Deeper Habits

Mel Win Khaw and Oskar Zorrilla *

November 27, 2018

Abstract

This paper offers experimental evidence for the microfoundations of the inertia observed in the aggregate consumption time series data. We design an experiment that is analogous to a consumption/savings problem in which our subjects simply decide whether to buy or decline individual assets in the absence of switching costs. Although all of the information relevant for optimal behavior is available at all times, we find that subjects nonetheless condition on past choices. We show that models of habit formation cannot account for the inertia in our data and argue that consumers condition on past actions as a way of economizing on cognitive resources. We develop a model of "rationally inattentive reconsideration." Importantly, our model implies that inertia is state-dependent. Within this framework the costs of inertia are estimated to be around one percent of consumption.

*Email: o.zorrilla@columbia.edu. First and foremost I have to thank my adviser Mike Woodford for all of his guidance. I also want to thank Hassan Afrouzi and Mark Dean for their continued support throughout. We would like to thank François Gerard, Stephanie Schmitt-Grohe, and Pierre Yared for their comments. All mistakes are our own.

1 Introduction

One of the most striking counterfactuals of economies populated by purely forward looking agents is the sensitivity of responses to shocks. As a result, a central concern of macroeconomic models in recent decades has been microfounding the presence of backward-looking terms—*inertia*—in the policy functions of agents. Although there has been a lot of progress, both empirical and theoretical, on the microfoundations of inertia on the firm side less attention has been paid to households. The purpose of this paper is to address the question of inertial behavior in consumption/savings decisions.

To achieve inertia in aggregate consumption, macroeconomists appeal to non-time-separable preferences: habits. Although estimates of habit parameters from aggregate data are positive and significant, there is little direct evidence for these types of preferences. Nor can the the consumption time series allow us to distinguish among different explanations for inertia. In fact, macroeconomists justify the use of habits in preferences *only* on the basis that they help to match empirical moments. As Nakamura and Steinsson (2018) put it, habits "allow the model to better fit the shapes of the impulse responses we have estimated in the data."

But evidence for inertia in consumption is not the same as evidence for habits in preferences. And distinguishing among the possible mechanisms that give rise to inertia is of particular interest to macroeconomics: understanding how agents *make choices* is crucial to evaluating the effects of alternative policies. Given the importance of teasing out this distinction and the difficulty of doing so with aggregate data, this paper runs a lab experiment in which we have subjects play a game that is analogous to a consumption/savings problem. We identify cognitive costs as a source of consumer inertia and develop a model to account for the inertia observed in our data. Importantly, agents in our model *choose* to condition on past choices optimally. We argue that these are *deeper* habits in the sense that they are not simply assumed—as is the case with non-time-separable preferences—but rather are a result of agents' economizing on cognitive resources. Consequently, these deeper habits are sensitive to changes in the state and therefore are not policy invariant.

Our experimental design has our subjects play a game that requires them to accumulate as many valuable assets as possible. Each period our subjects make a simple binary decision whether to accept or reject the asset on offer. The value of each asset is drawn from a known distribution. To pay for the assets our subjects are endowed with tokens which they receive randomly from time to time. In the case they decide to buy the asset on offer, they must pay one token. If they have no tokens, they cannot not buy the asset. To induce discounting, at the end of each period the game has a fixed probability of terminating. Once the game ends, subjects are paid only for the total value of all the assets they have bought; any residual tokens have no value. Although highly stylized, our framework captures the dynamic tradeoff at the heart of consumption/savings problem. Our subjects must decide whether to purchase an asset today versus the probability of being unable to buy a better offer in the future if they hit the credit constraint. In the same way that wealth can be turned into consumption, tokens can be turned into assets, but not the other way around. Tokens, like wealth, have value only because they represent future consumption.

By studying a consumption/savings problem in the lab, our paper makes two important empirical contributions to the consumption literature. First, we are able to find inertia in *choices* rather than simply as a feature of aggregate data. This first piece of evidence favors writing down models in which the households euler equation for consumption contains backward-looking terms. Further, we are also able to identify cognitive constraints as the source of this inertia in our subjects. This finding is important because it points toward a more rigorous and realistic microfoundation of households in general equilibrium models. A third contribution of this paper is theoretical. We propose a model of choice disciplined by the evidence in our data that can be adapted in a straightforward manner to DSGE models. An important advantage of our methodology is that lab-generated data can help us narrow the ways in which we model deviations from the rational expectations benchmark. In the model of rational inattention we propose, agents pay a fixed cognitive cost every time they "think through" a problem. As a result, they condition on past actions as a way to forego this thinking cost. Because our agents are rationally inattentive, they choose the probability of reconsideration conditioning on the state and their past actions. Our model also has the implication—at odds with the habits from preferences model—that the degree of inertia depends both on the policy rule and the realization of the state.

The fact that our experimental design shuts down any external motives to condition on past choices makes allows us to rule out some of the alternative explanations in the literature. We take care to eliminate external costs of switching, which is the mechanism that leads to inertia in models like Klemperer (1995). Under the commonsense assumption that our subjects seek to maximize their expected payoff from participating in our experiment their objective is time-separable. This feature rules out habits from preferences as in the models of Ravn, Schmitt-Grohe and Uribe (2004). Laibson (2001) develops a behavioral model were external cues may condition preferences and lead to habits. Like ours, his is a cognitive model, but the process he has in mind is quite different from ours. His model is in the spirit of Becker-Murphy rational addiction models. Our model is much more closely tied to the rational inattention literature, specifically, the model developed by Woodford (2009).

Although we model cognitive constraints broadly as a rational inattention problem our framework allows us to distinguish among several of the inattention models as explanations for our data. Our subjects have access to all of the information they need each period in order to make their choice and no time limit. Information is not sticky in the sense of Mankiw and Reis (2002) or Reis (2006); unlike their models were consumers forgo any new information for several periods the information is on the screen every turn. If anything it is harder to ignore than to read. Similarly, it is difficult to appeal to models of exogenous imperfect information such as Woodford (2002) and Lorenzoni (2009). Angeletos and Huo (2018) show that higher order beliefs can anchor outcomes on past behavior. This hinges on the joint assumption of imperfect information and strategic complementarity; our experiment has subjects solve a decision problem and thus there are no strategic considerations.

Our paper is most closely related to Matyskova *et al* (2018). Like us, they find evidence of state-contingent habits in an experimental setting. Their focus is somewhat different, yet complementary to ours. The purpose of their experimental design is to distinguish between two alternative cognitive explanations for habit formation. To that end their subjects are asked to identify the outcome of a binary random variable through time. In contrast we are interested in whether subjects display inertia in a specific context relevant to macroeconomics: while solving a consumption/savings problem where past actions are not directly payoff-relevant. Crucially, they conclude that habits in their choice data are not formed mechanically. Their findings that subjects condition on the past as a way to alleviate cognitive costs is in line with our own interpretation of the experimental results presented here.

To be clear, our experiment cannot rule out the existence of habits from preferences. We cannot take a stand either way on whether preferences are time-separable. Likewise, the fact that we make all relevant information available and eliminate any switching costs in the lab does not mean that we believe these conditions necessarily hold when households make consumption/savings choices. Rather than a weakness, however, this is precisely the strength of our approach. All of the channels that we shut down in our experiment are hard, if not impossible, to account for when looking at field data. The fact that the information set of the econometrician is not the same as that of the agents has long been an issue in the econometrics of rational expectations. This asymmetry is not present in the lab; we do not need any auxiliary assumptions about preferences, external constraints, or the stochastic nature of the data-generating process. Nor do we have to worry about measurement error of the relevant variables. Measuring the wealth of individuals in cross-section or panel data, for example, is notoriously difficult. By contrast, in our experiment it is trivial.

In a deeper sense, the strength of our experiment lies in that it sheds light on how subjects *make choices* when faced with a consumption/savings problem. We do not set out to measure inertia ex ante; actually, we make sure that there is no reason for our subjects to display it. The fact that they do is very strong evidence that cognitive limitations are at least part of the source of inertia in field data. Another way to interpret it is as a *lower bound* on the degree of inertia. Furthermore, our approach also uncovers an important result that is relevant for policy analysis: the degree of inertia depends on the stochastic process and the realizations of that process. Our model implies that inertia will be at its weakest when rare but large shocks occur. Consequently, models with state-independent inertia, like the standard habits-from-preferences models used in macroeconomics may overstate the degree of sluggishness of aggregate demand in “crisis” periods.

Section 2 outlines the experiment while section 3 solves the problem faced by our subjects. In Section 4 we establish our main empirical result—inertia. Section 5 introduces a model of inattentive reconsideration that matches our findings and within this framework we provide estimates for the cost of inertia. We consider some alternative explanations for our findings in section 6. Section 7 concludes.

2 Experimental Design

We describe the problem to our subjects as a real estate investment game where subjects’ objective is to accumulate as many “rents” from purchasing “properties” as possible. “Rents” in the context of our experiment are not the same as rents in the way economists usually understand them; we simply call them this because of how we frame the game to our subjects. In order to make the problem as simple as possible, the rental payments do not accrue over time so subjects do not have to calculate the discounted present value of each property. The “rent” on each property is paid to the subject at the moment it is bought and never again. If the subject buys the property on offer she receives the one-time rental payment but nothing otherwise. The rents of the properties range from 500 to 100,000 in increments of 500 and are drawn independently each period from a discretized log-normal distribution.

All properties have the same price of $1,000,000 in cash and subjects start each game with a cash stock of $11,000,000, which is the mean of the ergodic distribution of cash under the rational expectations policy. With a probability of 10 percent subjects receive –and are notified– of an income payment of $4,000,000 at the start of each turn. There is no interest nor any depreciation of cash so that the cash at the end of each turn carries unchanged into the following turn if the game does not terminate. At the end of each period the game ends with probability .002; this implies that the average game lasts 500 rounds. Upon termination of the game subjects are

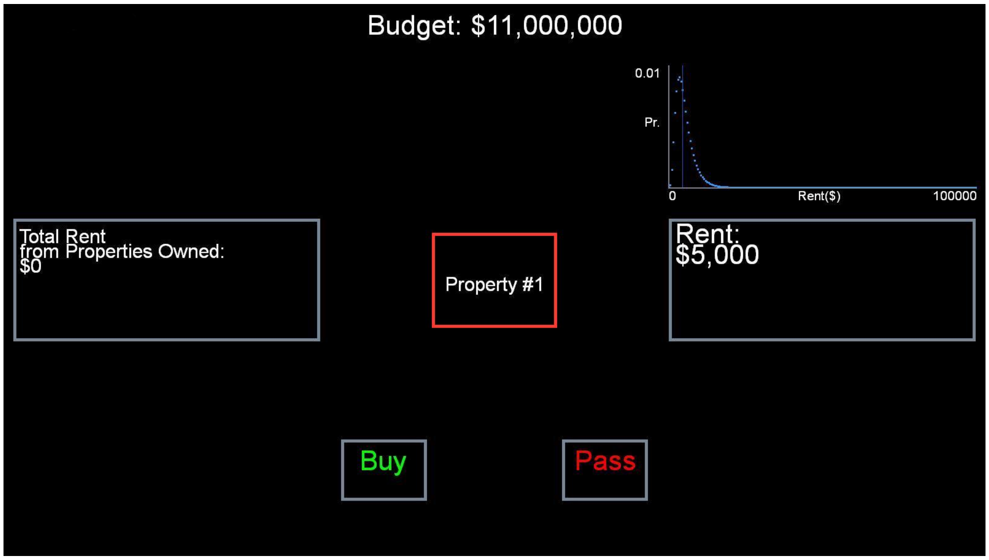

Figure 1: The computer screen during a single trial of our experiment. While subjects deliberate, we present to them information regarding the distribution of possible rent values, the current rent value, the accumulated rents of the subject, as well as the current budget available to be spent.

paid *only* for their total accumulated rents at an exchange rate of .00001. In dollar terms this implies that the rental payments on the properties range from half a cent to 1 dollar. Rents and cash, like consumption and wealth, are separate instruments. Subjects can turn cash into rents, but they cannot turn rents into cash.

