the pricing decision faced by subjects in the experiment. As with Fischer (1977), Taylor (1980), Calvo (1983), and others, this model is an abstraction from the underlying decision making setting. That is, our research process began with the creation of an environment that mimics the phenomenon in the naturally occurring world that we wanted to investigate and then we constructed a model to forecast behavior. Thus, our experiment is not meant to be a strict test of this model. The model is a state contingent pricing model in discrete time in which a producer must make a decision on how to allocate time between the production of one unit of a single good and a task to optimize the price at which the good is sold. Time is the only input for both activities, and it takes one unit of time to produce one unit of the good. The price setting task is time consuming, too; and it is also assumed to take one unit of time.⁴ The producer, then, faces a trade off between earning positive profits from the current production or suspending production for one period to update the current price and

³Moreover, Klenow and Kryvtsov (2008) note that one of the primary issues that arise when assessing price rigidity is whether to focus on the median or the mean of the price change and they stress that it is also important to consider how one must account for “sales” (i.e., short-term reduction in prices, often used to lure customers into the store).

⁴There is no theoretical reason why the time cost of production and price setting need to be the same; but observed behavior in the experiment is consistent with this simplifying assumption.

enjoy higher expected unit profits in the next period. The profit function, Π_t , is specified in terms of the deviations of the optimal price, P_t^* , from the current price, P_t , as

$$\Pi_t = \exp[-\mu \cdot |X_t|]$$

where $X_t = P_t - P_t^*$ and $\mu > 0$ is the parameter that regulates the rate of decay of profits. Given this profit function, we can express the Bellman equation of the producer's recursive problem as

$$V(X_t) = \max_{u_t} \{(1 - u_t) \Pi_t + \beta \mathbb{E}_t (V[X_{t+1}(X_t, u_t, \varepsilon_{t+1})])\}$$

where u is a binary variable representing the decision of the producer to update the price ($u = 1$) or produce in the current period ($u = 0$), and ε is a zero mean, *i.i.d.* exogenous innovation to X . The optimal price level P_t^* drifts away from current price P_t . The price deviations are assumed to follow either a random walk process conditional to the decision of not updating or to be a simple white noise shock if the producer decides to update. Hence, the stochastic process of X_t can be expressed as

$$X_{t+1} = \begin{cases} X_t + \varepsilon_{t+1} & \text{if } u_t = 0 \\ \varepsilon_{t+1} & \text{if } u_t = 1 \end{cases}$$

The producer makes the production decision after observing the current X_t . If the price is updated, Π_{t+1} is known up to the $t + 1$ realization of the random innovation ε_{t+1} . Finally, future profits are discounted by the subjective discount factor $\beta < 1$.

There are two important differences between this theoretical framework and our implementation of the experimental design below. First, the two tasks will last prolonged periods and not simply one. This implies a longer lag between the time when a task is started and the time when it is completed, with an increase of uncertainty that affects the updating and production decisions of the producers. In the model, this feature is only in part captured by the fact that X_{t+1} after the price adjustment still has a stochastic component represented by ε_{t+1} . Second, it is possible that the time necessary for subjects to complete the two tasks may not be exactly the same and could vary over time. While the structure of the model is time invariant, the calibration of β allows for some flexibility in the transformation rate between the two tasks since β can be used to modify the relative value of updating even if the two tasks require the same input of one unit of time. Nevertheless, the model provides a very good guidance in calibrating the parameters of the experimental treatments and a set

of very useful comparative statics to benchmark our experimental results.

The model is solved numerically by value function iteration for a discretization of the state space. The associated solution of the policy function $u_t = m(X_t)$ provides the inaction regions of the producers' updating decisions (the set of values of X_t for which it is more profitable to produce rather than updating the current price). In addition to the profit decay parameter μ , the solution of the model depends on the uncertainty about the future economic circumstance that is embedded in the variability of ε . We model ε_t as a discrete random variable which takes on two possible values with equal probability:

$$\varepsilon_t = \begin{cases} h & \text{with } p = .5 \\ -h & \text{with } p = .5 \end{cases}$$

The step size h of the increases (decreases) of ε_t represents also the standard deviation of the innovation, and it can be thought of as a measure of the volatility of the underlying economic fundamentals on which X_t depends. We set $\beta = .98$, and we study four calibrations of μ and h that are going to match the four treatments in our experiment.⁵ We choose fast/slow decay and small/large step size. Intuitively, the inaction region gets larger when the decay is faster because a faster decay of profits causes a reduction of the future value of updating; higher uncertainty affects the inaction region in a similar way. Figure 1 illustrates the inaction regions for the four combinations of the parameters; the proposed values are sufficient to obtain a substantial variation across different combinations.

From the numerical solutions, we were also able to obtain qualitative (i.e., directional) predictions concerning the relative magnitudes of total profits earned. Figure 2 shows that we should anticipate subjects earning the most profit in the slow decay/small steps treatment, approximately the same profits in the intermediate treatments, and the least amount of profit in the fast decay/large steps treatment, the most volatile of the four. We were also able to obtain directional predictions on the relative fraction of tasks completed that would be price updates. Figure 3 shows what we would predict intuitively: That subjects would update relatively fewer times in the first treatment, approximately the same number of times in the intermediate treatments, and update prices most frequently in the last.

⁵We adopt a value of the discount factor that is fairly standard for macroeconomic applications. However, a direct mapping of β into the discount factor of the experimental subjects is not straightforward. Therefore, we check the robustness of the theoretical predictions of the model for a large set of alternative calibrations of β , included between 0.80 and 0.99. Smaller values of β shifts the thresholds in Figure 1 to the left by only a few percentage points, while the comparative statics of the results are completely preserved. Since we are mostly interested in comparisons across treatments, different values of β would be equivalent.

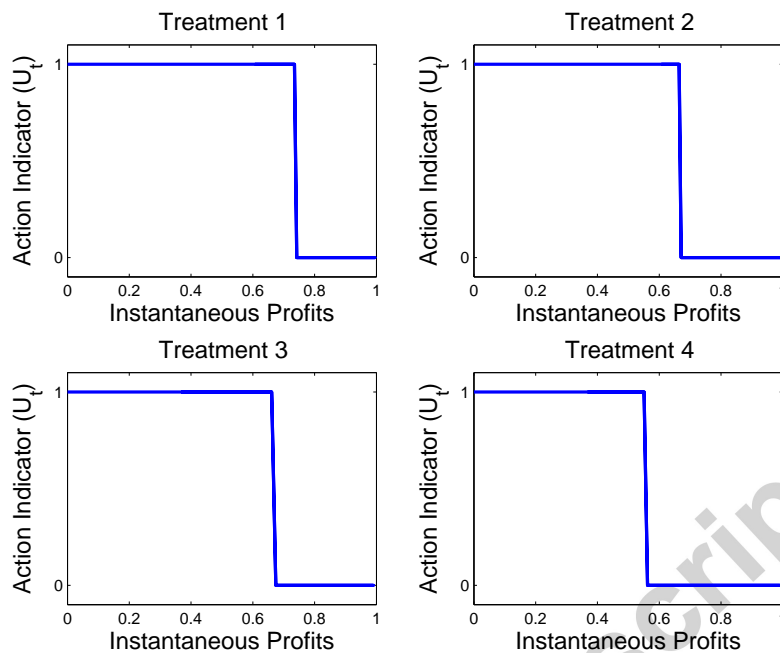


Figure 1: Inaction regions for the stylized theoretical model.




Figure 2: Theoretical relative total profits for the stylized model. Values standardized to profits in Treatment 1.




Figure 3: Theoretical share of price updates for the stylized model.

4 Treatment Design and Experimental Implementation

4.1 Experimental Framework

In order to examine the impact of implicit menu costs on price updating decisions, an experiment was devised to be administered in a laboratory setting. The experiment was developed using a combination of HTML5 and Javascript, and it was deployed using the Google Chrome internet browser.

The experiment put each subject in charge of a virtual pizzeria which sells its pizzas at a *Current Set Price*.⁶ The subject earns profits (denoted in Experimental Lab Dollars, or ELD) by fulfilling incoming orders by assembling pizzas with the correct ingredients.⁷ The upper left-hand corner of Figure 4 shows the pane into which an order comes. The five boxes

⁶If it is common in experiments to avoid using particular contexts that bias behavior, in our case, framing the task in the context of a business aids in helping the subjects comprehend the task at hand without either directing subjects as to how they should behave or providing them with an informative reference.

⁷1 ELD = 0.25 USD. The subjects were informed of the exchange rate before the experiment began.

at the top of the ingredient columns indicate the five ingredients for that “order.” These boxes are populated by randomly and independently selecting an ingredient from a list of ten possible (note that this means that any single ingredient could possibly occur multiple times on a single “order”). Below each of these boxes is a column of the ten possible ingredients. The ingredients in each of the columns are randomly arranged for each new order.

To fulfill an order the subject must use a computer mouse to select the box next to the ingredient in the column that matches each of the five boxes above it. Once a subject has selected ingredients, the subject then clicks the “Submit Order” button and earns some profit if all of the ingredients are correct.⁸ The subject earns profits determined by the difference between his *Current Set Price* and the *Optimal Price*, X_t . The problem faced by the subject is that the *Optimal Price* changes according to a random walk every 0.2 seconds, where with probability 0.5 the *Optimal Price* increases or decreases by step size h . Made available to the subject is a pane, shown in the lower left-hand Figure 4, where the subject may observe the realtime movement of the *Optimal Price* (the green line) and its relationship to his *Current Set Price* (the red line) enabling the subject to better understand the variability in a given treatment. This allows a visual representation of the X_t upon which each subject’s instantaneous profits depend. The instantaneous profit is also displayed on the screen, as is the “Cumulative Profit” the subject has earned that period.

The profit earned by the subject when the *Current Set Price* and the *Optimal Price* are equal (i.e., $X_t = 0$) is 1.00 ELD. Profit per pizza is calculated as $\Pi = \exp[-\mu \cdot |X_t|]$. If the subject determines that the per pizza profit that he is currently receiving is too low, then at any point that subject has the option to update his price—that is, set his *Current Set Price* equal to the *Optimal Price* at that moment in time. However, in order to do this, the subject must complete the task of entering an *Update Code* that is similar to the process of baking a pizza. The upper right-hand corner of Figure 4 shows the pane in which the *Update Code* is observed and entered by the subject. The top five boxes are populated randomly and independently with an integer in the set $[0, 9]$. The numbers in each of the columns below each of the five boxes are arranged randomly each time the “Update Price” button is clicked. Upon selecting the box next to each of the boxes next to the numbers that correctly correspond to the numbers in the five boxes, the *Current Set Price* is immediately set equal to the *Optimal Price* at that moment. To prevent a subject from selecting all of the ingredients in a pizza then updating his price and then submitting the already-prepared pizza, any time either the “Submit Order” or “Update Price” button was clicked, all of the

⁸For ease of exposition, we refer to this action as “baking a pizza” throughout the paper.

