Rationally Inattentive Consumer: An Experiment

Andrea Civelli*
University of Arkansas

Cary Deck
University of Alabama
Chapman University

Justin D. LeBlanc
University of Arkansas
RichContext

Antonella Tutino†
Federal Reserve Bank of Dallas

Abstract

This paper presents a laboratory experiment that directly tests the theoretical predictions of consumption choices under rational inattention. Subjects are asked to select consumption when income is random. They can optimally decide to reduce uncertainty about income by acquiring signals about it. The informativeness of the signals directly relates to the cognitive effort required to process the information. We find that subjects' behavior is largely in line with the predictions of the theory: 1) Subjects optimally make stochastic consumption choices; 2) They respond to incentives and changes in the economic environment by varying their attention and consumption; 3) They respond asymmetrically to positive and negative shocks to income, with negative shocks triggering stronger and faster reactions than positive shocks.

Keywords: Rational Inattention, Experimental Evidence, Information Processing Capacity, Consumption
JEL Classification: C91, D11, D8, E20.

*University of Arkansas, Walton College of Business, Department of Economics, Business Building 402, Fayetteville, AR 72701. Email: andrea.civelli@gmail.com.

†Federal Reserve Bank of Dallas. Email: tutino.antonella@gmail.com

1 Introduction

Cognitive limits in processing information have important implications in a variety of economic settings. Theoretical studies on stochastic choices (Matějka and McKay, 2015), investment decisions (see, among others, Mondria, 2010) and pricing decisions (Mackowiak, Bartosz, and Wiederhold, 2009), as well as empirical studies on consumers' choices,[1] financial markets and voting behavior (examples include DellaVigna and Pollet, 2007; Shue and Luttmer, 2009) highlight the role people's limitations to paying attention to changes in the economic environment, even if these changes are relevant for their decisions. While there are several approaches to information acquisition (see Hellwig, Kohls, and Veldkamp, 2012, for a survey), this paper is concerned with the theory of rational inattention pioneered by Sims (1998) and Sims (2003). The goal of the paper is to formally test the principles and implication of this theory for individual consumption decisions by using a laboratory experiment.

Sims postulates a model in which a decision-maker chooses optimally the amount of information for a given decision problem where the cost is based on Shannon's mutual information between prior and posterior beliefs. Following this approach, Tutino (2013) builds a consumption-saving decision problem where consumers, aware of their cognitive limitations, change the amount of information and attention in response to changes in the economic environment. Tutino (2013) explicitly considers the cost implied by Shannon's mutual information as a *cognitive* cost, i.e. a cost that captures the limits of people to map quickly and precisely all the available information about the economic environment into consumption choices. Moreover, the paper shows that for a given shadow cost of processing information, consumers react to exogenous economic changes by varying the informativeness of the signals they require to select consumption choices.

This concept, known as *elastic* information processing capacity, implies that different economic stimuli correspond to different consumption behaviors, according to whether the decision-maker deems it beneficial to save processing effort by accepting lower consumption or she prefers to incur the cost of paying additional attention to increase the informativeness of signals and make sharper consumption choices. Tutino (2013) shows that taking into account cognitive limitations accounts for a number of empirical regularities in the consumption literature and, more importantly, documents novel findings pertaining to rational inattention theory on randomness of choices and asymmetric responses to economic changes that are

[1]See, for instance, Chetty, Looney, and Kroft (2009) for an application to sales taxes and de los Santos, Hortasu, and Wildenbeest (2012) for an application to e-commerce.

hard to reconcile with standard choice theories.

This paper formally tests the predictions of a static version of the model in Tutino (2013) in a controlled laboratory experiment. In the basic task, a simple decision problem is faced by the participants, who make consumption decisions under income uncertainty. In each period, income is randomly drawn from a uniform distribution. Prior to selecting a consumption level, subjects can reduce the uncertainty about their income by acquiring signals that inform them of how much income they may have available for the period.