Figure 1 shows a screenshot of the experiment in the first period of the game. Center top of the screen is the budget available to the subject. In the upper right hand corner is the discretized log-normal distribution of the rents with a vertical line that intersects the x-axis at the rental payment of the current property on offer. Below the graph of the probability mass function is the value of the rental payment associated with the property on offer. The center of the screen shows the number of turns played, inclusive. Center left is the running total of rental payments that subjects have received. Recall that each property offers only a one-time rental payment, so this number only increases when a new property is purchased.

Once subjects press a button, we do not ask them to confirm their choice, but we

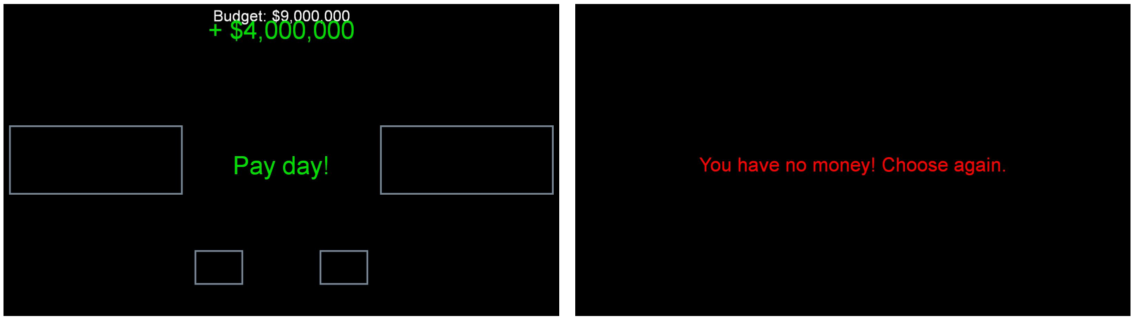

Figure 2: On trials where participants receive income, a notification of "Pay day" is displayed on screen for 3 seconds before subjects are allowed to make a decision (returning to the default interface in Fig.1). If subjects decide to purchase with no remaining money in their budget, subjects are displayed a notice in red for 1.5 seconds before they are returned to the decision screen.

verify the consequences of that choice. Upon making a decision the corresponding deduction from their budget (-$1,000,000 or $0) along with the addition to their rental payments (current rental value on offer or $0) are displayed for 1.5 seconds. After this the game either continues onto the next round or terminates. At the beginning of each turn the mouse cursor is reset to the midpoint between the "Buy" and "Pass" buttons.

Figure 2 shows the other two screens that subjects see during the experiment. In the event that subjects receive income they see the left panel of Figure 2 for three second before being taken automatically to the interface in Figure 1. The screen displays both the income payment, which is always $4,000,000 and the budget with which they ended the previous period. This is the only notification of income. In the case of Figure 2, for example, once the subject is returned to the main interphase, their budget would display $13,000,000.

When subjects exhaust their budget the main interphase continues to display, rental offers continue to be drawn and both the "Buy" and "Pass" buttons remain active. If subjects choose the "Buy" button, however, they are taken to the screen shown on the right panel of Figure 2. The notice displays for 1.5 second before subjects are taken back to the previous screen: the game *does not* continue. When credit constrained, subjects can only move on from their current round by pressing the “Pass” button.

Our experimental setup is designed to make it as easy as possible for subjects to enact the policy of a fully-informed agent with rational expectations. We provide agents with all of the necessary information to make their choice and make the two relevant variables – Rent and Budget – the same color and size. By resetting the cursor between the “Buy” and “Pass” buttons we remove any external costs of switching actions from period to period. We also allow our subjects as much time per round as they want to remove any possible external information processing constraints. Furthermore there are no time-keeping displays on the screen that might highlight to subjects the time it takes them to complete each round.

The main virtue of our experimental setup is that we can minimize the additional assumptions that are often needed to test hypotheses using field data, particularly in macroeconomics. We do not need to rely on conjectures about the objectives of our subjects, the information that is available to them, or the nature of the data-generating process. First, we have designed a game with an objective that rules out the need to condition on past choices. Second, our implementation of the experiment minimizes any external frictions that could impose additional costs or constraints beyond the rules of the investment game. This gives us confidence that problem misspecification on our part cannot account for possible deviations in our data from the full-information rational expectations policy. We can therefore interpret any such deviations as coming from cognitive constraints.

3 Agent's Problem

Consider an agent starting turn t with wealth $\tilde{x}_t$. She must choose sequences of actions $c \in \{0, 1\}$ contingent on each possible history of wealth and rental values that she may face in subsequent periods $t + k$. Formally, she seeks to maximize the following objective:

$$\max_{\{c_r(x_k)\}} \mathrm{E}_{f(x_t,r|\tilde{x}_t)}[rc_r(x_t)+\delta\mathrm{E}_{f(x_{t+1},r|\tilde{x}_{t+1})}[rc_r(x_{t+1})+...$$
$$+\ \delta^k\mathrm{E}_{f(x_{t+k},r|\tilde{x}_{t+k})}[rc_r(x_{t+k})+... \tag{1}$$

where r is the rental value and the expectations at $t+k$ are taken conditional on each possible value of wealth at the end of period $t+k-1$, $\tilde{x}_{t+k}$. Because the rental draws are *iid*, the marginal distribution of r does not depend on k. In fact, the joint distribution $f(x_{t+k}, r|\tilde{x}_{t+k})$ is the product of the marginals. The distribution of $x_{t+k}|\tilde{x}_{t+k}$ is $\tilde{x}_{t+k}+4$ w.p. α and $\tilde{x}_{t+k}$ w.p. $1-\alpha$. The sequence of conditional distributions is stationary $f_k(r', x_{t+k}) = f(r', y') = \Pr(r=r')\Pr(y=y')$. Note that to economize on zeros we have deflated wealth by 1,000,000 so that income, y is either zero or four. This also helps economize on variables since it implies that the price of each property is normalized to one.

Her contingent plans are subject to the budget and no-borrowing constraints, both of which must be satisfied in every period $t+k$:

$$\tilde{x}_{t+k} = \tilde{x}_t + y - c_t \quad \forall c \in \{0,1\}, \quad \forall y \in \{0,4\} \qquad \text{(BC)} \tag{2}$$
$$\tilde{x}_{t+k} \geq 0 \qquad \text{(CC)} \tag{3}$$

Equation (1) looks like it can be decomposed into the value of the contingent choices in the current period t and the continuation value given by $\tilde{x}_{t+1}$. As in a standard dynamic consumption problem, we can rewrite the problem above as a bellman equation. We denote the value of each pair (x, r) by

$$\tilde{V}(x,r) = \begin{cases} \max_{c\in\{0,1\}} cr + \delta V(x-c) & \text{if } x > 0 \\ \delta V(x) & \text{if } x = 0 \end{cases} \tag{4}$$

And denote the value of each level of wealth $\tilde{x}$ as

$$V(\tilde{x}) = \mathrm{E}_{f(y,r)}\left[\tilde{V}(x,r)\right] \tag{5}$$

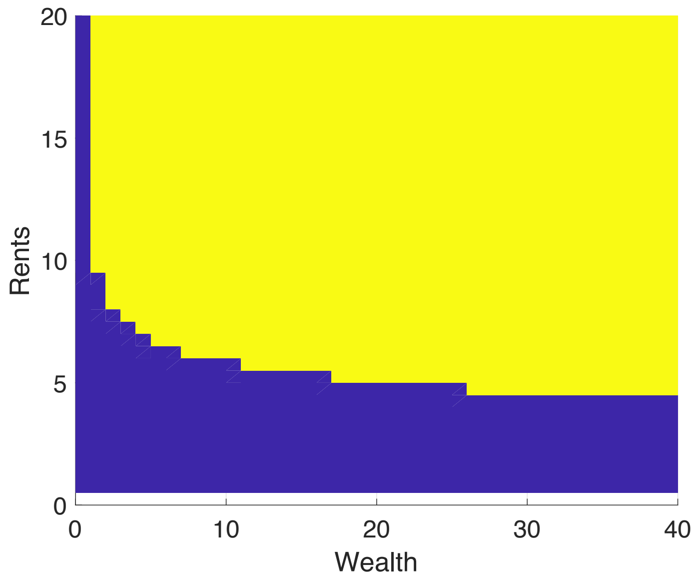

Figure 3: Threshold function $s(x)$. Agents buy in the yellow region and pass in the dark purple region.

Although we have described the problem as one of making a binary choice for each tuple (x, r) the solution to the agent's problem is a *threshold* policy $s(x)$ that depends only on x. Consider once again the agent starting the period with $\tilde{x}_t$. She does not need to wait until a property offer materializes in order to formulate her plan of action. Instead, for each level of wealth x she sets a reservation offer $s(x)$; she buys if and only if $r \geq s(x)$. Figure 3 shows the threshold policy. Higher wealth puts subjects farther away from the credit constraint so it also makes them more willing to accept lower offers.

We can write the problem directly as a choice of threshold. For tractability, the distribution from which we draw rental offers in the experiment is a discretized log-normal. Suppose we had instead allowed for the rental offers to be drawn from the positive reals. In that case, for every level of wealth x, the agent would choose a *unique* threshold s, similar to choosing a unique level of consumption for every level of wealth in a consumption/savings problem. The bellman equation for the problem of choosing $s(x)$ can be written:

$$V(\tilde{x}) = \mathrm{E}_{f(y)} \max_s \sum_{r=s(x)}^{\bar{r}} \Pr(r) r \\ + \delta \left(\sum_{r=\underline{r}}^{s(x)} \Pr(r) V(x+y) + \sum_{r=s(x)}^{\bar{r}} \Pr(r) V(x+y-1) \right)$$

In other words the agent is choosing an expected rent $u(s) \equiv \sum_s^{\bar{r}} \Pr(r) r$ along with its associated continuation value:

$$V(\tilde{x}) = \mathrm{E}_{f(y)} \left\{ \max_s u(s) + \delta \mathrm{E}_{f(r)} V \left(\tilde{x} + y - \mathbb{1}_{\{r \geq s\}} \right) \right\} \tag{6}$$

Even though our experiment requires to make a simple binary choice, the bellman equation above shows how this problem can be interpreted as an analogue to the familiar consumption/savings problem usually studied in macroeconomics. In our case, however, our subjects solve for savings rather than consumption. $V(\cdot)$ is a strong contraction and we can solve for $s(x)$ through value function iteration. Just as consumption is increasing in wealth, $s(\cdot)$ is weakly decreasing in x. Put another way, the probability of buying conditional only on r is increasing in r. Similarly, the probability of buying conditional only on x is also increasing in x.