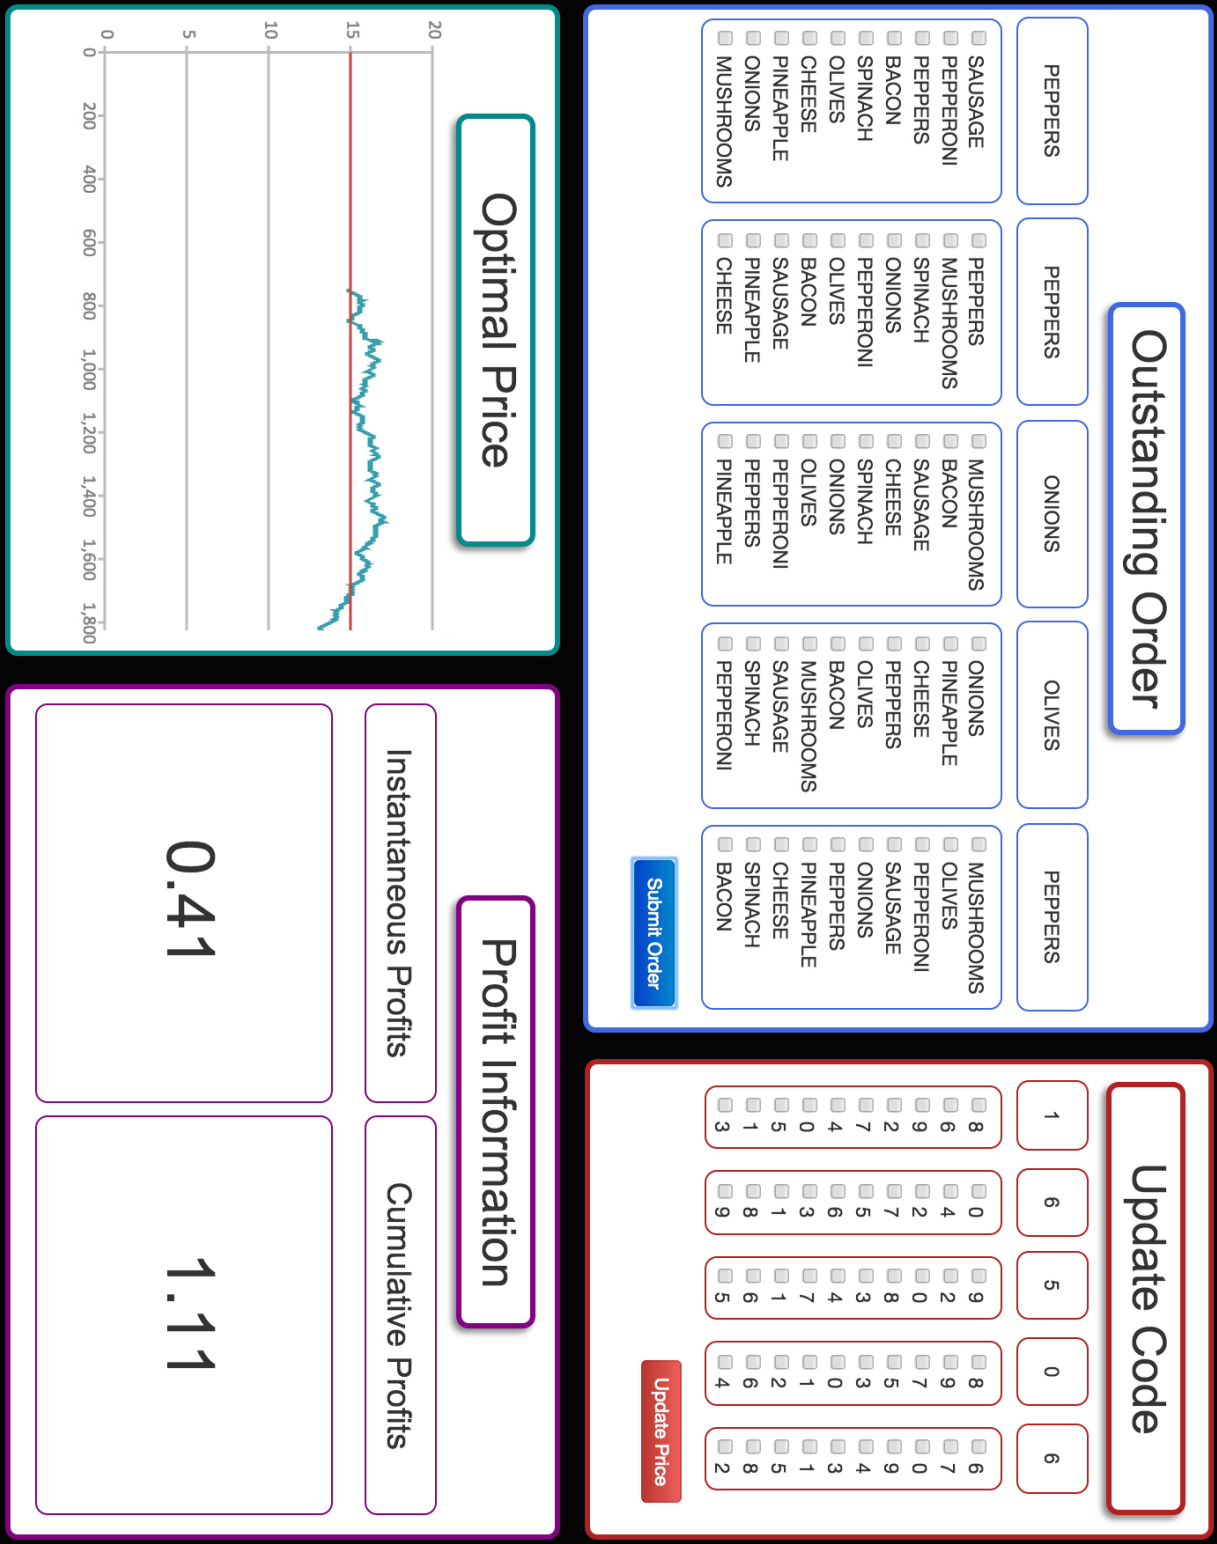


Figure 4: A screenshot showing the user interface with which the subjects interacted.

checkboxes in each of the columns for both tasks were cleared. That is to say, completing one task—either correctly or incorrectly—completely reset both tasks.

It should be noted that the decision node for the subject is made immediately after the completion of the previous task (except for the very first task, of course). Because of this, there is uncertainty in the profit that will be subsequently earned by baking a pizza. That is to say, the time necessary for the subjects to complete a task introduces a delay between the moment a task is started and when the task is concluded. After the beginning of the task, the *Current Set Price* can drift several steps away from the initial observed price, following a richer distribution over time than the simple Bernoulli that governs the price change at each discrete step of time. This is illustrated for a few steps in the probability tree in Figure 5 and the full distribution is represented for ten seconds from the beginning of a task by the two panels of Figure 6. The bottom panel illustrates the three dimensional distribution as a function of time, expressed in seconds, and the deviations of the *Optimal Price* from the initial price, measured by the number of discrete steps away from the price observed at the beginning of the task. The top panel of the figure shows a two dimensional aerial projection of the same distribution to aid in interpreting the three dimensional distribution. The distribution is symmetric around the initial price and, as time passes, it spreads out becoming more similar to a Gaussian distribution. However, it is important to notice that the distribution fundamentally corresponds to a random walk, and it remains centered around the initial price at any point in time, making the conditional expected value for the subjects independent of the time required to complete the task.⁹

Using the simulation data that we discuss in detail in Section 5.1, we can compute the average time that passes after someone updates their price until they face a decision at a particular instantaneous profit given they are following the optimal threshold. Figure 7 plots this relationship by treatment. One should keep in mind that time passes while engaging in baking so that the relevant instantaneous profit is the one they observed after, and not during, baking. Hence, one can encounter prices well below the instantaneous profit threshold if the previously encountered level was just above the threshold and the price continues to move further and further away from the optimal level while one is baking. The size of the markers in Figure 7 corresponds to the relative frequency of observing that instantaneous profit – higher profit levels are more commonly encountered. It is also worth noticing that these averages take into account that the probability distribution of the price movements

⁹The actual size of these steps depends then on the calibration of h in each treatment, but this does not have any bearing on the shape of the distribution in Figure 6.





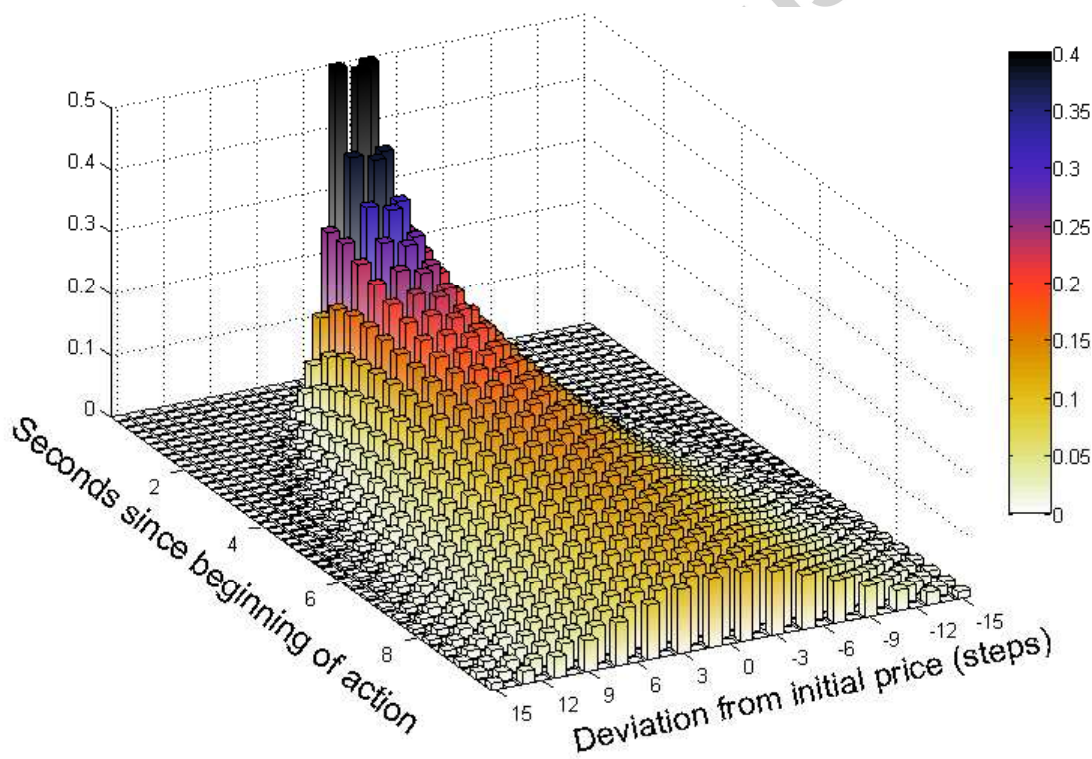
Figure 5: Probability tree of the Optimal Price deviations from the initial Current Set Price for four price changes.

experiences a truncation every time a subject crosses the optimal threshold. That is to say, starting at the optimal price, the instantaneous profit distribution has a precise theoretical form (derived from Figure 5) for any time horizon. However, the portion of this distribution actually spanned during a price spell will not always be the same because it depends on the duration of the spell itself.