The precision of the signal is directly proportional to the cognitive effort participants need to exercise in order to extract information. We capture cognitive effort by requiring participants to solve logic puzzles, with more difficult puzzles demanding more effort to be solved by the subject. Since rational inattention is at its core a theory of how hard people think, the logical puzzles allow us to measure the effort participants choose. We relate the difficulty of the puzzles to the informativeness of the signals, with harder puzzles corresponding to more precise information about income (Civelli and Deck, 2018). Difficulties of the puzzles range from trivial (uninformative) to hardest (perfectly revealing). Upon completion of a task, participants update their prior beliefs about income and select their desired consumption level based on their informed posterior.

If the selected consumption level does not exceed the income drawn that period, the subject obtains the selected consumption level, otherwise the subject moves on to the next period with zero consumption. Participants are paid on the basis of their accumulated consumption over the course of the experiment. Other than the opportunity cost of processing information versus consumption, there is no time limit on any given period.[2] By repeatedly exposing the subjects to each decision problem we can collect data on consumption choices as well as information/signal attempted and realized choices. These data are relevant in assessing whether the predictions of the rational inattention model are verified (see Caplin and Dean, 2015).

We focus on five main specific predictions from Tutino (2013) and design experimental treatments conducive to directly verifying whether laboratory evidence corroborates these predictions. The first two predictions we test concern the value of information and stochastic choices of consumption decisions. Consistent with the theory, the behavioral evidence reveals that, while on average a more informed subject receives a higher payoff, subjects often prefer to exercise low effort and process less information even if that choice implies lower

[2]The overall length of the task was random with an expected duration of 20 minutes. Subjects who neither acquired information nor consumed successfully experienced a 60 second reduction in the alloted task time as a penalty.

consumption rewards. Moreover, unlike standard choice theories, we verify the prediction of randomness in choices by documenting stochastic behavior, with consumption outcomes generating a non-spurious posterior distribution.[3]

The third prediction we verify concerns whether subjects respond to monetary incentives by varying their processing effort. To test this prediction, we modify our baseline experiment by varying the payoffs offered at specific intervals. In particular, we inform the subject that in this treatment the payoff will substantially increase every eleven periods while in other periods consumption values are a fraction of those in the baseline. We find that participants take advantage of the higher payoff by processing more information in periods where the consumption possibility are more sizable and reduce their cognitive effort with respect to the baseline treatment when consumption offers are small. Thus, unlike random utility model where attention is fixed, we document that subjects respond to incentives by varying information processing efforts when presented with lucrative alternatives. This finding is consistent with the notion of elastic capacity where individuals modulates their attention and cognitive effort according to what is at stake.

The fourth prediction of the theory posits that individuals tend to process less information and make less precise consumption decisions when the economic environment becomes more predictable. We test this prediction by implementing a treatment where the income process has persistence. In such an environment, we show that subjects change their consumption choice and signal precision infrequently. Generally, subjects fail to realize changes in their income possibilities as they process less information about income than they do in the baseline. In particular, we document that participants resolve the trade-off between cost of processing more precise signals with the benefit of higher payoffs by forgoing units of consumption rather than acquiring better information.

The fifth theoretical prediction we put to the empirical test is the asymmetric response of consumers to shocks, with negative shocks triggering faster and more sizable consumption reaction than positive shock. This prediction is a novel finding in Tutino (2013) which makes rational inattention theory observationally distinct from other theories of costly information acquisitions and attentiveness. We verify this prediction by feeding shocks of different size and sign to the predictable environment described for the third prediction. We found corroborating evidence to the finding of asymmetric consumption responses in the direction predicted by the theory. This finding is also supported by a host of empirical results on

[3]For other examples of this result in experimental data connected to rational inattention, see Dean and Neligh (2017), Khaw, Stevens, and Woodford (2016), Cheremukhin, Popova, and Tutino (2015).

consumption demand (see, for instance, Abaluck and Adams, 2017; Shea, 2009), and the impact of taxation (Broda and Parker, 2014; Johnson, Parker, and Souleles, 2006).