Crucially, (x, r) are not only necessary, but also sufficient statistics for the agent to implement her optimal policy. In period t all the agent has to do is check whether $r_t \geq s(x_t)$. Previous choices only affect the current value of the problem through x_t. And previous offers contain no information about future offers. Our experimental design also ensures that we are not missing potential costs of switching actions that are not part of the game but may be present in the lab. Since the cursor is reset between the "buy" and "pass" buttons every period, there can be no motive to economize on switching the cursor from one button to the other. This result gives us a straightforward testable prediction: the probability of buying in any given period should depend only on (x, r).

4 Results

We collected data for 24 subjects. Each subject played on average just under three games, 69 games in all. [1]We have a total of 27,638 observations. Because the length of the game is stochastic, some subjects played only one game while the maximum number of games played by a subject was four. If subjects played more than one game then one of the games was chosen at random for payment and this was explained to the subjects. On top of the earnings from playing the game, all subjects received $10 for showing up. We have several hundreds of data for each subject. The subject with the fewest data played 839 rounds over four games while the subject with the most data played 2,762 rounds over two games. We do not believe that fatigue is a major concern because although subjects played many hundreds of rounds, each round took subjects on average only 1.3 seconds to complete with a standard deviation of 2.35 seconds.

Figure 4 shows subject 20's probability of buying conditional on (x, r). The black line corresponds to the rational expectations threshold $s(x)$. Even though subject 20 does not buy only and always above the line the probability of buying is higher above the line than below it for each level of wealth x. This pattern is typical of our subjects and translates into our aggregate data. Since we observe stochasticity at the subject level we interpret stochasticity in our pooled data as a feature of behavior rather than as the result of aggregation. Another salient feature of Figure 4 is how little of the state space is actually covered. Although we have a lot of data, our state space is countably infinite. The maximum rental draw in our experiment was 36,000 and the maximum wealth level was 87 million. Even if we were to truncate the state space at these values, it would still have a cardinality of 6,336. Which, if uniformly distributed, would correspond to an average of four observations per (x, r) pair. Of course they are not evenly distributed which explains the white spaces in Figures 4 and 5. Due to the size of our state space, looking at aggregate data has the benefit of allowing for more precise estimates of the policy function and adds power to our formal hypotheses tests. All our results and figures are for our pooled data unless

[1]One game ended after a single round

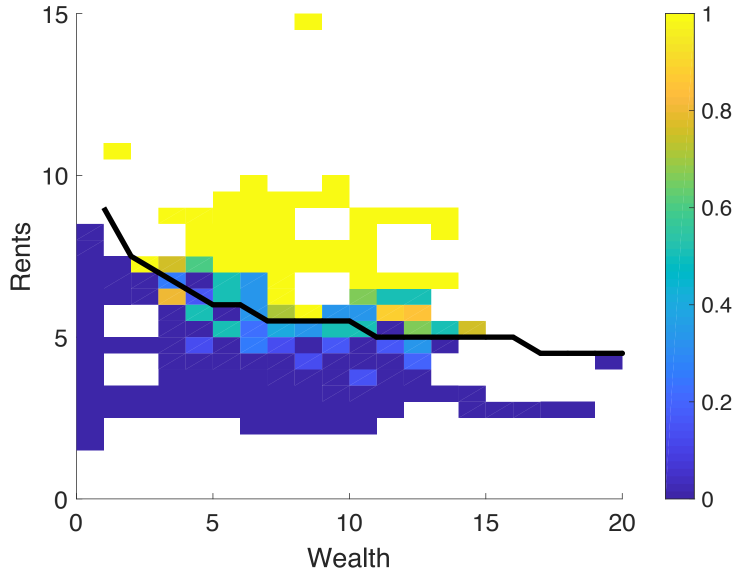

Figure 4: Empirical probability of buying conditional on the state (x, r) for subject 20. Rents are in thousands and wealth in millions. The threshold function $s(x)$ is in black. Probability scale on the right.

otherwise specified. We will revisit our data subject by subject as a robustness check of our main result in section 4.3.

4.1 Subjects Condition on the State

Before showing that our subjects condition their current choices on past choices we first show that they also respond to (x, r). Without this feature we would not be able to rule out that our choice data are simply noise. As a first pass we can compare the average earnings of our subjects with how much they would have earned under three alternative state-independent rules of thumb. On average our subjects earned \$12.63. If our subjects had used a fair coin to decide their action they would have earned on average only \$9.31. If they had naively tried to buy every property they were offered while not being credit constrained, average earnings would have been \$9.65. Finally we can look at how our subjects would have done if they had randomized using their actual unconditional probability of buying. In that case, on average, they

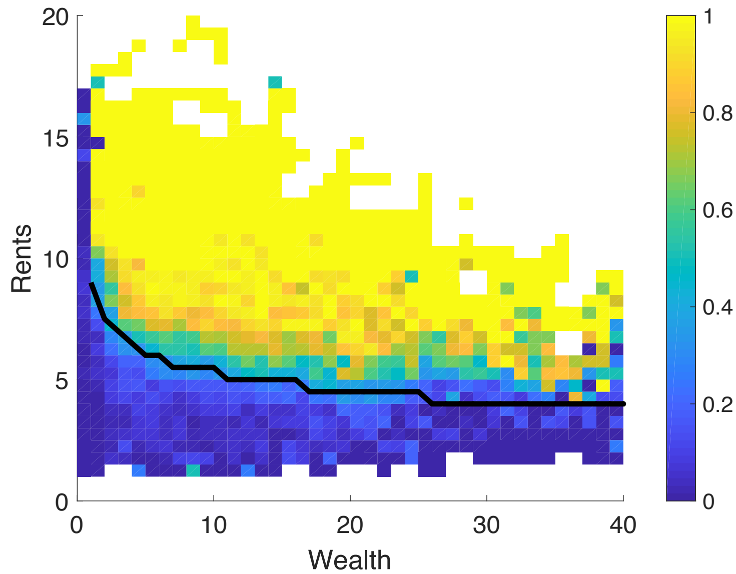

Figure 5: Empirical probability of buying conditional on the state (x, r) for all subjects. Rents are in thousands, wealth in millions. The threshold function $s(x)$ is in black.

would have earned \$8.68. The fact that our subjects earn so much more than if they bought with the same unconditional probability of buying suggests that they were conditioning on the state rather than just clicking away.

Figure 5 shows the probability of buying conditional on (x, r) with the threshold function $s(x)$ in black. Although we have truncated the figure at $x = 40$ million and $r = 20$ thousand, 95 percent of our data fall within these bounds. Outside of these bounds the conditional probability estimates are quite noisy and not very informative. Our data appear to be most noisy around the rational expectations threshold; although our subjects cannot implement the full rational expectations policy, they seem to have a fairly sophisticated understanding of the incentive structure of the game.

Figure 6 depicts more clearly the degree of sophistication of our subjects; panel (a) displays the probability of buying conditional on rents while panel (b) displays the probability of buying conditional on wealth. The probability of buying is increasing in both rents and wealth, respectively. There is also a pattern to the deviations from

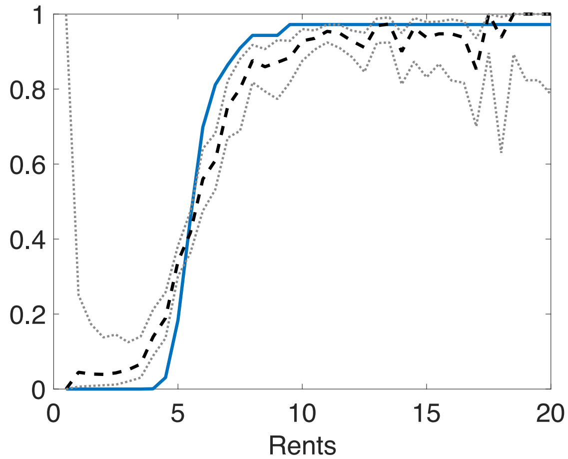

(a) Prob(Buy$|r$)

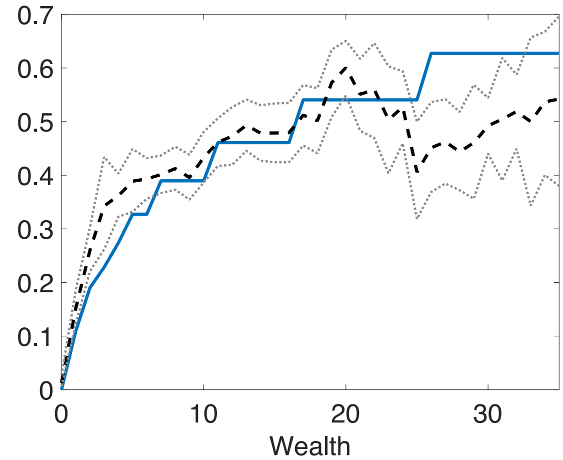

(b) Prob(Buy$|x$)

Figure 6: RE probability (solid blue); empirical probability (dashed black). 95 percent confidence intervals (dotted gray) are clustered by subject. Rents are in thousands, wealth in millions.

the rational expectations solution: subjects tend to overbuy for low levels of wealth and rents, but underbuy for high levels of both. Our data display more sophistication than simple "rule of thumb" behavior on the part of our subjects. We can rule out, for example, a single threshold s independent of x. This would imply a vertical line in Figure 6(a), a hypothesis that is clearly rejected by the data.

It is worth highlighting the fact that the probabilities implied by the rational expectations policy (in blue) are *below* the 95 percent confidence intervals (dotted gray) for low values of both r and x and similarly *above* for high values. This implies that the empirical and theoretical probabilities must cross; they do so around the means of each distribution, which are 5,600 for rents and 11 million for wealth. Our data display sophistication in the sense that the probability of buying not only *increases* as the state increases but the degree to which it does means that the *level* remains tethered to the one implied by the rational expectations policy. Subjects do not systematically underbuy or overbuy for all levels of wealth or all rental values.

4.2 Subjects Condition on Past Choices

We begin discussion of our main empirical result by considering the odds ratio. We first focus on the odds ratio because it allows for a more straightforward interpretation of the degree of inertia than regression coefficients. We of course follow this discussion with formal hypotheses testing of inertia in our data.

Let $\pi(\cdot)$ denote the probability of buying. Our analysis of the problem in Section 3 showed that π should depend *only* on (x, r); we can actually write the rational expectations policy as $\pi\left(r \geq s(x)\right) = 1$. The odds ratio, in turn, is defined as

$$\rho \equiv \frac{\pi}{1-\pi}$$

Like π, ρ should only depend on (x, r). In particular, it should be independent of the past action, c_{-1}. Similarly to Mayskova *et al* (2018), we say that our subjects condition on past choices if $\rho(x, r, c_{-1}) \neq \rho(x, r)$. While we find that ρ *does* depend on the state as implied by Figures 5 and 6, we find it *also* depends on c_{-1}. Moreover, the conditioning is *habitual*. The odds of buying go up if subjects have bought in the previous period: they display inertia.