As would be expected, a higher profit falloff rate and a larger price step size correspond to a faster profit decay and also a more rapid realization of the updating thresholds. The three horizontal red lines are the optimal thresholds in the four treatments. Treatment I has the highest threshold, but also a relatively smaller profit falloff than the other three treatments, which implies the largest average time to hit the threshold (about 80 seconds). We calibrate Treatments II and III to have nearly the same threshold and instantaneous profit time profile, while Treatment IV exhibits the lowest threshold and fastest time to reach it. Figure S13 in the online Appendix illustrates the full time distribution of profits in each treatment after a price update underlying the output in Figure 7. Consistent with this output, profit diffusion gets faster and falls deeper moving from the first to the last treatment.



(a) 2D Distribution



(b) 3D Distribution

Figure 6: Distribution of price deviations from the initial price observed when an updating task is started.




Figure 7: Average time at which a given level of instantaneous profit is expected to be observed by treatment. Simulation-based profit distributions, with ten seconds task duration. The size of the dots in each sequence of points indicates the relative frequency of realization of a profit level within each treatment. The horizontal lines indicate the optimal updating thresholds from simulation.

4.2 Experimental Treatments

The experiment was conducted in the Behavioral Business Research Laboratory at the University of Arkansas. We employed a 2×2 within-subject design, and treatments were ordered so as to mitigate order effects. Table 1 shows the four treatments and their associated parameter values, also relating them to the qualitative language used when discussing the stylized model of slow/fast decay and small/large step size. Further, in order to eliminate the possibility of backward induction by subjects, they were told that each treatment period had random termination with an average expected duration of approximately five minutes. In a given session, every subject faced each of the four treatments, but the order of the treatments was determined by random assignment.¹⁰

¹⁰Four termination times were determined in advance. For every subject, the first *period* lasted 350 seconds. The second, third, and fourth periods experienced lasted 380, 260, and 300 seconds respectively. This ensured that fatigue and experience with the interface were held constant.

Recall from Figure 1, that our theoretical predictions indicate that Treatment I—being the least volatile—should have the smallest inaction region. The most volatile treatment, Treatment IV, should have the largest, and Treatments II and III should be very similar and fall in between the other two. This is directly related to Figure 3, where we see that Treatment I should have the fewest number of price updates, Treatment 4 should have the most, and Treatments II and III being very close to one another and falling in between the other two treatments. Finally, we see in Figure 2 that we would expect subjects to earn the highest profits in Treatment I, the least profits in Treatment IV, and intermediate profits in Treatments II and III.

Treatment Parameters		
Profit Falloff Parameter (μ)	Step Size (h)	
	$h = 0.16$ (small step)	$h = 0.32$ (large step)
$\mu = 0.25$ (slow decay)	Treatment I	Treatment II
$\mu = 0.50$ (fast decay)	Treatment III	Treatment IV

Table 1: Profit falloff parameters and *Optimal Price* step size values for each of the four treatments.

4.3 Experimental Implementation

Before beginning the experiment, the subjects worked through a series of instructions explaining the mechanics of the tasks they were being asked to perform.¹¹ They were given the opportunity to read through these at their own pace and were also given the opportunity to ask questions for clarification. After reading through the instructions, the subjects were required to participate a practice round that lasted two minutes¹².

¹¹These instructions have been included in Appendix A.

¹²The profit falloff parameter and step size used in the practice round were different from those parameters used in the actual experimental treatments ($\mu = 0.375$ and $h = 0.32$). The reason for this was that each of the parameters were at exactly the midpoint of the two extrema of the parameters that the subjects would encounter in the experiment. The concern was that if the parameters were too drastic in either direction, it would condition the subjects to expect a certain level of volatility. Setting the practice parameter values directly in the middle of the actual experiment parameters gave them a good idea of a baseline, so to speak, and so what they experienced was somewhat of a perturbation from this initial exposure.

Upon completion of the practice round, the subjects were prompted to inform the experimenter. At this point, the subjects were required to complete a quiz¹³ to test their understanding of the experiment. Each quiz was evaluated by an experimenter, and for those questions that were answered incorrectly, the experimenter privately explained the correct answer to the subject. The experimenter then entered a pass-code enabling the experiment to commence. Prior to the beginning of each treatment, the subject was given a 60 second forewarning. During this time, the subject was presented with information about the treatment he was about to face. For instance, for Order A, prior to Period 1 (which corresponds to the parameters used in Treatment 1) beginning the screen displayed the following:

Your first period will begin in 60 seconds.

In this period the Optimal Price will move relatively slowly, and the profit fall off will be relatively small.¹⁴

Similar information was provided for 60 seconds before each period. At the 10 second mark the screen turned a bright red color and informed the subject that the next period began in 10 seconds.

Subjects were recruited from the undergraduate student body at the University of Arkansas, and a total of 70 subjects participated in the experiment over the course of 6 sessions in the month of June 2014. Subjects earned an average total profit, including a \$5 show-up fee, of \$18.04 USD. Each session lasted approximately 1 hour, which included working through the instructions, the practice round, taking the comprehension quiz, and participating in the actual experiment (i.e., all four treatments, each one lasting approximately 5 minutes). Payouts in excess of the show-up fee were determined by the subjects total cumulative profits across all four periods.

5 Results

In this section, we enumerate the main results of our experiment, and we discuss them in comparison to the quantitative predictions obtained from a benchmark simulation of the experiment. The results can be summarized in three points. First, although the subjects

¹³A copy of this quiz has been provided in the Appendix B.

¹⁴We adopted the language of “period” rather than treatment when communicating with the subjects. “Period” simply indicates the position in which the subjects encountered the various treatments. For every subject, there were four periods experienced chronologically; these periods, depending on the Order, could have different treatment parameters.

seem to recognize the state-dependent nature of their price updating decision, they typically identify smaller than optimal inaction regions. Second, subjects do not follow purely state-dependent rules and the updating decisions exhibit some degree of time dependence as well. Third, we observe a large heterogeneity in the thresholds across individuals. However, conditional on an individual threshold, subjects tend to update their *Current Set Price* too frequently relative to the profit-generating production tasks. For ease of presenting results, we will throughout this section present results corresponding to each treatment in the following form: (Result_I, Result_{II}, Result_{III}, Result_{IV}).

The predictions of the model are quite stark and it should not be surprising that people do not behave in strict accordance with those predictions. Rather the objective is to understand how people actually behave. Economic models are designed to be simplifications that capture core aspects of a decision problem and laboratory experiments provide a means of evaluating the predictive success of those models. In this, we are primarily concerned with the type of decision rules that people use to determine when they should update their prices and the behavioral responses to changes in the decision environment. We find that these aspects of the model predict the observed behavior reasonably well, although people's decisions do exhibit some time dependence. That people behave in a consistent manner and respond to environmental changes in an intuitive manner, we take as evidence that the subjects understand their tasks and act in a meaningful way. While it is always true that the participants might eventually change to some other strategy, perhaps to the optimal one, we find no evidence of learning within our experiment.¹⁵

5.1 Simulation

After the experiment was conducted and the data analyzed, it was possible to simulate what the optimal behavior would be taking into account the actual time people need to complete tasks. A simulation is necessary because the experiment adds a significant amount of complexity beyond the stylized model presented in Section 3. The simulated outcome of the experiment provides accurate theoretical predictions fully consistent with the experimental design, which can be directly compared to the outcome of the actual experiment. Looking at the data we deduce the time it took for a subject to bake a pizza by taking the

¹⁵As described in detail in section A of the online Appendix, we look for learning effects in two ways. The first is by comparing behavior in a treatment based on the order in which the subject experienced the treatment. Such an order effect, if it existed, would suggest that learning was being transferred across environments. Second, we compare a decision maker's behavior during the first portion of a treatment to her behavior later in the same treatment to look for environment specific learning.

difference between when the “Submit Order” button was clicked and when the last event happened—either the clicking of “Update Price” or “Submit Order.” The data show that the mode is stark and approximately 10 seconds, with a mean baking task time of (12.9, 12.8, 12.6, 13.7) seconds. If we also look at the time subjects took in order to update their *Current Set Price* to match the *Optimal Price*, we find qualitatively similar results. The data show that yet again the mode is very pronounced around 10 seconds, with mean updating times of (10, 10.5, 10.5, 10.7) seconds. We found that updating and baking times were not always statistically similar from treatment to treatment, but we find that the differences are economically negligible.