This paper contributes to the literature concerning stochastic choices under uncertainty and to the recent and growing literature aimed at measuring the role of information acquisition in rationalizing economic outcomes. Recent studies test different models of attention and information acquisition (Khaw, Stevens, and Woodford, 2016; Gabaix, Laibson, Moloche, and Weinberg, 2006; Manzini and Mariotti, 2015). Within this literature, a subset of papers (see PInkofskiy, 2009; Cheremukhin, Popova, and Tutino, 2015), experimentally test models of rational inattention using binary choices between gambles. Unlike these papers, our model is designed to directly quantify participants' choices of cognitive effort by employing logic puzzles and to relate the effort to consumption choices. A closely related paper that experimentally checks the predictions of rational inattention against experimental data with perceptual tests is Dean and Neligh (2017). We differ from their approach in that we design the experimental setting to test a particular set of predictions of the rational inattention theory with respect to consumption behavior. This paper corroborates the empirical tests of Dean and Neligh (2017) regarding incentives and randomness in rational inattention models. Moreover, we complement their results with the empirical findings of delayed and asymmetric responses of consumption and attention to stimuli. To our knowledge, this is the first paper that explicitly test for asymmetry in response to shocks of rationally inattentive agents.

The rest of the paper is organized as follows. Section 2 formally presents the theoretical consumption decision problem under rational inattention. It discusses the properties of the problem's solution and the testable predictions. Section 3 lays out the experimental setting and the treatments implemented to verify the theoretical prediction. A mapping between experimental set-up and the theoretical model is formally established. Section 4 shows the congruence of the experimental with the theoretical predictions. Finally, Section 5 offers some concluding remarks. Robustness checks are relegated to an Appendix.

2 Theoretical Framework

We briefly introduce the theoretical framework to provide a structure for understanding the consumption and attention allocation decisions faced by subjects in the experiment. Our experimental design implements a static version of the theoretical model by Tutino (2013). The model derives testable predictions of rationally inattentive agents' behavioral responses to changes in the economic environment during their lifetime. We use this model as a guide to discipline our experiment and to contrast the theoretical predictions with the laboratory evidence.

The model describes an optimization problem in which an information processing capacity constraint is added to an otherwise standard consumption-saving problem. We focus here on the key aspects and implications of these limits to information processing for the consumer's optimal decision. We formally outline the model with more technical details in Appendix A.

An agent with limited information capacity faces a consumption decision under uncertainty about her wealth. Since wealth is unknown, the agent cannot know precisely how much consumption she can afford. Hence, she must treat both wealth and consumption as random variables before deciding how much information about wealth she wants to process in order to consume. As a result, the optimization problem must be expressed and solved in terms of joint probability distributions over wealth and consumption. By contrast, in the infinite information capacity case, wealth is perfectly known and utility is directly maximized by choosing the optimal consumption stream.

More specifically, the optimization problem entails three main parts. First, the agent starts each period t with a prior distribution on her initial level of wealth; although wealth is unknown, this prior distribution is not. Let w_t indicate wealth, the prior is given by the distribution function $g(w_t)$. Before processing any information, consumption is also a random variable. This is because the uncertainty about wealth translates into a number of possible consumption profiles with various levels of affordability.

Second, the agent chooses how much information about wealth she wants to process in order to make informed consumption choices. One way of thinking about the information processing decision is that the consumer chooses a noisy signal on wealth, where the noise can take on any distribution selected by the consumer consistent with the information processing capacity. In other words, the consumer allocates attention to forming a new probability distribution for w functional to the consumption decision which improves on the prior $g(w)$.