Since rental draws are exogenous we can integrate over r and, abusing notation, denote the odds ratio in terms of wealth and the previous choice $\rho(x, c_{-1}) = \mathrm{E}_{f(r)}[\rho(x, r, c_{-1})]$. Figure 7 shows the increase in the odds of buying when subjects have bought in the previous period for wealth levels ranging between zero and 20. In any turn t, wealth is of course correlated with previous actions. However, c_{-1} contains no additional relevant information beyond x. So for any given level of wealth, the odds of buying should not change given the previous action. Instead, we see that for the 21 levels of wealth in figure 7 the point estimates (dashed black line) are always above zero and nine of them significantly so (red dots). On average, the odds of buying increase by around 50 percent if the previous action was "buy."

It is important to highlight the fact that this result rules out the interpretation that our subjects are simply behaving as if they were implementing $s(x)$ plus some noise. If subjects were only able to implement a stochastic choice rule there is still no reason why it would depend on anything other than (x, r), even if they internalize

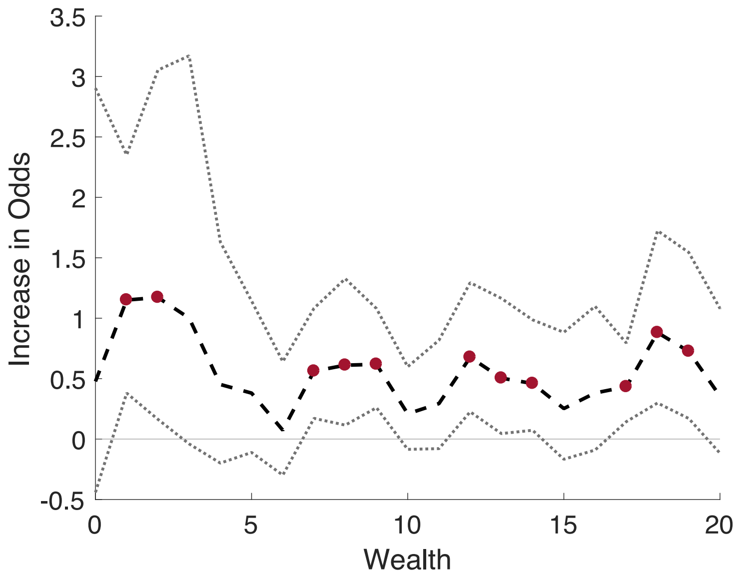

Figure 7: Increase in odds ratio. Wealth in millions. Red points denote significance at a five percent level. Bonferroni correction applied to hypotheses tests.

the stochasticity of their choice rule when designing a policy.

To test this formally we have to estimate a dynamic probit with unobserved effects. Formally, the non-linear model we would like to estimate is:

$$\Pr(c = 1) = G(r, x, c_{-1}, d^i) \tag{7}$$

where d^i is a subject fixed effect indexed by i. In linear models, the incidental parameters problem that arises from the fact that the coefficients on d^i cannot be estimated consistently is addressed by differencing out the fixed effects. In non-linear models we need to integrate over the d^i which requires a way to deal with the initial observation $c_{t=1}$. This is known as the initial conditions problem.

To address this problem we follow Wooldridge (2010) and estimate a random effects augmented probit that includes as regressors the initial choice of each subject, i, $c^i_{t=1}$ for every t. Our specification differs from Wooldridge in that he proposes including every pair (x^i_t, r^i_t) as a regressor. A large T makes this is infeasible since it would require a vector of regressors $z^i = (x^i_t, r^i_t)^T_{t=1}$ for each subject.[2] Instead we include leads and lags of (x^i_t, r^i_t). As an alternative we also draw a sample $\tilde{z}^i_\tau$ of fixed length $\tau = 40$ from z^i and include those as regressors. Finally, as further check we replace the initial condition $c^i_{t=1}$ with the average condition $\bar{c}^i$ as proposed by Chamberlain (1980). All our specifications include time polynomials. This allows us to control for learning and fatigue and ensures that the inertia we observe in the data is not driven by these—very plausible—factors. None of these alternative specifications affect our conclusions.

We estimate the following random effects probit:

$$\Pr(c = 1) = G(\beta_0 + \beta_1 r + \beta_2 x + \beta_3 x^2 + \beta_4 c_{-1} + controls) \tag{8}$$

where *controls* accounts for the incidental regressors discussed in the previous paragraph. The choice involves subjects weighting, however they may go about this process, an additional payment r and the *value* of reducing their wealth by one against the *value* of their current wealth. Since the value function is concave we include quadratic terms of wealth. We find that third and higher order polynomial terms of wealth are not significant and therefore do not report estimates of those specifications.

Table 1 reports our results. The headings in bold refer to the number of lags and leads of (x, r). Specification (A) follows Wooldridge and includes the initial condition $c^i_{t=1}$ while specification (B) follows Chamberlain and include the average condition $\bar{c}^i$. The first row reports our estimates of the coefficient on c_{-1} (β_4). The first thing to notice is that our point estimates are robust to either our choice of specification (A) or (B) as well as the lag and lead length. The average over our estimates is 0.840 with the smallest being 0.835 and the largest 0.844. Standard errors are clustered by subject and all six estimates are significant at the .001 percent level.

We do not report a table for the specification were we include the subsample $\tilde{z}^i_\tau$

[2] A large T also ameliorates the worry that the initial choice c_1 is correlated with the majority of the subsequent choices. Like the incidental parameters problem, the initial conditions problem is acute when T is small but less severe as T increases. Given our very large T, we are confident our results are not driven by this problem.

	4 Lags and Leads		**10 Lags and Leads**		**20 Lags and Leads**	
	(A)	(B)	(A)	(B)	(A)	(B)
Lag Action	0.844	0.843	0.840	0.838	0.837	0.835
	(0.188)	(0.188)	(0.186)	(0.186)	(0.200)	(0.199)
Rent	64.709	64.700	66.108	66.099	67.547	67.554
	(9.005)	(9.030)	(9.198)	(9.209)	(9.616)	(9.626)
Wealth	0.667	0.667	0.670	0.671	0.682	0.683
	(0.056)	(0.056)	(0.060)	(0.060)	(0.060)	(0.060)
$Wealth^2$	-0.004*	-0.004*	-0.004*	-0.004*	-0.004*	-0.004*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Cons	-4.671	-5.533	-4.246	-5.220	-4.030	-5.239
	(0.543)	(0.746)	(0.540)	(0.731)	(0.529)	(0.761)
N	27152	27152	26348	26348	25103	25103

$^*p < 0.05$, $p < 0.001$ otherwise. Standard errors clustered by subject.

Table 1: Panel Data

instead of leads and lags since there is no single regression. Instead we draw 100 subsamples $\tilde{z}^i_\tau$ and run a regression for each. The average estimate $\hat{\beta}_4$ is 0.85 with a 95 percent confidence interval [.76 .93].

Our regression analysis also confirms that our subjects respond to the state. The probability of buying is increasing in r and we also find evidence that the *subjective* value of wealth is indeed concave with a “large” positive coefficient on x and “small” negative coefficient on x^2. We temper this finding by pointing out that the coefficients on x^2 are significant only at a five percent level. Similar to β_4, our estimates of the other coefficients are significant and invariant to the regression specification.

One of the challenges of estimating non-linear regressions is the *economic* interpretation of the parameters. We do not intend to give a structural interpretation to the probit model in (8). The aim of the formal statistical analysis is to get qualitative rather than quantitative results. The meaningful result is that inertia is a feature of choice. Quantifying the degree of inertia is left for future research. We are thus more interested in the *signs* of our estimates than in their values. The main takeaway from Table 1 is that it confirms the claims of the two subsections thus far. First, the coefficients on rent and wealth are all positive at a significance level of .001: *subjects condition on the state*. Second, and the central claim of this paper, the coefficient on the lag action is positive and significant at a .001 level: *subjects condition on past actions.*

4.3 Robustness Checks

To check the severity of the bias in our estimates we conduct Monte Carlo simulations on the counterfactual were subjects do not condition on c_{-1}. For each of our subjects we estimate a policy function $\pi^i = g(\beta_0^i + \beta_1^i r + \beta_2^i x + \beta_3^i x^2)$. We do not have to worry about the size of our 24 subsets of data since each contains on average more than 1000 observations. By construction, these estimated choice rules do not condition on past actions.

For each of the 69 games played we draw one of the choice rules from $\{\pi^i\}_{i=1}^{24}$ with replacement and generate a new history of play. So for each of the games the vectors of exogenous variables—income, rental and continuation—remain the same and only the wealth and choices are simulated. This generates a dataset with the same number of observations as our original dataset but with simulated choice and wealth data. Importantly, our simulated choice data depend only on wealth, rents, and a subject fixed effect but not on past actions.

For each of 1200 Monte Carlo simulations we estimate regression (8) and generate a distribution of estimated parameters $\hat{\beta}_4$.[3] Figure 8 shows the kernel density estimate of $\hat{\beta}_4$. Although the estimates are upward biased, the average bias of .075 is an order of magnitude smaller than the estimates reported in Table 1. The 95 percent confidence interval ranges from .01 to .14. The black vertical line in figure 8 intersects the x-axis at .84, the average of the $\hat{\beta}_4$ reported in Table 1. The standard deviation of the kernel density is .036, which puts the black line nearly 20 standard deviations away from the upper boundary of the confidence interval. Even after we

[3]Given the stability of our estimates we only report the specification with 4 lags and leads.

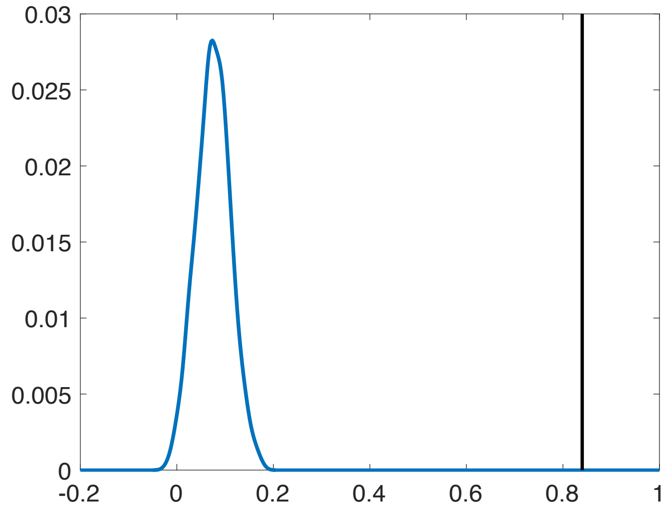

Figure 8: Kernel Density of $\hat{\beta}_4$ (blue). Average value of $\hat{\beta}_4$ from table 1 (black).

apply this bias correction to our estimates they remain significant at the same .001 level.

As a second robustness check we can run time series probit regressions for each of our 24 subjects. Table 2 reports the estimates of β_4 for each of the 24 subjects. All but two of the estimates are positive. The two negative point estimates are extremely small and not significant; in both cases the *p-value* is over 50 percent. Seven of the remaining 22 point estimates are significant at the .1 percent level and one is significant at the five percent level.[4] About a third of our subjects convincingly display inertia while we estimate a positive residual autocorrelation in actions for over 90 percent of them.

One of the drawbacks of disaggregating our data is that our state space is very large so the coefficient on the lag term is not very precisely estimated. We are interested in average behavior rather than individual heterogeneity. We only want to consider the results in 2 as a robustness check that the unobserved effects d^i are not driving the estimates we report in table 1.