For the simulation, we chose 10 seconds for both the pizza baking and the updating task.¹⁶ To simulate this decision mechanism a MATLAB code was written to produce an optimal price vector according to the same generating function as discussed in Section 3. The code initializes assuming that the first task an individual will perform is the baking of a pizza because every period in the experiment starts with the instantaneous profit at its maximum. The code then moves forward in “time” along the optimal price vector the amount of time it takes for the baking task to be performed. Then it calculates the instantaneous profits earned by baking the pizza at this “time” and adds it to total profits. Next, if the instantaneous profit is *below* a given threshold, then the code moves forward in “time” the amount of time it takes to perform the price updating task and resets the *Current Set Price* equal to the *Optimal Price*. If the instantaneous profit is *above* the given threshold, then a baking task is initialized, meaning the code moves forward in 10 seconds “time” along the optimal price vector, and the instantaneous profit at that point is added to the total profits. The code does this until 322 “seconds” have been simulated.¹⁷ This entire process is referred to as one *simulation*.¹⁸ The code collects data on the total profits, the number of pizzas baked, and the number of price updates in that simulation. The code calculated the average total profit across 10,000 simulations performed at each decision threshold in the range $[0.01, 0.02, \dots, 0.99]$. The optimal inaction regions are defined by the thresholds that maximize the average total profits.

Table 2 summarizes the pertinent results of the simulation. Of particular note are the optimal decision thresholds, in ELD, of (0.52, 0.38, 0.38, 0.20). Further, the average number

¹⁶We looked at several variations of this task timing choice (mode, median, mean, and introducing noise). We decided to use 10 for ease of communication and because we found that the variation of the task time had little to no bearing on the outcome of the simulations.

¹⁷This is the average duration of the four treatments.

¹⁸An illustrative example of this process, and a pseudo-code have been provided in Appendix C.

of price updates at the optimal threshold are (2.9, 4.3, 4.3, 5.3); the average number of pizzas baked in each of the four treatments are (29, 27.7, 27.7, 26.7); and the average total profits were (21.4, 16.8, 16.8, 11.6) in ELD.

5.2 Inferring Thresholds from the Data

While the updating thresholds are clearly defined in the simulations, the inaction regions adopted by the subjects in the experiment are less easy to identify due to the mistakes the subjects make in establishing the correct thresholds. Therefore, the subjective thresholds of these regions must be indirectly inferred from the observable decisions of the subjects, accounting for the uncertainty they face in making their decisions. Our strategy is to rely on a comparison of the distributions of the profits observed by the subjects when they make a decision of either updating the price or baking a new pizza. Using only the distribution of the profits for the updating decisions would not be efficient because the distribution of profits for baking decisions provides valuable information about where it is more likely a subject places his implicit threshold. Intuitively, if the subjects understand they should follow a state dependent rule, they would start with a prior belief of the position of the threshold and search for the most profitable threshold as they conduct the experiment. Obviously, the process of uncertainty reduction is not going to be flawless, and we can suppose that the subjectively optimal threshold would lie inside a rather imprecise area of transition from the inaction region into the action region. Figure 8 illustrates the profit realizations (green line) and the baking/updating decisions (red and blue squares, respectively) for four of the subjects in the experiment (such plots for all 70 subjects may be found in Section B.4 in the online Appendix). Each row represents a subject, columns are the treatments, the numeration above each panel reflects the order in which that subject observed the treatments, the dashed horizontal lines show the optimal thresholds found in the simulation stage. Instantaneous profits are on the vertical axis, and seconds are on the horizontal one. As expected, the baking decisions are typically clustered in the high-profit region, while the updating decisions occur in a region of lower profits; however, the two regions are only rarely completely disjointed. Updating decisions are often made for very high profits as well, and some baking decisions can occur at profits below the average point of price updating. This figure also illustrates well the point emphasized in Section 4.1 that there is unresolved uncertainty between the decision node and its completion. We see in Figure 8 that immediately after the completion of a given task (i.e., the decision node point), the optimal price begins to deviate from where it was when the decision was made. Looking at the completion of that action, we can see that

the optimal price is as at a different level, showcasing the enduring uncertainty throughout the duration of the chosen real effort task.¹⁹

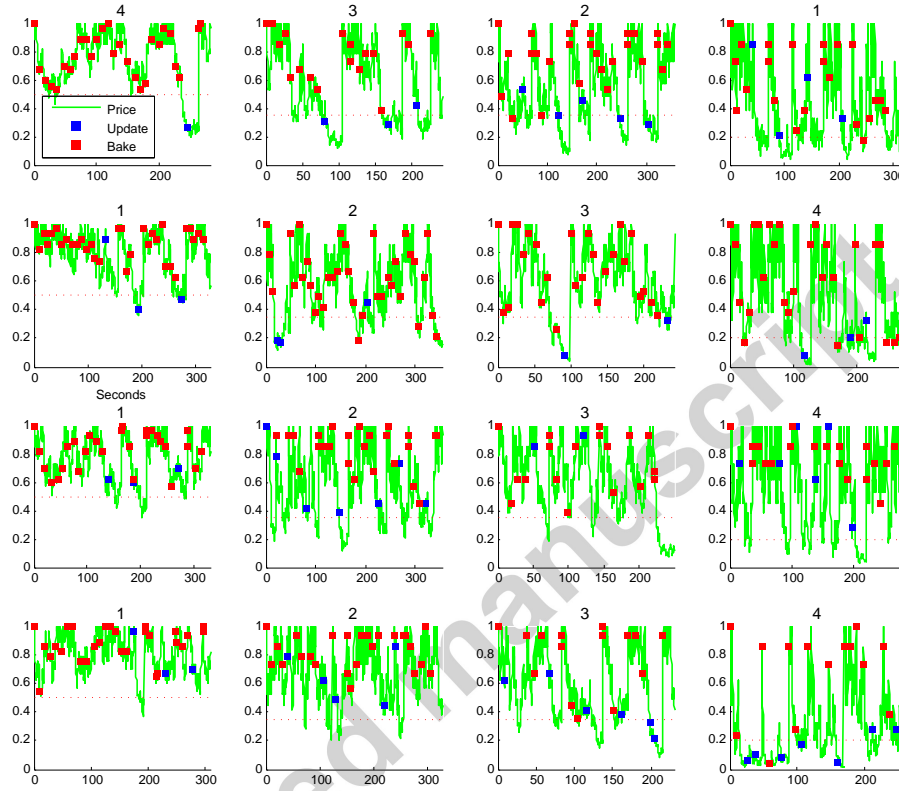


Figure 8: Price history (green line), price updating task decisions (blue dots), and baking task decisions (red dots) for a selected set of subjects. Each row represents a subject, columns are the treatments from T1 on the left to T4 on the right, the numeration reflects the sequence of the treatments for the subject, and the dashed horizontal lines show the optimal thresholds. Instantaneous profits are on the vertical axis and seconds on the horizontal one.

The characteristics of these plots suggest a simple approach to estimating the individual thresholds. First, we estimate the kernel density distributions of the two task decisions by using a normal kernel (that is, of the blue and the red squares in Figure 8) and compute the

¹⁹For comparison with Figure 8 and the companion plots for the entire set of subjects in Section B.4 of the online Appendix, Figure A2 in Appendix D reports the histories of profit realizations and decisions made by seven virtual subjects under optimal thresholds simulation. Although the thresholds are obviously always respected by these virtual subjects, this Figure confirms the volatility of profits at the moment of the price updating decisions that reflect the lag between the beginning of a task and its completion.

respective cumulative distribution functions.²⁰ The survivor function of the baking decisions (defined as $1 - CDF$) would typically be greater than that of the price updating decision on the entire support. This can be interpreted as a sort of stochastic dominance which reflects the fact the density distribution of baking is shifted to the right of the $[0, 1]$ support of profit realizations relative to the updating distribution, but it does not exclude the possibility of overlapping of the two distributions. We illustrate this point in general terms for the aggregate set of decisions across subjects in Figures 9 and 10. At the level of individual subjects the same result typically holds, with the exception of the few cases in which a subject never updates during the experiment or updates only once. For these cases, the density of the updating decision degenerates, and we attribute the threshold simply to that single observation.

Finally, the threshold is estimated as the profit level that maximizes the distance between these two $CDFs$, which basically corresponds to the point where the two densities switch relative position and updating becomes more likely than baking (coming from the right hand side of the support). We report the thresholds computed based on the aggregate observations for the four treatments in Figure 11. It is important to note that the thresholds are indeed following the relative ordering found in the simulation, with the lowest threshold for the last treatment and similar thresholds for the intermediate treatments. Nevertheless, as discussed below, the subjects seem to perceive their thresholds above the simulated ones, uniformly across treatments.

Turning to the analysis at individual level, we observe a higher degree of heterogeneity across subjects. The relative ordering of the thresholds is not often exactly satisfied by the individual subjects and the distribution of thresholds is quite spread over the profit support as illustrated in Figure 12. As an illustrative example, Figure A3 and Figure A4 in Appendix D report the survivor functions and the estimated thresholds for the same set of individual subjects for which Figure 8 is constructed.

We use the data of the optimal simulations to assess the validity of our approach to infer the thresholds from the distribution of the updating and baking decisions. The results of the assessment are reported in Section C of the online Appendix. The assessment confirms that our empirical strategy correctly recovers the thresholds used by the subjects in the experiment.

²⁰Equivalent results are obtained when other popular kernel functions are adopted.

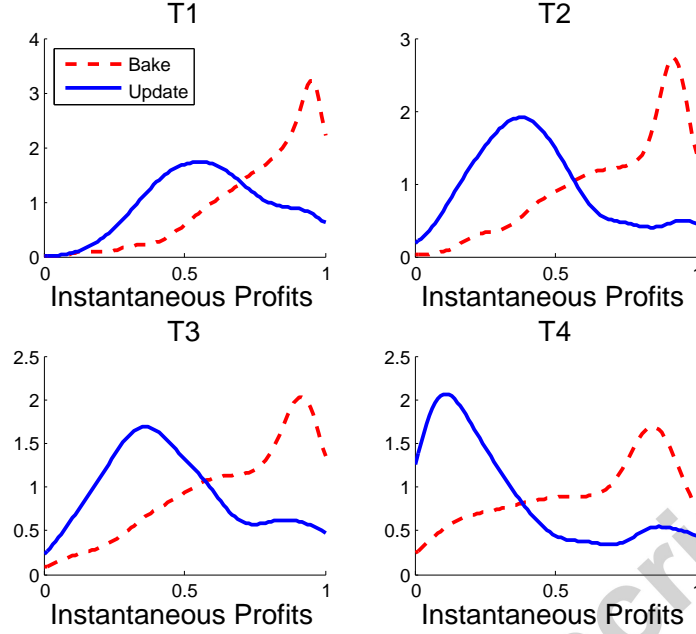


Figure 9: Empirical kernel distribution of the decisions of price updating and baking tasks. Gaussian kernel function for the aggregate set of decisions across subjects.