In more precise terms, the reduction of uncertainty about wealth and the consumption choice, indicated by c_t, must be seen as two sides of the same utility maximization problem which occur simultaneously. Hence, when information cannot flow at infinite rate, the choice of the consumer is actually the joint distribution $p(w_t, c_t)$, as opposed to the stream of consumption $\{c_t\}_{t=0}^{\infty}$ we would have in the full information case. Given that the agent has a probability distribution over wealth, choosing the joint distribution $p(c_t, w_t)$ is akin to choosing a signal on wealth. This is easily illustrated by applying Bayes' rule to the joint distribution

$$p(w_t|c_t) = \frac{p(c_t, w_t)}{\int p(c_t, w_t) dw_t}$$

where the type of signal chosen corresponds to the conditional probability $p(w_t|c_t)$.[4]

The optimal choice of the joint distribution, $p^*(c_t, w_t)$, depends on the constraint on the amount of information the consumer can processes. $p^*(c_t, w_t)$ makes the distribution of consumption conditional on wealth as close to wealth as the limits imposed by Shannon capacity allow. Before processing any information, uncertainty about wealth is measured by the entropy of the prior of wealth $\mathcal{H}(w_t) \equiv -\mathbb{E}[\log_2(g(w_t))]$. Since consumption and wealth are related in the consumer's decision, knowledge of consumption provides information about wealth. The reduction in uncertainty about wealth that can be extracted by knowing consumption is expressed as the conditional entropy (or residual uncertainty) about wealth once consumption is known, $\mathcal{H}(w_t|c_t)$. The information flow, $I(c_t, w_t)$, associated with a signal is the maximum reduction of entropy induced by the signal, which is bounded by the information that a selected signal conveys, κ,

$$I(c_t, w_t) = \mathcal{H}(w_t) - \mathcal{H}(w_t|c_t) \leq \kappa.$$

The bound on the information flow depends on the effort the consumer chooses to exercise to track her wealth. Paying attention to reduce uncertainty requires spending some time and utility to process information. The task of thinking is modeled by augmenting the optimization problem with a Shannon channel. Limits in the capacity of the consumers to process information are captured by the fact that the reduction in uncertainty conveyed by the signal cannot be higher than a given number κ_t, which explicitly limits the set of available signals by introducing an upper bound to the feasible information flows.

The third part of the maximization problem regards the transition law from one period to the next, which in this setup corresponds to how the consumer's information on wealth evolves over time rather than the standard law of accumulation of wealth. At the end of each period, after the realization of c_t, the consumer is endowed with a new income draw, y_{t+1}, from a known distribution which corresponds to the prior of the initial wealth $g(w_0)$, since $w_0 = y_0$ by construction. The law of accumulation of wealth reads

$$w_{t+1} = R(w_t - c_t) + y_{t+1}, \tag{1}$$

[4]The conditional distribution $p(w_t|c_t)$ represents the signal of wealth the consumer receives after consumption is realized, where consumption is a draw from the optimally chosen $p(c_t, w_t)$.

where R is the interest on savings. The stochasticity of w derives from the stochasticity of y.

The way beliefs about wealth transit across states is described by the update law of the prior $g(w_t)$ into its posterior $g(w_{t+1})$. This law takes into account how the choice in time t, $p(w_t, c_t)$, affects the distribution of the consumer's belief after observing c_t as well as the stochastic accumulation of wealth in (1). Given the initial prior state $g(w_t)$, the next period belief state $g'(w_{t+1}|c_t)$ is determined by revising each state probability as displayed by the expression known as Bayesian conditioning

$$g'(w_{t+1}|c_t) = \int \tilde{T}(w_{t+1}; w_t, c_t) p(w_t|c_t) dw_t,$$

where $\tilde{T}(\cdot)$ is the transition function embedding (1). $g'(w_{t+1}|c_t)$ is used, in turns, as the new prior for the period $t+1$, $g(w_{t+1})$.

Let $u(c)$ be the utility of the household defined over consumption c. Combining these three components, the program of the consumer under information frictions can be written as:

$$\max_{p(w_t, c_t)_{t=0}^{\infty}} \quad \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \int u(c_t) p(c_t, w_t) \mu(dc_t, dw_t) | \mathcal{I}_0 \right\} \tag{2}$$

$$\text{s.t.} \quad \kappa_t = I_t(p(c_t, w_t)) \tag{3}$$

$$g'(w_{t+1}|c_t) = \int \tilde{T}(w_{t+1}; w_t, c_t) p(w_t|c_t) dw_t \tag{4}$$

$$p(c_t, w_t) \in \mathcal{D}(w, c) \tag{5}$$

$$g(w_0) \text{ given} \tag{6}$$

where $\mu(\cdot)$ is the Dirac measure that accounts for discreteness in the optimal choice of $p(c, w)$, while $\mathcal{D}(w, c)$ restricts the choice of the agent to be drawn from the set of distributions, $\mathbb{E}_0$ is the conditional expectation defined with respect to the σ-algebra $\mathcal{I}_0$.