We have already argued at length that our experimental design allows us to rule

[4]Since we are testing multiple hypotheses, we apply the Bonferroni correction before testing for significance.

(1)	(2)	(3)	(4)	(5)	(6)
0.454	0.614***	0.035	0.430	0.350	1.407***
(0.568)	(0.135)	(0.323)	(0.201)	(0.242)	(0.222)
(7)	(8)	(9)	(10)	(11)	(12)
1.375***	-0.156	1.998***	0.396	1.138***	-.033
(0.221)	(0.216)	(0.117)	(0.202)	(0.192)	(0.200)
(13)	(14)	(15)	(16)	(17)	(18)
1.226***	0.661	0.246	0.776***	0.109	0.604
(0.257)	(0.246)	(0.186)	(0.162)	(0.183)	(0.358)
(19)	(20)	(21)	(22)	(23)	(24)
0.322	0.133	0.501	0.583	0.369	0.776*
(0.115)	(0.214)	(0.249)	(0.237)	(0.399)	(0.253)

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Robust standard errors.

Table 2: Lag Action Coefficient by Subject.

out preferences or external costs as the sources of inertia. Having firmly established inertia as a feature of behavior, we can only conclude that it must come from cognitive biases or limitations. Our experimental design allows us to determine that the problem as laid out in equations (4)-(5) is *not* the problem that subjects are solving. Although we cannot observe cognitive processes directly, we can nonetheless use our choice data to discipline our modeling of those processes. In what follows, we will argue that inertia is best understood as a way of *economizing on cognitive resources* rather than as a behavioral bias that creates attachment to goods or actions.

5 Inattentive Reconsideration

Why would our subjects condition on past choices when x and r are freely available to them at all times? Our preferred explanation is consistent both with introspection and models of cognition from psychology and neuroscience: thinking hard about a problem is slow and requires effort. It is well known that there is no single mechanism through which the brain makes decisions.[5] Choices from deliberation, which requires a lot of cognitive resources, are slow compared to *habitual* actions. Subjects might therefore condition on previous actions as a way to economize on thinking through the problem every round. In some rounds they may choose to simply continue doing what they did before, while in others they may decide it is worth it for them to think more carefully before choosing.

The hypothesis we have sketched above introduces an interim choice. We are not proposing that the conditioning on past actions is automatic. Quite the opposite, we are proposing that subjects take account of the state when they choose whether to reconsider. Alternatively, we can interpret this interim choice as a decision whether to *continue* deliberating. What we have in mind is a subject who sees her budget and the rental offer and then decides whether to just choose as she did in the previous period or keep thinking through whether to buy or not. If she continues to think she may very well choose the same action as she has in the past. However, changes in actions necessarily imply that she has reconsidered.

We find suggestive evidence of this process in our reaction time data. Reaction times for action switches are longer than reaction times when actions remain the same. Recall that the mouse cursor is reset each turn so there are no external adjustment reasons for this time delay. We estimate a random effects regression of reaction times on an indicator of action switches

$$rt = \gamma_0 + \gamma_1 \mathbb{1}_{\{c \neq c_{-1}\}} + controls$$

Our controls are a third order polynomial for the number of turns as well as the state (x, r). Since neither wealth nor rents affect reaction times or have a meaningful effect on our estimate of γ_1 we drop them from the specification we report here. A fixed effects model yields nearly identical results.

The point estimate $\hat{\gamma}_1$ is .08, the 95 percent confidence interval, clustered by subject, is [.04 .12]. The associated *p-value* is less than .1 percent. As a comparison $\hat{\gamma}_0 = 1.72$. The interpretation is that action switches on average lead to reaction

[5]See, for example, Sanfrey and Chang (2008).

times that are 80 milliseconds longer than the 1.72 seconds that it takes to choose the same action is before.

To be sure, this is only a small delay. However, the 1.72 seconds include processes which subjects must undergo regardless of their choice, such as reading the screen and moving the mouse. It is therefore hard to directly compare time spent thinking. Reaction times are not the focus of this paper so we need to be careful not to over-interpret $\hat{\gamma}_1$. While reaction times are not a direct test of whether our subjects are "thinking more," the fact that subjects do take longer in rounds where they change their choice from the previous round is evidence in favor of the hypothesis for which we have argued above.

5.1 Model

In the following sections we formalize the story sketched out above in a model of *inattentive reconsideration*. We begin with a discussion of two important details of our modeling strategy. We then proceed to write down the model and solve for the policy function of our cognitively-constrained agents. This model follows a model proposed by Woodford (2009) to explain *pricing* behavior. A variant of this model was also used by Khaw *et al* (2017) to account for pricing-like behavior in an experimental setting. The model combines rational inattention with the two-step decision of whether to reconsider and, conditional on reconsideration, whether to buy.

First proposed by Sims (2003), the rational inattention model captures cognitive limitations in a reduced form as informational costs. Rather than conditioning on the state, agents must now condition on signals of the state. There is an information cost function that makes it costlier for signals to be more precise and therefore agents must weigh the benefit of making choices with more precise signals against the cognitive cost of those signals. Importantly, however, agents are free to choose whatever signal structure they want.

In a dynamic environment—as is the case in our experiment—one must consider whether agents would also choose to condition their actions on the history of signals they have received rather than simply the current signal. If memory itself is costly, accessing the history of signals up to time t is not free and should itself be transmitted by an informative, yet not perfectly revealing signal. But then it is better for the agent to get a more precise signal today than to devote cognitive resources on accessing the history of signals up to today. In settings where knowing the state is very difficult, it makes sense to model agents as conditioning their signals on beliefs rather than on the state itself. This is not the case in our framework. Agents do not have to guess what their current wealth is or what the rents on offer are in any given period. Both are displayed on the screen and so it would be strange for subjects to devote time and effort recalling *past* values of the given that the only reason they would do so would be to update their beliefs about the current values of the state, which they can easily read on a computer screen in front of them. We therefore find it sensible to assume that signals will be conditioned only on the current state and not on past signals. Further, as shown by Caplin and Dean (2015) and Steiner *et al* (2017), agents will only choose at most one signal for each available action. It is without loss of generality that we restrict the signal space to one of recommended actions.

Information costs generally lead to stochastic rather than deterministic choice rules. Although this is certainly in line with what we observe in our data, we have to think carefully about the implications of hitting the credit constraint. To do this we have to distinguish between actions and consequences. For the fully informed rational agent of Section 3 there is no distinction. For an inattentive agent this is also the case everywhere except when $x = 0$.

Recall that the experiment allows subjects to click “buy” even when they are credit constrained. The rational agent would never click on “buy,” unlike our subjects whose probability of trying to buy when they are credit constrained is 1.3 percent while .6 percent is the lower bound of the 95 percent confidence interval. This is the experimental equivalent behavior of having your credit card declined; it doesn’t happen often, it doesn’t happen to everyone, but it happens from time to time. We want to capture this feature of behavior in our data because it is behavior that we observe outside of the lab. To do this we decouple the action “buy” from the consequence of buying at $x = 0$ by introducing a fixed cost κ of trying to buy when credit constrained. In the context of our experiment κ can be interpreted as the wasted time that comes from not clicking "pass."

5.2 To Think or not to Think

In our model agents choose whether to reconsider and pay a fixed cost. If they do not reconsider they repeat the previous period's action. If they decide to reconsider they now choose whether to "buy" or "pass". Importantly, this choice is subject to an information cost. In addition to wealth and rents, the agent's previous choice is now relevant to her decision and becomes another state variable. The value of the triple (x, r, c_{-1}) is given by:

$$\tilde{V}(x, r, c_{-1}) = \begin{cases} \max_{\mu}(1-\mu)(rc_{-1} + \delta V(x - c_{-1}, c_{-1})) + \mu(\bar{V}(x, r) - \gamma^{rec}) - \phi^{-1}I(\mu) & \text{if } x > 0 \\ \max_{\mu}(1-\mu)(-\kappa c_{-1} + \delta V(x, c_{-1})) + \mu(\bar{V}(x, r) - \gamma^{rec}) - \phi^{-1}I(\mu) & \text{if } x = 0 \end{cases} \tag{9}$$

Equation (9) is the analogue of the rational expectations $\tilde{V}(\cdot)$ in equation (4). The value function, $V(\cdot)$ now depends on c_{-1} and, similarly, is the analogue of equation (5):

$$V(\tilde{x}, c_{-1}) = \mathrm{E}_{f(r,y)}\left[\tilde{V}(x, r, c_{-1})\right] \tag{10}$$

where $\mu(x, r, c_{-1})$ is the probability of reconsideration, ϕ parametrizes the marginal cost of information, $I(\cdot)$ is an information cost function, γ^{rec} is the fixed cost of reconsidering and $\bar{V}(\cdot)$ is the value of reconsidering. In this section we want to focus on the choice to reconsider so for now we take $\bar{V}(\cdot)$ as given. In the next section we focus on how this value is determined.

Equation (9) says that the agent must choose a probability of reconsideration for each (x, r, c_{-1}) by balancing the benefit of the more precise signal (the first terms in brackets) against the cost of more precision (the last term). If she does not reconsider she faces the value of that outcome in the first term inside the brackets: she takes her past action, c_{-1} and carries it forward into the next period. If she does reconsider, she pays a fixed cost but gains the value of reconsideration given the state, $\bar{V}(\cdot)$. Notice that the value of reconsidering does not depend on c_{-1}; the only reason for the agent to condition on c_{-1} is to avoid paying the cost γ^{rec}. Once she incurs the cost and commits her cognitive resources, she only pays attention to the payoff-relevant variables (x, r). Our agent is forward-looking, but optimally chooses to condition on c_{-1}.

The second line in equation (9) captures the cost of being credit constrained. Suppose the agent hits her credit constraint and does not reconsider. If she does not buy her payoff is zero and her continuation value is $V(0, 0)$. If she does click "buy," however, she pays a cost κ independent of the rental value on offer. Even though the agent incurs this cost if she tries to buy at $x = 0$ it does not necessarily mean that she will reconsider with probability one when $x = 0$ and $c_{-1} = 1$. She must balance this agains the cost of discriminating the state perfectly.

To operationalize our model we follow the literature of rational inattention and assume that the cost of information is the mutual information function. The mutual information function is the expectation taken with respect to the joint distribution of the action and the state of the log of likelihood ratio of the joint over the marginals. Simple algebra shows that this implies $I(\cdot)$ is of the following form:

$$I = \mu \log\left(\frac{\mu}{M}\right) + (1 - \mu) \log\left(\frac{1 - \mu}{1 - M}\right) \tag{11}$$

where μ is the conditional probability of reconsideration while M is the unconditional probability. $I(\cdot)$ captures the tradeoff of paying attention by penalizing deviations from expected behavior. There is no penalty to choosing $\mu = M$. There is only a cost if you deviate from what you do in expectation.