5.3 Experimental Results

We now turn to observed subject behavior. First, we examine the thresholds subjects use. Figure 11 shows histograms for the average decision threshold for all subjects in each treatment. The mean thresholds are (0.70, 0.57, 0.57, 0.38).²¹ All of these are *highly* statistically different from one another ($p < 0.001$) except for when comparing Treatments II and III.

As can be seen in Table 2, we find that in all four treatments, the average decision thresholds used in the experiment are in all cases substantially higher than those determined to be optimal by the simulations. Figure 12 shows a histogram of the subjects' decision thresholds in each of the treatments. The vertical red line indicates the simulated optimal threshold. As can be seen, the mass of the distribution is well above the simulated optimal.²²

²¹We test for the within-subject design by subtracting the average number of the relevant measured values for individual i in Treatment j to i 's average measured value in Treatment k . Then the average of those differences was taken for all individuals, and a paired *tw*-test was performed to test whether or not the mean of that difference vector was different from zero.

²²We would ideally test for statistical difference in the variance of these distributions to test the hypothesis that more volatile underlying economic fundamentals would result in a greater variance in the individual

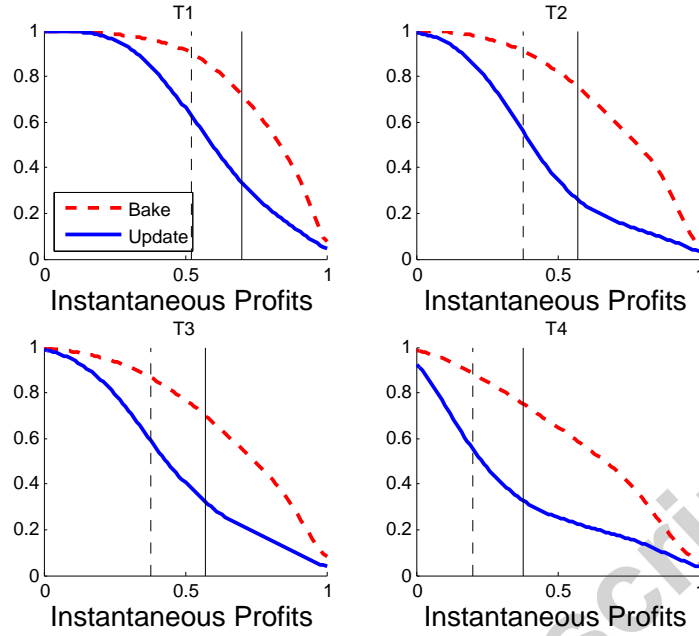


Figure 10: Survivor functions (1-CDF) of the two types of decisions. The dashed vertical lines correspond to the simulated optimal thresholds; the solid vertical lines are the estimated thresholds implied by the empirical distributions of the two task decisions. Aggregate set of decisions across subjects.




Figure 11: Estimated thresholds in the four treatments; aggregate set of decisions across subjects.

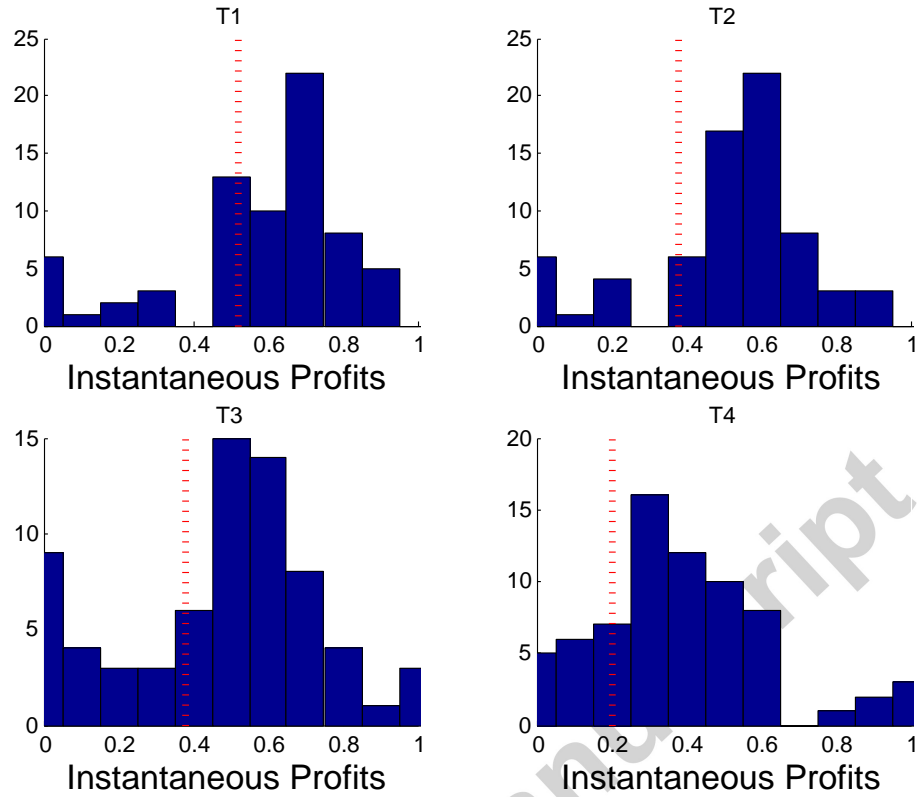


Figure 12: Dispersion of the individual estimated thresholds by treatment. Vertical dashed lines correspond to the simulated optimal thresholds.

We find the average number of updates per treatments to be (3.4, 4, 3.9, 4.8). Relative to each other, these averages follow the predictions of our stylized model: Treatment I (slow price movement, small profit falloff) has fewer price updates, while Treatment IV (fast price movement, large profit falloff) has more price updates, with Treatments II and III being less extreme. Statistically, we find that there is significant difference ($p < 0.001$) between any two treatments except for Treatments II and III, which do not differ statistically. We find the mean number of pizzas baked to be (24.7, 22.5, 22.9, 20.7). Again, we find that these averages are statistically different to a very high degree of significance ($p < 0.001$) in all cases except for the comparison of Treatments II and III, which do not differ statistically.

thresholds, but there are some limitations in our case. An F -test was performed, and the results showed that there were no statistical difference between any of the given treatments' variances. This should be heavily caveated with the fact that the F -test is highly sensitive to the assumption that the data follow a Gaussian distribution. Moreover, it should also be emphasized that the F -test requires that observations be independent. Since each subject in our experiment experienced each of the four treatments, this assumption is not met.

	Treatment I	Treatment II	Treatment III	Treatment IV
Experiment (means)				
Number of Price Updates	3.4	4.0	3.9	4.8
Number of Pizzas Baked	24.7	22.5	22.9	20.7
Decision Threshold [ELD]	0.70	0.57	0.57	0.38
Total Profits [ELD]	17.4	13.3	13.1	8.8
Simulation (means)				
Number of Price Updates	2.9	4.3	4.3	5.3
Number of Pizzas Baked	29	27.7	27.7	26.7
Decision Threshold [ELD]	0.52	0.38	0.38	0.20
Total Profits [ELD]	21.4	16.8	16.8	11.6

Table 2: Summary of pertinent data collected from the experiment and the simulation

Even though we find that subjects have a higher decision threshold in each of the treatments compared to the simulated optimal, the average number of updates in each treatment of the the experiment (3.4, 4.0, 3.9, 4.8) is quite similar to the number of price updates per treatment in the simulations (2.9, 4.3, 4.3, 5.3). At the same time, as can be seen in Figure 13, we find that the subjects bake statistically significantly fewer pizzas relative to what the suggested optimal would be in each of the four treatments (24.7, 22.5, 22.9, 20.7 compared to the 29, 27.7, 27.7, 26.7 as implied as optimal by the simulations). As a consequence, the total number of tasks performed by the subjects falls several units short of the total number of tasks completed in the simulations. On average, subjects perform approximately 4 to 6 fewer tasks in each treatment. Finally, considering the fact that the subjects can complete a smaller number of tasks relative to the simulated optimal, it is not surprising that we find that subjects earned statistically significantly less total profits than the simulations (Figure 13). On average, we find that subjects earned (17.40, 13.30, 13.10, 8.80) in ELD. However, the total profits earned vary as predicted by the stylized model from one treatment to the next and these averages are statistically different to a very high degree of significance ($p < 0.001$) in all cases except for when we compare Treatments II and III, which do not differ statistically. Again, all of these data are summarized in Table 2.

Two main effects can explain this set of results. First of all, given the total number of tasks completed, the price updating decisions are relatively more frequent than expected. This reflects the higher inaction thresholds identified by the subjects, which induce people to update “too early” in state and more often than optimal. The implicit trade-off cost of price

updating affects the subjects' perception of the actual value of an updating decision, causing them to overestimate the benefits from updating. In other words, subjects do not realize the correct cost in terms of profit losses they face when updating and seem to put more importance on producing at higher rates per pizza, thus hastening the price adjustment. This result can be interpreted in light of the large experimental literature that show that increased cognitive loads make people more impatient and lead to poorer economic decisions (see Deck and Jahedi, 2015, for a more complete survey of this literature).

Second, the subjects are not as efficient at making decisions as machines, and they will naturally complete a smaller number of total tasks. However, their efficiency is also affected by the uncertainty of the environment as the increasing gap between simulated total tasks performed and actual tasks shows (this gap goes from 4 for Treatment I to 6.5 for Treatment IV). Magnani et al. (2015) document a large dispersion of subjects' actions around optimal thresholds, which increases in the volatility of the environment. In their setup the adjustment cost is known and the average thresholds are quite close to the optimal ones. We also find a large heterogeneity in the thresholds across individuals, but allowing for implicit adjustment cost we illustrate a limited ability of the subjects to correctly identify the thresholds. Furthermore, the mistake they make in assessing the thresholds is quite uniform across treatments and depends only marginally on the uncertainty of the economic environment. The new feature we introduce in our experimental design, namely the trade-off between production and price setting tasks, has a direct negative impact on the cognitive capacity of the subjects. The implicit cost of adjustment forces subjects to infer their cost of resetting prices in terms of forgone profits while they accomplish the updating task, heightening their cognitive cost. Similarly, the trade-off adds a real source of distraction for the subjects that diverts their attention from screening the state of the economy to executing the production task. While it is hard to precisely disentangle these two channels, they both operate in the same direction by increasing the likelihood of optimization errors even in the baseline treatment.