Let θ indicate the *Lagrange multiplier* associated with constraint (3), the total cost of an agent of choosing a signal and processing information is then given by $\theta\kappa_t$. We allow κ_t to vary over time to capture the possibility that subjects in the experiment may want to process different amount of information within and between treatments. This formulation is known theoretically as subject's *elastic capacity* of information processing. Subjects' own θ is not directly observable in the experiment. However, we impose a structure of signals with different precision levels representing different information flows (κ), and let the subjects select among the alternatives according to how much time and effort they want to invest to sharpen their knowledge of wealth. In turn, this choice determines their optimal total cost.

The program in equations (2)-(6) is a well-posed mathematical problem, but with state and control variables that are infinite dimensional. However, Tutino (2013) shows that a discretization of this framework is also a well-posed problem and returns as a solution the ergodic distribution $p^*(c_t, w_t)$ which, from Bayes' rule as well, can be represented as

$$p^*(c_t, w_t) = p^*(c_t|w_t)g(w_t),$$

where the marginal wealth distribution is equal to the prior, $\int p^*(c_t, w_t)dc_t = g(w_t)$, to satisfy model's internal consistency. The conditional distribution $p^*(c_t|w_t)$ embeds the effects of more accurate information about wealth provided by the selected signal to sharpen consumption choices.

Our experiment studies the implications of the equilibrium solution of the static version of this model. The model is simplified in a static setting since savings are not possible, and the unused wealth is forgone. In fact, in a static setting $w_t = y_t \ \forall t$, and the evolution of beliefs about wealth (4) is simply replaced by the prior itself $g(w_t) = g(w_0) \ \forall t$; similarly, the objective function turns into the one-period version of (2).

In this case, the optimal solution is readily computed and can be expressed as:

$$p^*\left(c, w\right) = g\left(w\right)\left(e^{\left(\theta + \frac{u(c)}{\theta}\ln 2\right)} - 1\right), \tag{7}$$

which illustrates the solution depends on the shadow cost of processing information, θ, the prior distribution of wealth, $g(w)$, and the functional form of utility, $u(c)$.

In Section 4, we discuss the main predictions obtained from the optimal solution of the static setting which can be tested with our experimental data.

3 Experimental Design and Implementation

This section discusses our experimental design. We first introduce the building blocks of the basic experiment and explicitly link them to the theoretical model in Section 2. We then describe the cognitive tasks the subjects engage with during the experiment and treatment variations we study.

3.1 Experimental design

We devise an experiment that aims to closely replicate in a laboratory setting the static version of Tutino (2013)'s model. Throughout the experiment, we present the subjects with a choice of 256 prizes whose values are expressed in experimental dollars. Prizes are sorted by increasing cost ranging from 1 to 256. The value of the prizes is directly related to its cost, i.e. how many experimental dollars it takes for the subject to claim a particular prize. The subjects attempt to maximize the total sum of prizes accumulated in the experiment which represents their final payoff. Subjects are endowed with a income in each period of the experiment. Each period, income is drawn from a distribution known to the subjects; however, the realization of the draw is unknown. Thus, the decision problem of the participants is to select prizes under income uncertainty. Thus, their decision problem boils down to selecting prizes under uncertainty about their available budget.