Taking the first order conditions we can write the solution to problem (9)-(11)

recursively as: [6]

$$\mu(x, r, c_{-1}) = \frac{M' \exp\{\phi \bar{V}(x, r)\}}{\Delta_\mu(x, r, c_{-1})} \tag{12}$$

$$\Delta_\mu(x, r, c_{-1}) = M' \exp\{\phi \bar{V}(x, r)\} + (1 - M') \exp\{\phi(rc_{-1} + \delta V(x - c_{-1}, c_{-1})\} \tag{13}$$

$$V(\tilde{x}, c_{-1}) = \frac{1}{\phi} \mathrm{E}_{f(r,y)} \left[\log \Delta_\mu(x, r, c_{-1})\right] \tag{14}$$

where $M' \equiv M \exp\{-\phi\gamma^{rec}\}/(M \exp\{-\phi\gamma^{rec}\} + [1 - M])$. Conditional on ϕ, γ^{rec} pins down M'. We can dispense with γ^{rec} and interpret M' as the reference probability of reconsideration so that it is costly to choose probabilities that differ from M' rather than M. Under that interpretation it is costly for agents to deviate from what they would *prefer* to do (reference probability M') instead of what they expect to do (unconditional probability M).

Equation (12) says that the agent rescales her reference probability of reconsidering, M', by the relative value of reconsidering given the state: $\exp\{\phi\bar{V}(\cdot)\}/\Delta_\mu(\cdot)$. This ratio is the value of discriminating the state with high precision. If this ratio is larger than one she sets her probability of reconsideration above her reference probability and vice versa. For example, if for some reason there was a pair (x, r) for which the value of reconsidering were the same as the value of not then the agent would reconsider with probability M'. Note, also, that the value of her choice is scaled by the marginal cost of information. When ϕ is high, the marginal cost of information is low and so the term in the exponent rises. Because these terms are exponentiated, ϕ amplifies the differences between the values of reconsidering versus not. This is exactly the intuition of the information cost: the lower the cost, the cheaper it is to distinguish between the two choices and the more sensitive μ is to the state rather than to the reference probability M'.

[6] We omit the case of $x = 0$ for brevity. in that case we would simply replace r in the last term of equation (13) with $-\kappa$ and c_{-1} in the first argument of the value function with zero.

5.3 To Buy or not to Buy

We now turn to the agent's choice of action conditional on her having reconsidered. We assume that this choice is also subject to an information cost. We also allow for a bias γ^{buy}; if positive it captures a bias toward spending while if negative it captures a bias toward saving. Although not crucial to our results, we include it to get cleaner estimates of the information cost parameters.

$$\bar{V}(x,r) = \begin{cases} \max_\pi \pi(r + \delta V(x-1,1) - \gamma^{buy}) + (1-\pi)\delta V(x,0) - \theta^{-1} I(\pi, \Lambda) & \text{if } x > 0 \\ \max_\pi \pi(-\kappa + \delta V(x,1) - \gamma^{buy}) + (1-\pi)\delta V(x,0) - \theta^{-1} I(\pi, \Lambda) & \text{if } x = 0 \end{cases} \tag{15}$$

where Λ is the probability of buying conditional on reconsidering and θ is the marginal cost of attention. Agents now choose the probability of buying $\pi(x,r)$ taking into account that they will condition on their action in the future. We can write the solution recursively as:

$$\pi(x,r) = \frac{\Lambda' \exp\{\theta(r + \delta V(x-1,1))\}}{\Delta_\pi(x,r)} \tag{16}$$

$$\Delta_\pi(x,r) = \Lambda' \exp\{\theta(r + \delta V(x-1,1))\} + (1-\Lambda') \exp\{\theta \delta V(x,0)\} \tag{17}$$

$$\bar{V}(x,r) = \frac{1}{\theta} \log \Delta_\pi(x,r) \tag{18}$$

where $\Lambda' \equiv \Lambda \exp\{-\theta\gamma^{buy}\}/(\Lambda \exp\{-\theta\gamma^{buy}\} + [1-\Lambda])$ has a similar interpretation as M'; it represents the reference probability of buying, conditional on reconsideration. Equations (16)-(18) should look familiar. They are the counterparts of equations (12)-(14) for a different binary choice: whether to "buy" or "pass." Equations (12)-(18) characterize the solution to the full two-stage decision problem. We can solve for the policy functions $\mu(x,r,c_{-1})$, $\pi(x,r)$ via value function iteration.

The policy equations (12)-(13) and (16)-(17) also suggest an alternative interpretation of the attention costs $I(\cdot)$. In equation (9), for example, the natural interpretation of ϕ is as the marginal cost of reducing the noise with which agents can observe the state. Similarly for θ in equation (15). This is the traditional interpretation of the Shannon information cost in the rational inattention literature. In the policy functions for $\mu(\cdot)$ and $\pi(\cdot)$, however, it is more natural to interpret ϕ and θ as governing the ease of distinguishing the values associated with each action. Under the first view, the cost function captures an inability to process inputs (the value of the state) without adding noise; under the latter—and, we would argue, more appealing—view, the cost function captures an inability to process the outputs (the value of the action) without noise. This is a story in which agents cannot always tell which action is more valuable, even if they know the state perfectly. This is why we argue that ours is a reduced-form model of cognitive constraints rather than a model of costly information acquisition.

We should underscore how this model is consistent with the evidence on reaction times. According to (9) action switches must be the result of reconsideration. Although we cannot observe agents' interim decision, we can isolate cases when we know the outcome of that decision. And in those cases, we see evidence of deliberation in the form of slower reaction times.

5.4 Estimating the Model

We estimate the five parameters of our model jointly via maximum likelihood. We report the estimates in Table 3. Since the estimates are in the arbitrary units of rental values in our experiment we convert them to U.S. cents. For example, the estimate of the cost of buying when credit constrained, $\hat{\kappa}$, is .012 or just over one hundredth of a cent.

The estimated reference probability of reconsidering, $\hat{M}'$ is less than one, which implies there is a cognitive cost of thinking. If γ^{rec} were zero the reference and unconditional probabilities of reconsideration would both be equal to one. The intuition is that by reconsidering subjects now get the value $\bar{V}(\cdot)$ without having to pay any cost γ^{rec}. We know that $\bar{V}(\cdot)$ is at least as high as the value of not reconsidering because, upon reconsideration, agents can always choose the same action as before.

Table 3: Estimated Parameters

$\hat{\theta}$	$\hat{\phi}$	$\hat{\Lambda}'$	$\hat{M}'$	$\hat{\kappa}$
114.3	49.65	.103	.878	.012

Therefore agents always reconsider. Given our empirical results in Section (4), it is unsurprising that the estimates of our structural model imply that agents condition on past actions with positive probability.

The estimated information cost of reconsidering $\hat{\phi}^{-1}$ is just over twice as large as the estimated information cost of choosing whether to buy $\hat{\theta}^{-1}$. Remarkably, this is roughly the same ratio that Khaw *et al* (2017) find in their experimental data as well as Stevens (2015) in her parametrization of micro price data. Three seemingly unrelated datasets generated by different subjects in different settings have now all estimated that the cognitive cost of whether to reconsider is around twice as high as the cognitive cost of what to do upon reconsidering. That the estimate of this ratio remain stable across these three different decision problems executed by different people is significant. It suggests that this model *is* capturing part of the cognitive process that culminates in choice. This points to an avenue for further research to more rigorously measure the ratio of these two attention costs.

The estimated reference probability of buying upon reconsidering is only 10 percent. This is low relative to the unconditional probability of buying in our data which is 40.3 percent. This implies that subjects are conservative in the sense that they are slightly biased toward not buying. This conservatism, however, does not imply a counterfactually low probability of buying. The unconditional probability of buying implied by our model is 39.89 percent. As a reference, the unconditional probability of buying had our subjects implemented the rational expectations policy is 40.4 percent.

Overall, our model is successful in capturing the main features of our pooled data. Figure 9 contrasts the empirical probability of buying conditional on (x, r) in panel (b) with the probability of buying conditional on (x, r) implied by our estimated model in panel (a). Our estimated model matches the data particularly well for

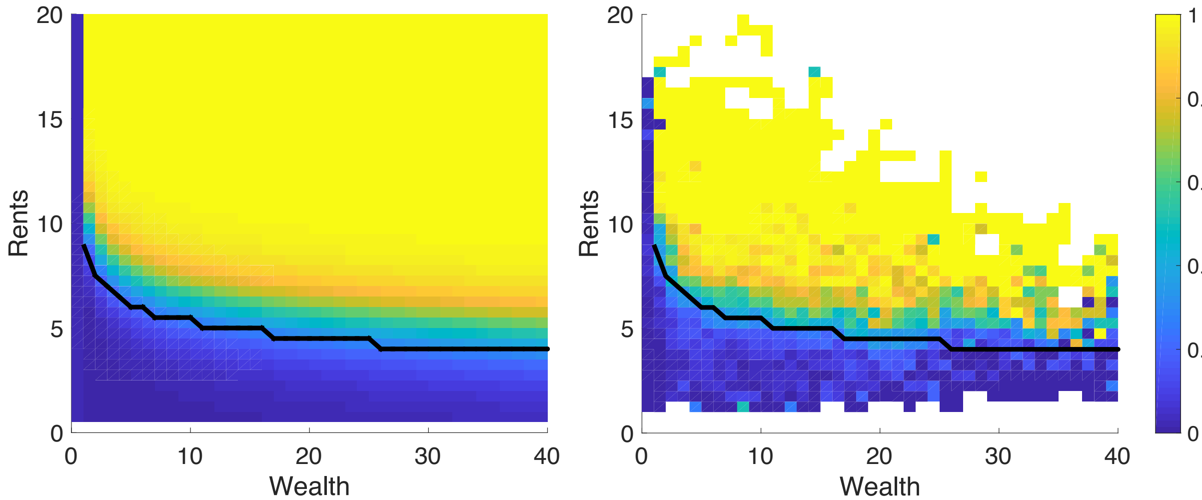

(a) Model (b) Empirical

Figure 9: Probability of buying conditional on wealth and rents. Threshold function $s(x)$ in black. Probability scale on the right.

levels of wealth up to around 25 million. At higher levels of wealth the empirical probability of buying properties with rents at or below 5000 is very close to zero whereas in our model they are closer to 40 percent. This, however, is likely due from fewer data; 88.9 percent of all our observations occur for wealth levels at or below 25 million.