5.4 The Empirical Inaction Regions

A person who is following a state dependent strategy, should always update the price when the instantaneous profit is below the threshold value (the action region) and should never update the price when instantaneous profits exceed the threshold value (the inaction region), as shown in Figure 1. For each subject in each treatment, we have already empirically identified the implied threshold value. Figure 14 examines how much variability each subject




Figure 13: Experimentally-determined Vs simulation-determined average number of baked pizzas per treatment (top), number of price updates (center), and total profits (bottom). Vertical bars show two standard deviations from the mean and are only suitable for comparison of means within treatment (but not across treatments).

exhibits when making choices in their respective action and inaction regions. Specifically, the vertical axis indicates the frequency with which someone updated the price when the instantaneous profit indicates they were in the action region—that is, the vertical coordinate indicates the percentage of times a person updated their price when they should have updated it. The horizontal axis indicates the frequency with which someone updated their price when the instantaneous profit placed them in the inaction region—that is, the horizontal coordinate indicates the percentage of times a person updated their price when they should not have updated it. A person who perfectly follows a state dependent strategy would appear at the top left of Figure 14, while someone who is equally likely to update their price when they are in the action or inaction region would appear on the 45 degree line.

From Figure 14 it is clear that subjects generally update more frequently when they are in the action region than when they are in the inaction region as almost every point lies above the 45 degree line in all four plots. In fact, many subjects never update when they are in the inaction region as indicated by the concentration of points along the vertical axis. However, when subjects are in the action region with low instantaneous profits they frequently do not update even though they should. Overall, Figure 14 suggests that subjects behave as if they use a state dependent rule but have difficulty identifying and implementing a stark threshold. Also this result can be linked to the effects of cognitive load and bounded rationality that the subjects face while inferring the optimal thresholds under the endogenous opportunity cost of updating prices.

5.5 Time Dependence Analysis

The degree to which our subjects exhibited state-dependent or time-dependent behavior is of particular importance. We assess the presence of time-dependent elements in subjects' decisions in two ways. The first way is by looking at the hazard function of price updates. We construct the hazard function computing the probability that an individual will decide to perform a price updating task after t periods, conditional on not having changed price for t periods. In order to construct these hazard rates, we first compute the duration of price spells since the most recent updating task for all subjects in a treatment, and we sort them in one-minute intervals. The last interval is not bounded in order to accommodate the non-predetermined end of the experiment for each subject. The hazard rates are then defined as the ratio of updating decisions in an interval to the total number of updating decisions made in that and all the subsequent intervals.

As discussed in Nakamura and Steinsson (2007), the shape of the hazard function depends

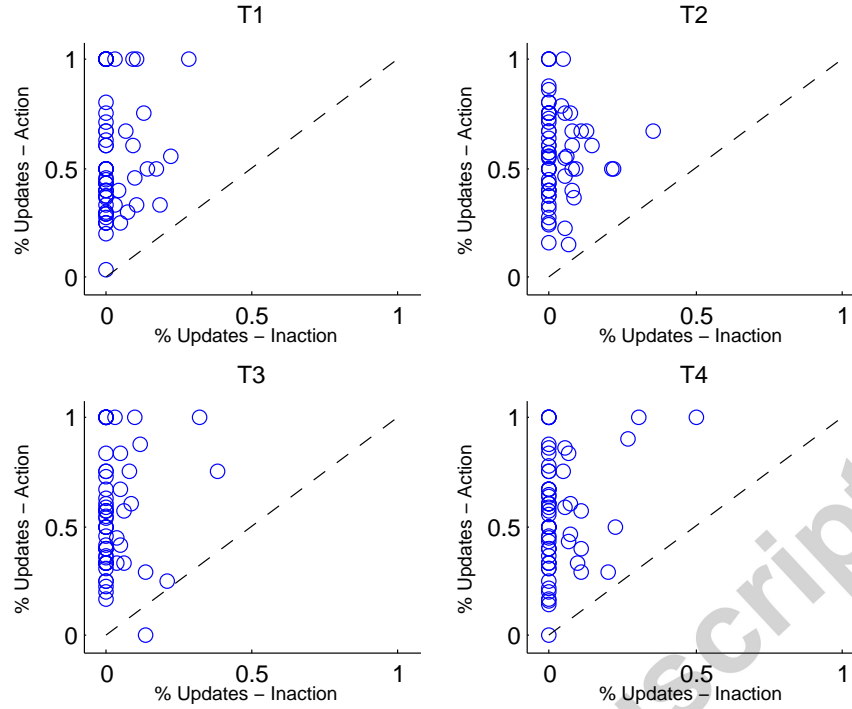


Figure 14: Scatter plot of the updating frequencies inside and outside the action region identified for each individual.

on the underlying pricing strategy followed by the price setters. In a menu cost model, the slope of the hazard function is related to the persistence of the shocks to the marginal cost. An upward sloping hazard function corresponds to non-stationary marginal cost, whereas less persistent shocks determine flatter (or even partially downward sloping) hazard rates. On the contrary, a flat hazard function reflects a constant probability of updating in every period and would correspond to Calvo pricing. The analysis starts from the hazard functions obtained from the simulated data, which we compare to the empirical hazards based on the data of the experiment. Figure 15 illustrates that the simulation-based hazard functions are indeed upward sloping, as typically expected for this type of state dependent model. Although similar in magnitude to the theoretical ones, the empirical hazard rates reported in Figure 16 definitely exhibit more irregular shapes, with essentially flat profiles for Treatment III and IV and large drops for some of the rates in the other two treatments. Moreover, contrary to the theoretical hazards, there are no observations for the longest durations in Treatment II to IV. This result reveals the existence of clear deviations of the subjects' decisions from the benchmark state dependent model and echoes the results in favor of time dependence found by Magnani et al. (2015) as well.

We further refine the empirical hazard measures by considering the updating decisions

relative to the inaction threshold of each individual. This allows us to separately study the updating decisions made inside and outside the estimated action region of an individual, controlling for potential differences in state dependence for those actions that may occur earlier if done inside the inaction region. Similarly, using individual thresholds helps us eliminate the “survivor bias” that may arise when mixing flat individual hazards for subject with heterogeneous thresholds, as documented by Klenow and Kryvtsov (2008).²³

Figure 17 reports the adjusted hazard functions for the four treatments of the experiment. The conclusions analyzed above are confirmed by the hazard functions of the decisions made inside the action regions, and time dependence does not seem to be determined by heterogeneity across subjects in our experiment. Furthermore, the hazard rates of the decisions made outside the action regions are downward sloping and, as expected, they show that these actions correspond to shorter durations. Subjects who update inside the inaction region are likely also more impatient. Finally, time dependence in the hazard functions is not particularly affected by high volatility treatments (the two treatments on the right hand side column of these Figures). Magnani et al. (2015) note that time dependence “increases substantially” when the volatility of the underlying economy is increased, but this result is not corroborated by our study.

For the second part of this time dependence analysis, we study the ratio of updating decisions to the total number of decisions made at each point in time in function of the time elapsed since the last price update occurred. This function represents how the probability of making errors in performing a price updating changes over time. As for the hazard analysis, we sort the observations by one-minute intervals and we separate them between action and inaction regions based on the individual threshold estimates. If price setters are strictly state dependent, the function would be equal to one and completely flat inside the action region and zero outside the action region. Allowing for uncertainty and errors in the identification of the thresholds, we would expect an essentially flat shape in the two regions (although the function would not be equal one or zero anymore) if time does not have a systematic bearing on the subjects’ errors in making decision to update prices.

Figure 18 reports these error probability functions for the four treatments of the experiment, which strongly confirm the presence of time dependent traits in the behavior of the subjects of the experiment. The functions are clearly downward sloping in every treatment for the decisions in the action region. On the contrary, in the inaction region, the functions

²³The survivor bias in this case would reflect the fact that individuals with lower subjective thresholds would take on average longer time to reach their action region compared to those with higher threshold values, causing a falling aggregate hazard function even when all individual hazards are flat.




Figure 15: Hazard plots for each of the four treatments – simulated data.




Figure 16: Hazard plots for each of the four treatments – experimental data.




Figure 17: Hazard plots for each of the four treatments dividing actions made inside and outside the identified action regions – experimental data.

start with an initial spike in the first minute and then become quite small and flat in the following intervals, which is more in line with expectations under a state dependent strategy. These results suggest two interesting observations on the types of errors made by the subjects. First, the initial spikes in the functions outside the action region indicate that some subjects are underestimating their cost of updating, and this leads them to update prices too “early” with respect to the state, when the deviation from the optimal price is not that large yet. As illustrated by the profits diffusion plots in Figure S13 of the online Appendix, the levels of profits for which this may occur are more likely observed at shorter durations. Second, the low persistence of the error functions inside the updating region represented by the steep negative slopes of the functions can be interpreted as evidence of habit formation in baking. Subjects become less likely to perform updating tasks as the price spells get longer, which will require subjects to be engaged in the production activity for longer durations. A prolonged focus on a series of baking tasks diverts the attention of the subjects from updating, and reduces the ability of the subjects to recognize the need of a price change even after entering the action region. This mechanism would be consistent with the view that the production task operates as a distracting factor, which increases the cognitive load of the subjects and the likelihood of mistakes in the updating process.

In conclusion, the evidence of time dependence elements in the updating rules of the




Figure 18: Error plots for each of the four treatments – experimental data.

subjects presented in this section are broadly consistent with the results of Magnani et al. (2015). However, we find that these elements are mostly related to the cognitive load mechanism introduced by our experimental design, rather than to the underlying volatility of the economic environment. In relation to the empirical data from the field, our hazard functions have shapes more similar to the flat hazard rates computed by Klenow and Kryvtsov (2008) than the downward sloping hazard functions estimated by Nakamura and Steinsson (2007). This type of hazard function would primarily point to models in which the probability of updating is constant over time, such as the Calvo pricing, and therefore completely independent of the state of the economy. In this respect, our experiment suggests that, although this type of hazard function does not reflect a purely state dependent pricing rule, it might still come from subjects who recognize the validity of a state dependent approach to some extent, but do not always successfully follow it because they experience attention diversion or limited information processing capacity.

6 Conclusion

The issue of price setting has been an area of intense macroeconomic research for decades. With the abundance of possible explanations for observing macroeconomic price rigidity and the emergence of the use of laboratory testing to evaluate the validity and applicability of theoretical models, a logical endeavor is to evaluate the efficacy of these price setting models

in a laboratory setting. This paper uses an experiment to understand if individuals behave in accordance with a strictly state-dependent pricing framework. We do not expect that subject behavior will perfectly match such a theoretical model. Rather, our objective is to identify how well an Ss model captures the decision rules that people actually use and how their responses vary with changes in the environment.