Subjects can reduce uncertainty about how big of a prize they can afford each period by choosing a signal of their endowment before selecting a prize. We provide them with an array of signals with varying levels of precision. The signals provide more information by narrowing the range from which the endowment was drawn. The precision of the signals is tightly linked to the cognitive effort the subjects need to exert in order to reveal the signals' information. The higher the precision of the signal is, the narrower the interval of possible endowments it provides, and the more significant the cognitive effort a subject must exert to achieve the signal. The precision levels available to the subjects and the associated average difficulty levels are illustrated in Table 1. More details about signals, intervals, and cognitive tasks are discussed in Section 3.2. With the introduction of these signals, the decision problem of the subjects becomes how much information they want to acquire before selecting the prizes.

We can explicitly map these features of experimental designs to the building blocks of the theoretical framework in Section 2. The cost of the prizes is the equivalent of the amount of consumption c chosen by the optimizing agent in the model. The value of the prizes represents the experimental counterpart of the utility derived by the agents from consumption. Thus, choosing prizes is akin to choosing consumption. The endowment in the experiment is income y, which in the static setup is equal to wealth in any point in time, $y = w$. The prior on wealth is the known distribution over the income/consumption support from which endowments are drawn, $g(y)$. For the sake of concreteness, in the baseline treatment of the experiment $g(y)$ takes on the form of a discrete uniform distribution on the 1-256 interval.

To simplify the experimental design, the structure of the signals we provide belongs to the uniform family. In terms of the model, this simplification maps into the constraint (5) requiring that the distribution is chosen not from any family, but from the discrete distribution. We acknowledge the fact that a uniform distribution is unlikely to be the optimal for the problem in (2)-(6), since it is the distribution with the highest entropy. However, we decide to forgo optimality for tractability of the experimental set-up and, more importantly, to increase clarity for the subjects about the experiment thereby avoiding mistakes due to misunderstandings of the experimental structure. Under this additional constraint, subjects can choose any feasible signal precision consistent with their information processing limits. By allowing the subjects to choose the precision level of the signals, we can infer subjects' information processing shadow cost, θ. As mentioned earlier, θ is not observed; however, subjects target their optimal total cost of information processing, $\theta\kappa_t$, by choosing the cognitive effort associated with a signal.

The interval revealed to the subjects by an acquired signal defines the corresponding optimal conditional distribution $p^*(c|w)$ in the theoretical framework. By narrowing the support of the income draw, subjects can modify their prior $g(y)$ with a more concentrated distribution reflecting lower uncertainty about income. For instance, in the baseline treatment, subjects would go from an initial discrete uniform prior on the full support 1-256 to a discrete uniform distribution over a restricted support, corresponding to a smaller interval. The observed realization of a particular prize is a draw from the optimal distribution $p^*(c|w)$. Finally, since in the experiment periods are not statistically linked in the baseline given the i.i.d. nature of income y, the prior $g(y)$ is identical in each period, simplifying the constraint (4) to $g'(w_{t+1}) = g(w_t) = g(y)$.

The experiment is deployed in the laboratory using the interface shown in Figure 1. On the top portion of the screen, the subjects see the amount of time elapsed thus far in the treatment (in seconds) in the left corner, the period currently being played in the middle, and the accumulated total prizes on the right. The central part of the screen conveys information about the possible values of the prizes, while the bottom row of aqua blue buttons displays the nine levels of precision form the available signals from which the subject can choose. The buttons describe the size of the possible income intervals corresponding to each precision level, along with the difficulty level of the cognitive task expressed in terms of expected rate of success. For a signal to reveal the information about income, subjects must successfully complete the cognitive tasks associated to the signal selected. Both the expected success rates and the cognitive tasks will be described in details in Section 3.2.