To compare the degree of inertia implied by our model to the one we estimate from the data we simulate 1200 datasets using the estimated policy functions of the model $\hat{\mu}(\cdot)$ and $\hat{\pi}(\cdot)$. These simulations take the histories $\{r_t^i, y_t^i\}$ as given so that only wealth and actions differ from our real dataset. For each of the 1200 simulations we estimate the increase in the odds ratio $\Delta\rho(x, c_{-1})$ for all levels of wealth in the same way as we did to construct Figure 7. This yields a sequence of distributions $(\Delta\hat{\rho} - \Delta\rho)(x, c_{-1})$ indexed by x where $\Delta\hat{\rho}$ is the estimated increase in the odds shown in figure 7. Figure 10 shows the mean (dashed black line) along with the 95 percent confidence interval (dotted gray) of these distributions. Our model captures the degree of inertia in the data extremely well. We cannot reject the hypothesis

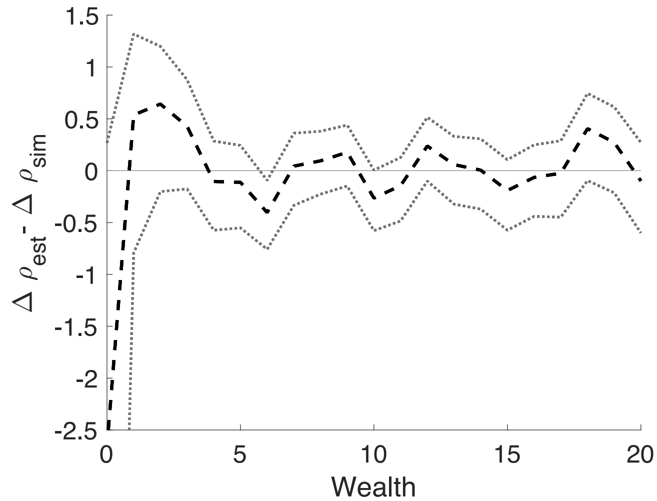

Figure 10: Difference between estimated increase in odds and simulated increase in odds. Average difference in dashed black. 95 percent confidence intervals in dotted gray. Wealth scale in millions.

that the degree of inertia implied by our model is different from that in our data. Only at $x = 0$ is the increase in the odds predicted by the model qualitatively higher from the one implied by our data. The level of uncertainty around this estimate is also high, however, and thus not statistically significant. At $x = 0$ our model, on average, implies a much higher increase in the odds of buying than for other levels of x. This is due to the credit constraint. As shown in Figure 11(b), the probability of reconsidering at the credit constraint is actually relatively low because the estimated costs of buying when credit constrained $\hat{\kappa}$ are not very high. Though the average difference is high, we cannot reject the null that they imply the same degree of inertia.

Beyond matching the *degree* of inertia in our data, our model implies that inertia is *state dependent*. Figure 11 displays the estimated probability of reconsideration conditional on the state and past actions $\hat{\mu}(x, r, c_{-1})$. Panel (a) shows the probability of reconsideration when the past action is "pass" whereas in panel (b) the past action is "buy." Unsurprisingly, this probability is nearly one for very high rents when the previous action is "pass" and for very low rents when the previous action is "buy."

In macroeconomics we are especially interested in how inertia changes with the

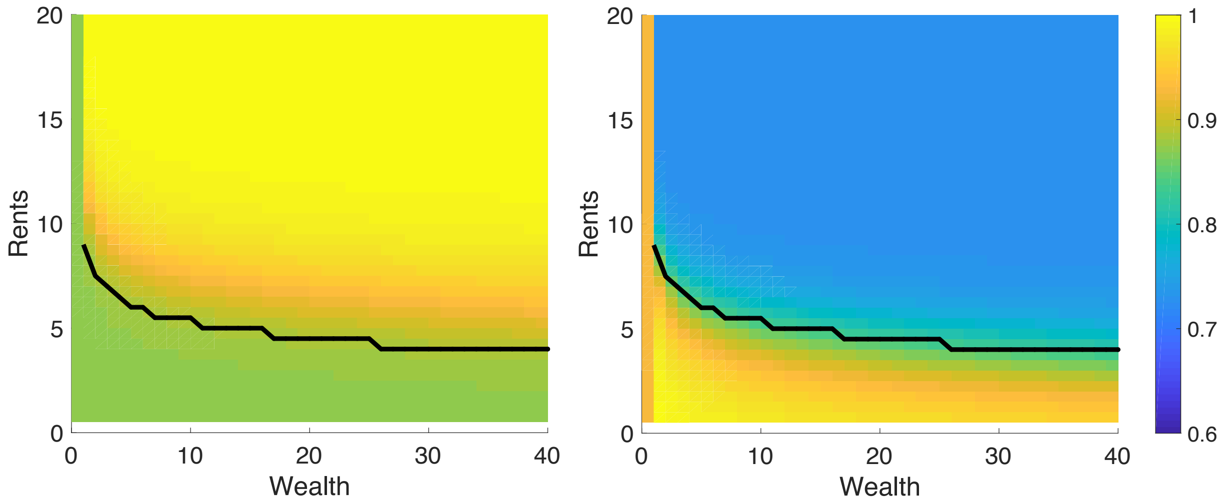

(a) $\hat{\mu}(x, r, 0)$ (b) $\hat{\mu}(x, r, 1)$

Figure 11: Probability of Reconsideration. Panel (a): previous action is "pass". Panel (b): previous action is "buy". Rents are in thousands, wealth in millions. Threshold function $s(x)$ in black. Probability scale on the right.

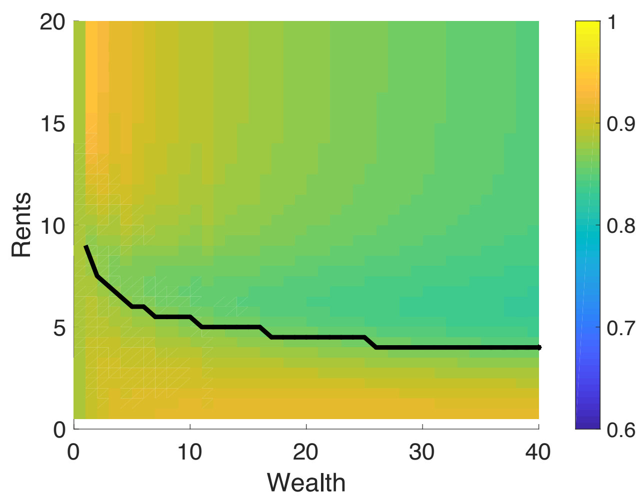

Figure 12: Probability of reconsidering conditional on wealth and rents only: $\hat{\mu}(x, r)$. Rents are in thousands, wealth in millions. Threshold function $s(x)$ in black. Probability scale on the right.

state. Figure 12 shows the probability of reconsidering conditional only on (x, r). The probability of reconsidering is highest close to the axes, forming a kind of interrupted L shape. The intuition is simple: the probability of reconsidering is highest where action switches are most likely to be valuable. For low levels of wealth agents are more likely to "pass" than to "buy." So when rare, but very high rental offers materialize, agents' expected losses from not reconsidering are high. These are states worth being able to distinguish precisely—hence the high probability of reconsideration. A similar argument applies for the lowest rental offers. These are properties that the agent would not want to buy, regardless of her level of wealth. Therefore it is also valuable for the agent to distinguish these states precisely so that she can make sure not to buy these properties when they go on offer. Inertia is thus highest in the non-extreme regions of the state space, which are also the most likely. Recall that the average rent is around 5,600; it is just above this value, especially as wealth increases, that we see the highest degree of inertia.

This contrasts sharply with standard habits models. In those models the stock of habits is predetermined so inertia is independent of the state. The result is an inelastic term in the household's consumption rule. Our model has no such term. The deeper habits that we propose will be strongest when the gains from reconsideration are small; but they will almost disappear when changes in the state make it too costly to act habitually. This has an important implication for the propagation mechanism of aggregate shocks. In our model, the larger the shock, the more responsive aggregate demand will be to that shock. This suggests, for example, that impulse responses with non-time-separable preferences overstate the degree of inertia in aggregate demand to large and infrequent shocks.

Figures 11 and 12 also show that the level of inertia implied by our parameter estimates is not too severe. The probability of reconsideration ranges from .729 to .996. This is not surprising given the environment in which our subjects completed the task. They had no outside distractions competing for their attention and could see both x and r on the screen at all times. It seems sensible that most of the time they focused on the problem and occasionally found it preferable to act "habitually." The unconditional probability of not reconsidering is 12.54 percent. And yet even

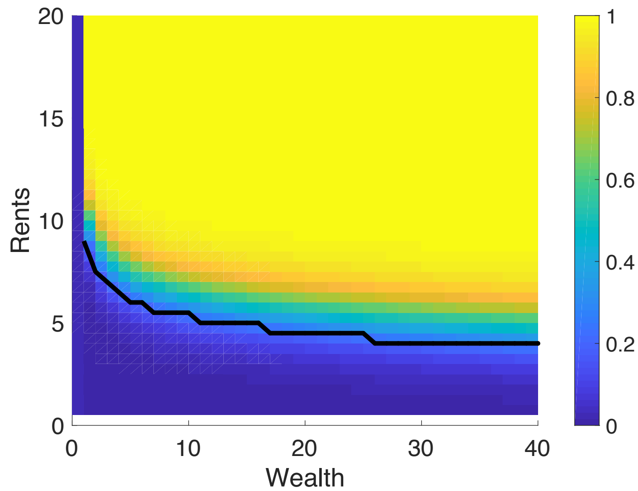

Figure 13: Probability of buying upon reconsideration $\hat{\pi}(x, r)$. Rents are in thousands, wealth in millions. Threshold function $s(x)$ in black. Probability scale on the right.

with this modest degree of inertia Figure 7 shows that the odds of buying go up by around 50 percent if the subject bought in the previous period.

The high probability of reconsideration also means that the bulk of stochasticity in our data is explained by the probability of buying conditional on reconsideration. Figure 13 shows this probability, $\hat{\pi}(x, r)$. Qualitatively it looks very similar to the probability of buying conditioning on only the state (x, r) as shown in Figure 9. This highlights the importance of information costs at both stages of choice. The information cost is higher for the reconsideration choice, and so agents do not distinguish among states too finely. The lower the attention costs, the wider the range of $\mu(\cdot)$. In the limit, if attention costs were taken to zero, $\mu(\cdot)$ would become a threshold; agents would reconsider if and only if the value of reconsideration exceeded the fixed cost. By contrast, in the case of the buying decision, where we estimate a lower attention cost, we see that the probabilities vary much more with the state, ranging from close to zero to close to one.

5.5 Costs of Inertia

It is difficult to understand the magnitude of the attention costs by looking at $\hat{\phi}$ and $\hat{\theta}$ alone. To get a better sense for two costs we decompose the gap between the average earnings from behaving according to our estimated policy functions and the rational expectations counterfactual. The average earnings implied by our model are \$12.63, which is exactly the same as the actual average earnings of our subjects. The average earnings from implementing the rational expectations policy would have been \$13.32. This is a small gap in dollar terms, but amounts to over five percent of earnings.

Our objective is to decompose this 69 cent gap into the costs of reconsideration and the costs of attention once subjects decide to reconsider. For clarity, we will refer to the first as the cost of inertia and the second as the cost of inattention. To achieve this decomposition we solve for the policy functions of two auxiliary models. In both models we use the parameter estimates reported in Table 3 while shutting down one of the two types of cognitive costs present in our full model.

In the first auxiliary model we get rid of inertia so that our subjects reconsider every period. This is equivalent to setting $M' = 1$, which implies $\gamma^{rec} = 0$. This counterfactual answers the question: how much more would our subjects have earned *given* $[\hat{\theta}, \hat{\Lambda}', \hat{\kappa}]$ if they did not have any reconsideration costs?

The second model shuts down the inattention costs upon reconsideration. This is equivalent to $\theta \to \infty$. In this case subjects do not reconsider every period, but when they do they behave according to a threshold policy. Importantly, this will not be the same as $s(x)$ since subjects now take into account the fact that they may be locked into this choice in future periods. This counterfactual asks: how much more would our subjects have earned *given* $[\hat{\phi}, \hat{M}', \hat{\kappa}]$ if they did not have any inattention costs?