Our work extends Magnani et al. (2015) by internalizing the costs incurred by subjects to update their prices and thereby introducing an intrinsic element of additional cognitive demand on the subjects. In this richer environment, we find that in general subjects behave as if they recognize the importance of a state dependent pricing strategy, but they identify noisy inaction regions and they also exhibit a substantial degree of *time* dependence. As a result, subjects adopt decision thresholds higher than optimal and they end up updating relatively too often instead of engaging in production as compared to the simulated optimal behavior.

Our results are consistent with the conclusions of Magnani et al. (2015) and their explanation for time dependence based on cognitive or attention costs. However, we also highlight a second source of errors related to bounded rationality, which has a very large impact on the pricing decision of our subjects. When the updating cost is not imposed exogenously, people will tend to underestimate the actual cost they face to update and they systematically identify higher thresholds even in low uncertainty environments. The introduction of the real trade-off between production and price updating in our experimental design allows us to further explore the link between attention and cognitive load in price setting on one side and some of the resulting errors on the other. We make a few interesting observations about this point.

The first is that the literature on the negative effects of cognitive load on the quality of economic decisions provides a possible interpretation of the main result. It has been shown that an increased cognitive load makes people more impatient; in our framework, early updating can be seen as the effect of impatience on subjects' decisions. As a second point, our experimental design introduces a direct and simple micro-foundation of a source of distraction of the subjects' attention through the production activity. It would be relatively easy to design modifications of this framework and study how pricing errors respond to changes in the type of cognitive cost in order to assess different theories of attention. For instance, the evidence of deviations from optimal price setting due to subjects' limited ability to fully analyze the underlying economic conditions could possibly lend credence to the strand of models that focus on firms' information processing bandwidth or even "rational

inattention.” Specifically, such a model would have to entail explicit decisions about attention allocation across signals of different precisions, for example associated to updating tasks of different lengths, to infer the underlying marginal cost of production; one of the testable predictions of the model would be that the precision of the selected signal should increase in response to higher returns of an incoming order.

Related to the discussion of possible modifications of our design, the third observation we make is a caveat on the mechanism implemented by our approach. While our experimental design endogenizes the opportunity cost of updating prices, the updated price was always set optimally. In practice, if sellers are not given the optimal prices and have to invest effort to estimate the optimal price level, then the cognitive costs associated with updating prices would be relatively greater than in our design. This increased cost, coupled with the risk of earning sub-optimal profits as a result of an error in setting prices optimally, would likely lead sellers to update prices less frequently than what we observe.

The last observation is that our research also provides some evidence on the behavioral support of different theoretical price models as well as of the empirical evidence obtained from field data on price changes. By directly testing subjects behavior in a state dependent environment, the first conclusion is that strictly state dependent models are not sufficient to represent pricing decisions; a relatively significant time dependent component emerges as well. From the flat hazard functions found for the experimental data, we can conclude that simpler staggered price models as in the Calvo pricing might still represent a valid alternative to other models to capture these deviations from full state dependence. Klenow and Kryvtsov (2008) find that flat hazard rates are typical with field data as well. Finally, as just discussed above, our entire analysis is impregnated by the relevance of cognitive and attention costs for price updating decisions and it calls for further, promising exploration in this direction.

References

- Daniel Benjamin, Sebastian Brown, and Jesse Shapiro. Who is ‘Behavioral’? Cognitive Ability and Anomalous Preferences. *Journal of the European Economic Association*, 11(3):1231–55, 2013.
- Mark Bills and Peter J Klenow. Some Evidence on the Importance of Sticky Prices Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy*, 112(5):947–985, 2004.
- A. S. Blinder, E. R. D. Canetti, D. E. Lebow, and J. B. Rudd. *Asking About Prices*. Russell Sage Foundation, New York, New York, 1998.
- GA Calvo. Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics*, 12(1978), 1983.
- A Caplin and J Leahy. State-dependent pricing and the dynamics of money and output. *The Quarterly Journal of Economics*, 106(3):683–708, 1991.
- Dennis W. Carlton. The Rigidity of Prices. *American Economic Review*, 76:637–658, 1986.
- Stephen G. Cecchetti. The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines. *Journal of Econometrics*, 31:255–274, 1986.
- Cary Deck and Salar Jahedi. The Effect of Cognitive Load on Economic Decision Making: A Survey and New Experiments. *European Economic Review*, 78:97–119, 2015.
- Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde. Are Risk Aversion and Impatience Related to Cognitive Ability? *American Economic Review*, 100(3):1238–60, 2010.
- Stanley Fischer. Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule. *The Journal of Political Economy*, 85(1):191–205, February 1977.
- Drew Fudenberg and David K. Levine. A Dual-Self Model of Impulse Control. *American Economic Review*, 96(5):1449–76, 2006.
- F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.

- Alan Kackmeister. *Has Retail Price Behavior Changed Since 1889? Evidence from Micro-data*. PhD thesis, University of California, Berkeley, 2001.
- Daniel Kahneman. *Maps of Bounded Rationality: A Perspective on Intuitive Judgment and Choice*, 2002.
- Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, New York, 2011.
- Anil K. Kashyap. Money, Credit, and Prices in a Real Business Cycle. *Quarterly Journal of Economics*, 110:245–274, 1995.
- Peter J Klenow and Oleksiy Kryvtsov. State-dependent or Time-Dependent Pricing: Does it matter for recent U.S. inflation? *The Quarterly Journal of Economics*, CXXIII(3), 2008.
- Daniel Levy, Mark Bergen, Shantanu Dutta, and Robert Venable. The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains. *Quarterly Journal of Economics*, 112(3):791–823, 1997.
- Robert E Lucas. Some International Evidence Tradeoffs. *American Economic Review*, 63(3): 326–334, 1973. ISSN 0002-8282.
- James N. MacDonald and Daniel Aaronson. How Do Retail Prices React to Minimum Wage Increases? *manuscript, Department of Agriculture*, 2001.
- Jacopo Magnani, A Gorry, and Ryan Oprea. Time and State Dependence in an Ss Decision Experiment. *American Economic Journal (forthcoming)*, pages 1–50, 2015.
- N. Gregory Mankiw. Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly. *The Quarterly Journal of Economics*, 100(2):529, May 1985. ISSN 00335533. doi: 10.2307/1885395.
- N. Gregory Mankiw and Ricardo Reis. Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics*, 117(4):1295–1328, 2002. ISSN 0033-5533. doi: 10.1162/003355302320935034.
- Filip Matejka. *Rationally Inattentive Seller: Sales and Discrete Pricing*. 2010.
- George A. Miller. The Magic Number Seven Plus or Minus Two: Some Limits on our Capacity to Process Information. *Psychological Review*, 63(2):81–97, 1956.

- Kanchan Mukherjee. A Dual System Model of Preferences Under Risk. *Psychological Review*, 117(1):243–255, 2010.
- Emi Nakamura and Jon Steinsson. Five Facts About Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics*, 123(4):1415–1464, 2007. ISSN 0033-5533, 1531-4650. doi: 10.1162/qjec.2008.123.4.1415.
- Christopher a. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003. ISSN 03043932. doi: 10.1016/S0304-3932(03)00029-1.
- John B Taylor. Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy*, 88(1):1–23, 1980.

Appendix

A Instructions

Introduction

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a considerable amount of money, which will be paid to you in cash at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is private. To ensure best results for yourself and accurate data for the experimenters, please **DO NOT COMMUNICATE** with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will come to you.

In the experiment you will make decisions over four periods. At the end of the last period, you will be paid \$5, plus earnings based on your profits earned in each period.

Task Details

You are in control of a company that bakes and sells pizzas at a Current Set Price. Your job is to fulfill incoming orders by assembling pizzas with the correct ingredients. When submitting an order, you will earn profits based on your Current Set Price and the Optimal Price. However, market conditions are constantly changing, so the Optimal Price is also constantly changing.

The further your Current Set Price is away from the Optimal Price (either too high or too low), the less money you make when submitting an order correctly. At any moment (and as often as you like) you can change your Current Set Price to the Optimal Price; but in order to do this, you must enter a code to update your price. After doing so, the Optimal Price will continue to fluctuate just as before, possibly moving away from the Current Set Price again.

Panel Details: Outstanding Order

In the upper left-hand corner of your screen, you will find a panel that looks like the one above. This is where orders come into your pizzeria. In order to fulfill an order, you must

select the button next to the ingredient that is in each box at the top of each column. Each column is populated with ten possible ingredients, and the ordering of each of these columns is random for every order. After you have correctly selected each ingredient, you will click the "Submit Order" button. Doing this will earn you profits and will bring in the next order. The more orders you complete correctly, the more profits you earn.

It is important to note that only one ingredient may be selected at a time. If you submit an order with more than one ingredient selected in a given column, you will not receive any profits, and you will have to try again.

Panel Details: Optimal Price

Remember that the market conditions are constantly changing, which means that the Optimal Price for your pizzas is also changing. The Optimal Price changes randomly 5 times per second. In the lower-left, you will find a panel that shows you the current Optimal Price, like the one in the image above. The green line indicates the Optimal Price, while the red line indicates your Current Set Price.

Panel Details: Update Code

Remember the further the Optimal Price deviates from your Current Set Price, the lower your profits per order. At any point you can update your Current Set Price. In the upper right-hand panel of your screen, you will see a 5-digit code that must be entered in order to update your Current Set Price. Simply select the button next to the appropriate number in each column to match the number at the top of that column and click the "Update Price" button, and then your Current Set Price will equal the Optimal Price at that moment, thereby maximizing your profit per order. But remember: The Optimal Price is always moving, so it may be the case that shortly after updating your Current Set Price, your profit per order starts to change.

Again, only one number in each column may be selected in order to enter the code correctly. Also, the ordering of each column is randomized each time a new code is generated.

Panel Details: Profit Information

The fourth and final panel on your screen will give you information regarding your profits. On the left, you will find a field called "Instantaneous Profits." The number in this box will update in realtime, and it indicates the amount of profit you will earn if you were to submit

an order at that exact moment. On the right you will see a box called "Cumulative Profits." This box will display the total profits you have earned throughout the period.