We also use the screenshot in Figure 1 to visualize the mapping between theory and

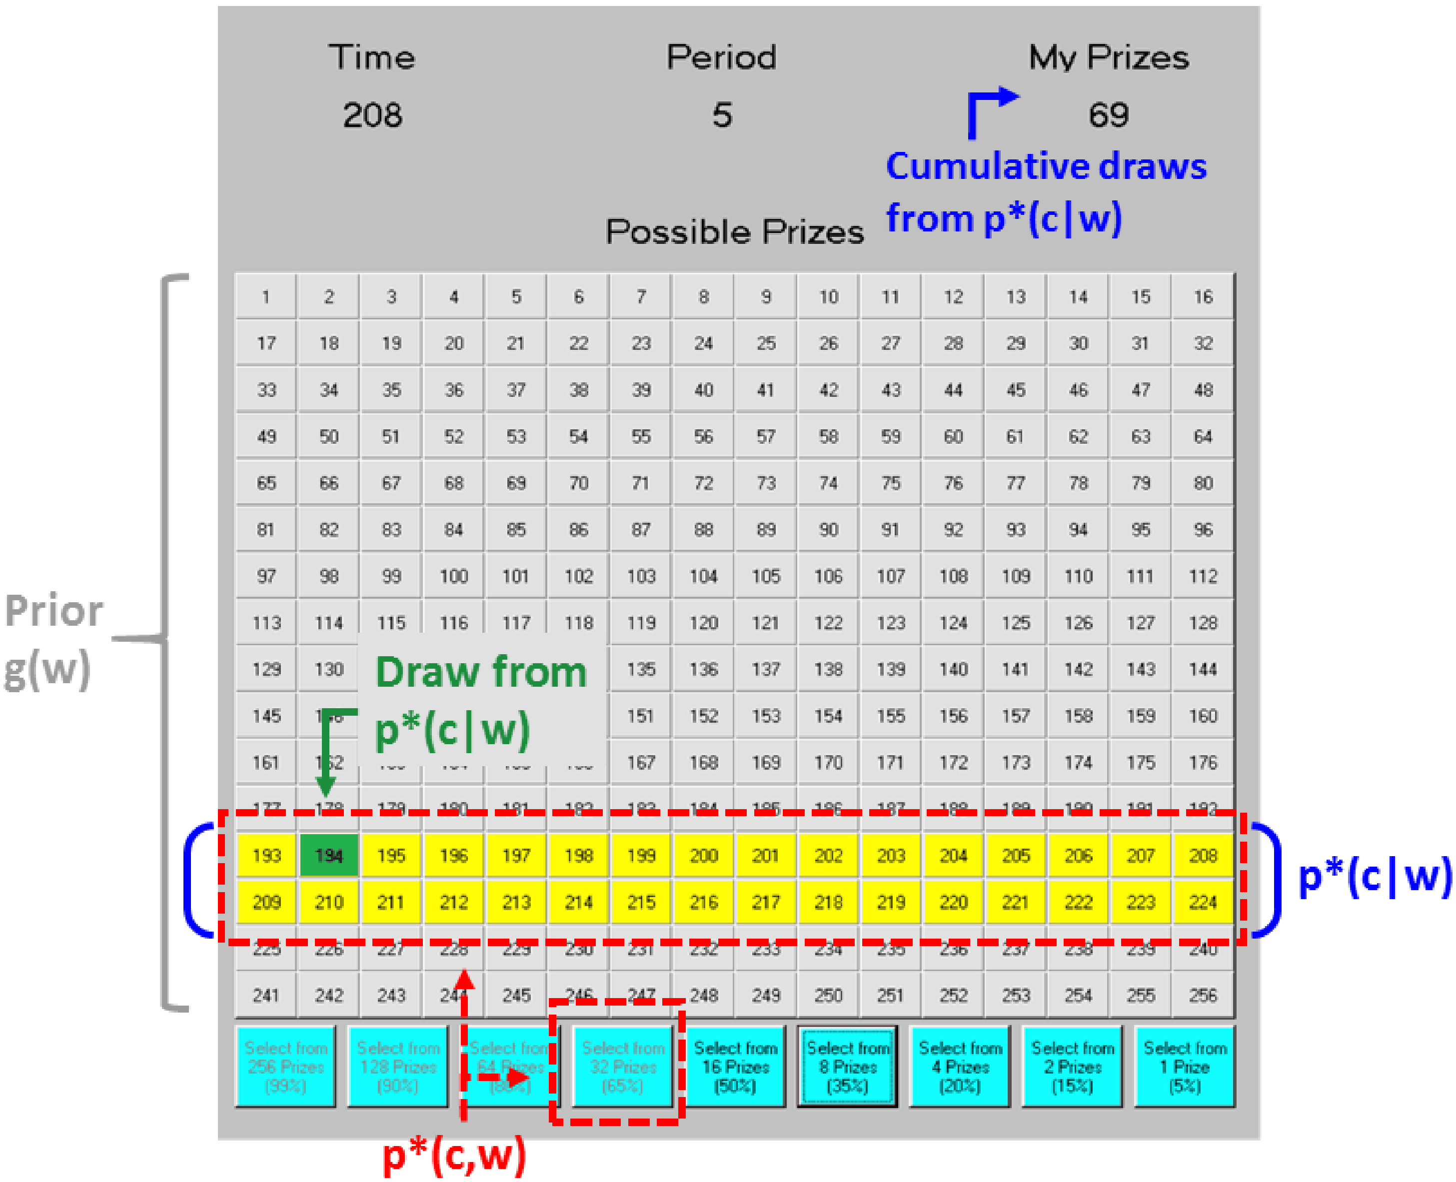

Figure 1: Mapping from the theory to the experiment: a screenshot of the experimental interface over which we superimpose the theoretical framework.

experiment. Subjects start each period with the grid of all potential prize values that reflects the diffuse prior on income $g(w)$. In this example illustrating the baseline, the value of prizes goes from 1 to 256 experimental dollars and there is a one to one transformation rate from the cost of prizes to their value. If subjects wants to proceed with no further information, they would select the first signal (with precision 0 in our classification) and the full set of prizes would turn yellow, indicating that the space of possible prize outcomes spans the whole support [1 256].

To make a concrete example, suppose a subject chooses a precision 3 signal and successfully completes the associated cognitive task. The signal reveals that income lies in the interval of 32 prizes from 193 to 224 and the two rows corresponding to this interval are in yellow. Thus, this particular signal selection allows the subject a reduction of income uncertainty of 1.5 bits – from the uninformative prior's entropy of 2.4 bits to the posterior's entropy of 0.9, as Table 1 illustrates.

The mapping between theory and experimental interface can be readily seen by interpreting the combined structure of prize grid and signal buttons as $p(c, w)$ *before* any decision is made by the subject. The optimal policy function of the subject is a strategy that select a signal whose information is functional to acquiring the prizes. This strategy is defined as $p^*(c, w)$, and in the experimental setting it encompasses