For each auxiliary model we run 1200 Monte Carlo simulations to estimate average earnings.[7] We estimate that 13 out of the 69 cents are due to inertia while the remaining 56 come from insensitivity. In other words, around one fifth of the gap

[7] These are done in the same exact way as we have already described in previous sections.

between the rational expectations payoff and the actual payoff can be attributed to inertia while four fifths can be attributed to inattention. Inertial behavior cost our subjects around one percent of expected earnings. This amount is small, but significant, especially considering that it comes from cognitive costs *alone*. To put it in perspective, one percent is the the average growth rate of real consumption expenditures in the United States.

6 Alternative Explanations

We conclude with a discussion of alternative explanations for our data. Since alternative models have different numbers of parameters, we use the Bayes Information Criterion (BIC) to score our models. The BIC weighs improvements in the log likelihood against a penalty for additional parameters. It is defined as $BIC \equiv -2LL + K\log(n)/n$. Where LL is the log likelihood, K is the number of parameters and n is the number of data. Model (a) would be selected over model (b) according to the BIC if it had a lower score. As an example, the rational expectations model has a BIC of infinity; because although it is a zero parameter model the likelihood that our data are generated by that model is zero. Table 4 reports the BIC scores for a few alternative models that cannot be otherwise rejected as possible explanations. The right hand column reports the BIC for models that do not condition on past actions (No Inertia), while all the models in the left hand column condition on c_{-1}. Our model of inattentive reconsideration (IR) at the bottom of the left-hand column has to lowest score among the models we look at.

We include models that do not condition on past actions to allow for the possibility that capturing this feature of the data leads to overfitting. We consider two econometric models—probit and logit—the rational expectations model, the rational expectations model with additive noise and a rational inattention model like the one proposed by Sims (2003). The probit and logit models include only x, x^2, and r as regressors as well as a constant. Since the deterministic threshold function $s(x)$ will obviously be rejected by the data we relax this strong assumption by allowing for additive noise in the implementation of the RE policy, this is the RE+noise model.

Table 4: Model Selection: Bayes Information Criterion

	INERTIA	NO INERTIA
Logit	19,819	20,551
Probit	20,064	20,799
RE		∞
RE+noise		23,516
RI		20,419
RI+switch cost	19,851	
IR	19,623	

Just like the RE model, however, only x and r enter into the policy function. In the RI model subjects choose the probability of buying subject to the mutual information cost function. The solution is a logistic distribution that once again only depends on the state. The RI model does best among this class of models. It is interesting how poorly the RE+noise model does. It is even beat by statistical models. This suggests that subjects are not just trembling in their implementation of the threshold $s(x)$. None of these models, however, have higher scores than any of the models that condition on past choices. In particular, both the logit and probit models' score improves by including c_{-1} as a regressor. Yet another piece of evidence pointing toward the importance of inertia in our data. We now turn to alternative explanations of this fact.

One could argue that habits result from cognitive biases that we would otherwise avoid if we could help it. If this were the case, then the fact that our experiment makes the objective time separable would not mitigate this type of behavior. Consider the case of external habits. Under external habits, agents derive utility from the quasi difference of their own consumption and *others'* consumption. It may be that agents would prefer not take other's consumption into consideration but simply cannot help feeling envy. Our subjects could not have conditioned on each others' outcomes. They did not observe each other's payoffs or history of play.

Habits could also be formed by subject's own past consumption. The fact that our subjects only got paid at the end makes this implausible. Yet it could be that the same biases that would lead subjects to become attached to goods make them somehow attached to the rental offers in our game. Suppose that in any given turn our subjects did not consider r_t but rather $r_t - \lambda R_t$ where the stock of habits, R_t is a linear combination of the elements in the history of rents accrued $\{c_{t-k} r_{t-k}\}_k$. Let $R_t = r_{t-1} c_{t-1}$. This is the simplest form of habits; the habit stock depends only on the previous period's consumption. These types of preferences would *not* lead to inertial behavior. The reason is that habits from preferences relies on *two* assumptions: non-linear utility and non-time-separable preferences. If the objective is linear, as in our experimental setup, previous consumption is a sunk cost. The disutility $r_{t-1} c_{t-1}$ is independent of the choice and therefore not relevant for today's decision. We can also test this in the data by including a the term $r_{t-1} c_{t-1}$ in our regression. This term is not significant.

Under a richer specification, R_t may depend on further lags of accrued rentals. In this case inertia would be generated because R_t is part of the state. In the previous case, only today's choice carries forward so past actions do not enter into continuation value of today's choice. If R_t has more lags then past actions will now be part of the continuation value of today's choice. R_t is still a sunk cost, but inertia is embedded in the continuation value of the bellman equation. The intuition is that the agent anticipates how the habit stock will affect future utility and takes that into account when making choices today. If the habit stock going forward contains past choices, then through R_t those will influence today's choice. For simplicity, let R_t be the total rents accrued up to time t. During the experiment the screen displays this total, making it easy for subjects to potentially condition their decisions on this statistic. We can test this by including R_t as a regressor in equation (8). The coefficient on cumulative rents is not significant, with a *p-value* of .8. The point estimate is $-2.2 \cdot 10^{-8}$ with a standard error of $8.8 \cdot 10^{-8}$. Meanwhile the estimates for the other coefficients remain unchanged.

One could alternatively argue that though there are no external switching costs, there is a behavioral attachment to past actions that is costly to break. This is similar to the habits model described above, except that the stock of habits is built over actions rather than over rental payments. In this case a subject would pay a fixed cost λ every time her actions in period t are different from her actions in period $t-1$. We combine these costs in with rational inattention—*RI+switch cost* in Table 4—and estimate the model via maximum likelihood. The key difference between this and our model of inattentive reconsideration is that inertia here is not state-dependent. In this model the costs of switching create a wedge in the odds ratio equal to $\exp\{\theta\lambda\}$, where θ is the marginal cost of attention. It has a higher score, and thus a worse fit, than our model of state-dependent inertia even though it has one fewer parameter. The key is state-dependence; the model with fixed switching costs will either over-predict inertial behavior in states were it should not matter, as shown in Figure 12, or under-predict it everywhere else.

7 Conclusion

We have run an experiment to replicate the consumption problem studied in macroeconomics. We have concluded that the failure of people to implement the rational expectations policy must come from cognitive limitations. We have identified two such types of limitations: costs of reconsideration and, conditional on reconsidering, costs of deliberation. The first leads to inertia while the second to—partial—insensitivity to the state. Although our data are generated in highly stylized environment, they nonetheless display these two features which have been documented not only on aggregate consumption data, but increasingly at the household level as well.

Our estimates show that even small cognitive costs can have stark behavioral implications. While state-dependent habitual behavior may have cost players less than one percent of the winnings they would have had without it, inertia of this type nonetheless has crucial implications for the reactions of aggregate demand to shocks. State-dependent inertia implies that aggregate consumption will be more forward-looking in times of crisis than during normal times.

At the micro level, our model also makes predictions about individual household consumption choices. It suggests that these choices may be akin to how firms set prices. In our model, inertia in aggregate consumption comes from the extensive margin. Individual consumption in our model need not react to shocks even if aggregate consumption does. This is, of course, an empirical question and an obvious direction for future research.

Finally, we want to highlight how our findings contribute to a growing literature that models short-run macroeconomic fluctuations as arising from some type of "bounded rationality." One of the truisms associated with this literature is that while there is only one way to behave rationally there are an infinite number of ways to deviate from the rational expectations benchmark. Our paper takes this concern seriously and makes use of experimental data to discipline the way we model bounded rationality. Going forward, experiments may prove a useful tool for macroeconomists as a way to discriminate among behavioral models. The fact that we adapt a model originally developed to explain *pricing* behavior is also significant. Calibration exercises of general equilibrium models usually involve adding many different ingredients ad hoc to match the inertia we observe in the data. Our paper is yet another piece of evidence that perhaps we can move toward a more parsimonious model where the same cognitive limitations—present in all agents—can be responsible for price sluggishness and consumption inertia.

References

Abel, A. (1990). Asset Pricing Under Habit Formation and Catching Up With the Joneses. *American Economic Review, PP*. 80(2), 38-42.

Angeletos G-M., Z. Huo (2018). Myopia and Anchoring. *NBER Working Paper* No. 24545.

Camerer, C., P. Landry, R. Webb (2018). The Neuroeconomics of Habit. *Unpublished*.

Caplin. A., M. Dean (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*. 105(7), 2183-2203.

Chamberlain, G. (1980). Analysis of Covariance with Qualitative Data. *Review of Economic Studies*. 47, 225-238.

Deaton, A. (1992). *Understanding Consumption*. New York: Oxford University Press.

Handel, B., J. Schwartzstein (2018). Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care? *Journal of Economic Perspectives.* 23(1), 155-178.

Hebert, B., M. Woodford (2017). Rational Inattention with Sequential Information Sampling. *NBER Working Paper* No. 23787.

Hortaçsu, A., S.A. Madanizadeh, S.L. Puller (2017). Power to Choose? An Analysis of Consumer Inertia in the Electricity Market. *American Economic Journal: Economic Policy*. 9(4), 192-226.

Kahneman, D. (2011). *Thinking, Fast and Slow.* London: Penguin Books.

Khaw, M.W., L. Stevens, M. Woodford (2017). Discrete Adjustment to a Changing Environment: Experimental Evidence. *Journal of Monetary Economics.* 91, 88-103.

Klemperer, P. (1995). Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade. *Reveiw of Economic Studies*. 62, 515-539.

Laibson, D. (2001). A cue-theory of Consumption. *Quarterly Journal of Economics*. 116(1), 81-119.

Lorenzoni, G. (2009). A Theory of Demand Shocks. *American Economic Review*. 99(5), 2050-2084.

Mankiw, G., R. Reis (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*. 117(4), 1295-1328.

Martin, V., S. Hurn, D. Harris (2013). *Econometric Modeling with Time Series: Specification, Estimation and Testing*. New York: Cambridge University Press.

Matejka, F., A. McKay (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review*. 105(1), 272-298.

Matyskova, L., B. Rogers, J. Steiner, K-K. Sun (2018). Habits as Adaptations: An Experimental Study. *Unpublished*.

Nakamura E., J. Steinsson (2011). Price Setting in Forward-Looking Customer Markets. *Journal of Monetary Economics*. 58(3), 220-233.

Nakamura, E., J. Steinsson (2018). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*. 133(3), 1283-1330.

Ravn, M., S. Schmitt-Grohe, M. Uribe (2006). Deep Habits. *Review of Economic Studies*. 73, 195-218.

Redish, A.D., (2013). *The Mind Within the Brain*. New York: Oxford University Press.

Reis, R. (2006). Inattentive Consumers. *Journal of Monetary Economics*. 53(8), 1761-1800.

Sanfrey, A.G., L. Chang (2008). Multiple Systems in Decision Making. *Annals of the New York Academy of Sciences*. 1128(1), 53-62.

Schmitt-Grohe, S., M. Uribe (2012). What's News in Business Cycles. *Econometrica*. 80(6), 2733-2764.

Sims, C.A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*. 50(3), 665-690.

Steiner. J., C. Stewart, F. Matejka (2017). Rational Inattention Dynamics: Inertia and Delay in Decision-Making. *Econometrica*. 85(2), 521-553.

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

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

Wooldridge, J. (2010). *Econometric Analysis of Cross Section and Panel Data, second edition*. Cambridge, MA: MIT Press.