Period Details

There will be 4 paid periods. Each period's length is determined randomly, but on average each one is approximately 5 minutes in length. There will be a 1 minute break between each period. Ten seconds before the next period begins, your screen will turn red. Each period will be similar, but can differ in two ways: (1) The movement of the Optimal Price over time can change; and (2) How instantaneous profit depends on the difference between the Optimal Price and the Current Set Price can change.

If you want to go back and get instructions, you may do so now by clicking the red "Back" button below. If you are comfortable with the layout of the experiment, you may proceed to a 2-minute practice round by clicking the blue "Start" button below. This practice round will have no impact on your total profit and it is intended entirely as a means by which you may familiarize yourself with the operations of the experimental interface.

Ready To Begin?

If you have any questions regarding the experiment, raise your hand now, and an experimenter will come to you and answer them.

If you are ready to begin, please raise your hand, and an experimenter will bring you a short comprehension quiz to ensure that you fully understand the experiment. Once you have successfully completed the quiz, an experimenter will enter the password to allow you to begin the experiment. Note that once the experiment begins, you cannot take a break.

B Comprehension Quiz

This will not affect your payoff in any way. It is designed to make sure that you understand the experiment before it begins. Please answer the following questions and then raise your hand so that a researcher can come check your answers. If you have a question at any point, please raise your hand and a researcher will approach you.

1. The only way you make money in this experiment is by making a pizza with the correct toppings.
 - (a) True
 - (b) False
2. How much money you earn making a pizza depends on
 - (a) How long ago you updated your Current Set Price.
 - (b) The Optimal Price when you last updated your Current Set Price.
 - (c) The difference between your Current Set Price and the Optimal Price when you click Submit Order.
3. Your Current Set Price will adjust
 - (a) Automatically every 0.2 seconds.
 - (b) When you correctly enter the Update Price code.
 - (c) Never
4. Each period will last
 - (a) Exactly 5 minutes.
 - (b) A random amount of time that on average will be about 5 minutes.
5. From period to period
 - (a) The Optimal Price will change at the same rate.
 - (b) The reduction in profit per pizza as Optimal Price and your Current Set Price drift apart is the same.
 - (c) Both (a) and (b) are correct.
 - (d) Neither (a) nor (b) is correct.

C Simulation Pseudo-code




Figure A1: An example of an optimal price vector (blue line) and the current set price (green line).

Using the figure above, we will walk through a simple example of how the logic flows for one simulation.

1. The code initializes (point *a* in Figure A1) knowing that at point *a* the *Current Set Price* is equal to the *Optimal Price*. Because of this, it logically knows that it should bake a pizza.
2. The code then steps 10 seconds into the future (which corresponds to point *b* in Figure A1) because it takes time to bake the pizzas, the instantaneous profits earned by baking it are computed at this point.
3. The code then looks at the instantaneous profits available to it at point *b*, and:
 - (a) If the instantaneous profits are below the optimal decision threshold, then the code would set the *Current Set Price* to the value of the *Optimal Price* at point *b*.
 - (b) Else, the code decides to bake another pizza.

4. Regardless of the outcome of the previous step, the code then moves 10 seconds along the optimal price vector to point c in Figure A1, and
 - (a) If the previous action was a price update, then the code performs Step 3 again.
 - (b) Else, the previous action was the baking of a pizza, then the code computes the instantaneous profits for baking the pizza, and earns the value for point c .
5. The code then cycles through decision tree represented in Steps 3 and 4 for each point along an price vector that simulates 5 minutes.

This entire process is *one* simulation. The process then ran for 10,000 simulations.

D Additional Figures

We present here some additional Figures mentioned in the main body of the paper.

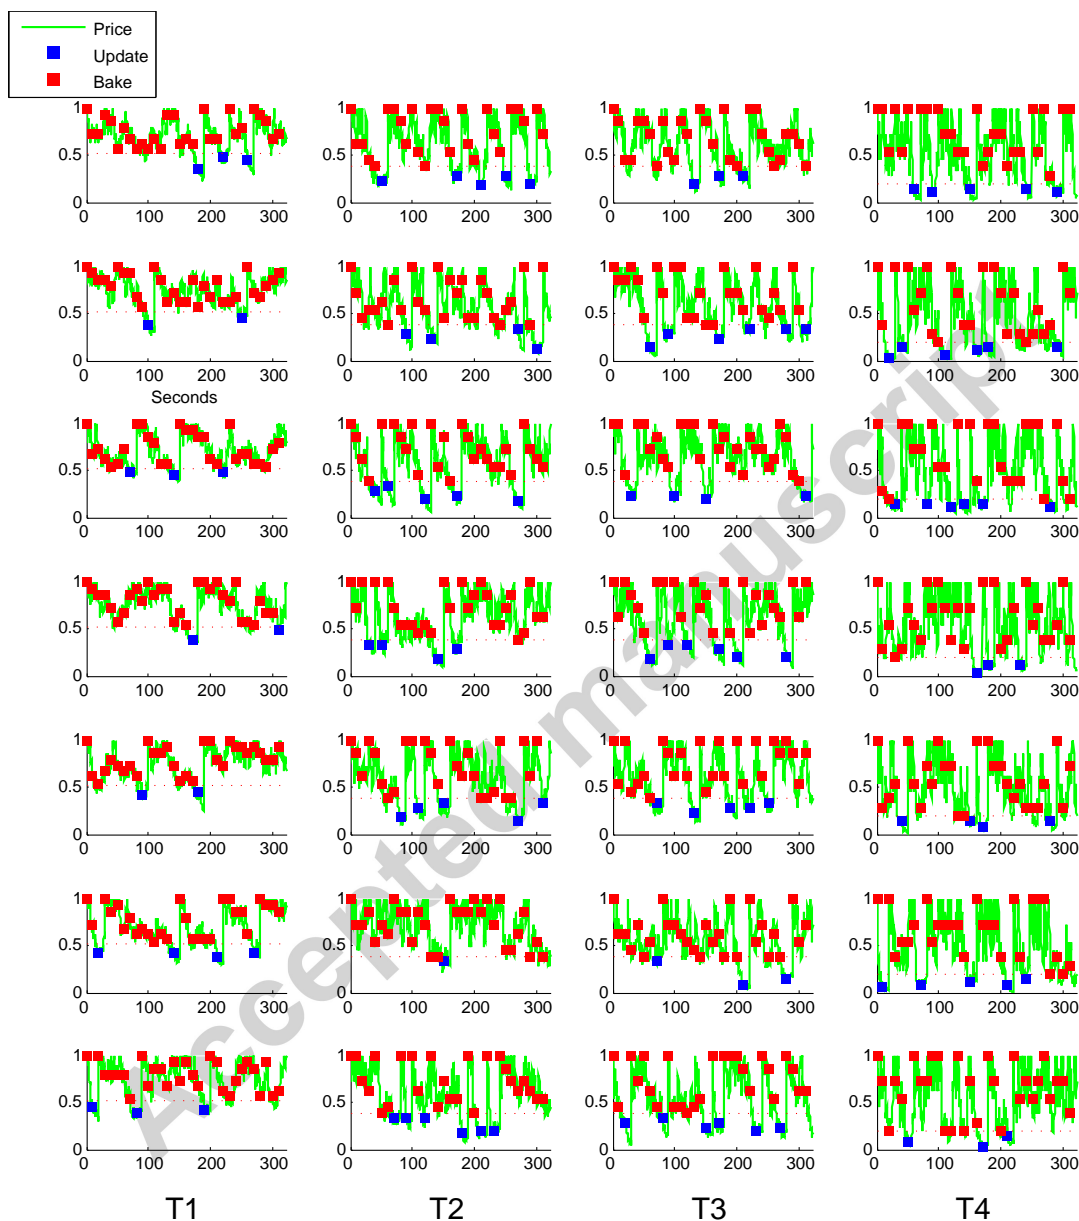


Figure A2: Price history (green line), price updating task decisions (blue dots), and baking task decisions (red dots) for a set of seven "subjects" from the optimal simulation. Each row represents a subject, columns are the treatments from T1 on the left to T4 on the right, and the dashed horizontal lines show the optimal thresholds. Instantaneous profits are on the vertical axis and seconds on the horizontal one.

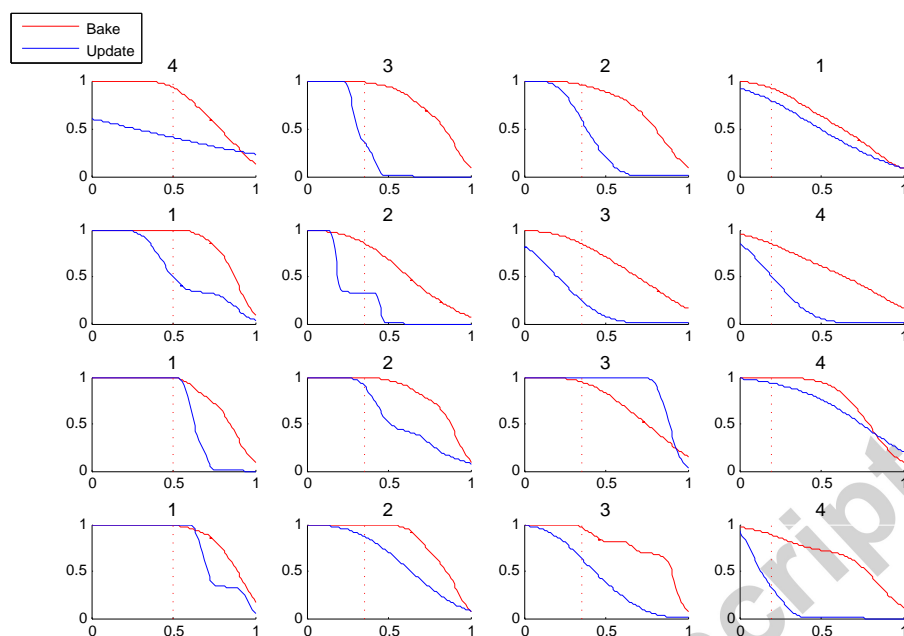


Figure A3: Survivor functions (1-CDF) of the two types of decisions for the same subjects in Figure 8 with instantaneous profits on the vertical axis. See caption of that figure for details.




Figure A4: Estimated thresholds in the four treatments for the same subjects in Figure 8. Each panel represents a subject, bars correspond to the four treatments, and the numeration reflects the sequence of the treatments for the subject. Instantaneous profits are on the vertical axis and treatments on the horizontal one.