both the choice of a particular signal on w and the information content about consumption possibilities revealed by the signal (as shown by the dashed red rectangle in Figure 1). Upon choosing the signal and successfully completing the cognitive task, the optimal posterior $p^*(c|w)$ maps into the highlighted interval of prizes in the grid, corresponding to the narrower support for the prizes, or, equivalently, to the reduction of uncertainty about income provided by the signal. The chosen prize, 194, can be treated as a random draw from the optimal posterior $p^*(c|w)$. If the prize drawn is smaller than the available income in that particular period, the prize box turns green and the prize value is added to the cumulative prize total.[5]

3.2 Signal characteristics and cognitive tasks

The information processing limits embedded in constraint (3) in the theoretical framework are implemented in the experiment through the structure of signals of different precisions that the subjects use to reduce uncertainty about income draws. We describe in this section

[5]Subjects are free to choose any prize even after receiving a signal. Knowing the more precise interval, though, subjects are expected, but not constrained, to choose from the prizes within the interval, as in the example discussed here. Section 3.3 describes how the experiment handles the cases in which the chosen prize is bigger than the income draw.

Precision Level	Intervals Description		% Task Correct	Information Flow (bits)
0	1 interval of length 256:	[1 256]	99%	0
1	2 intervals of length 128:	[1 128][129 256]	90%	0.3
2	4 intervals of length 64:	[1 64] ... [193 256]	80%	0.6
3	8 intervals of length 32:	[1 32] ... [225 256]	65%	0.9
4	16 intervals of length 16:	[1 16] ... [241 256]	50%	1.2
5	32 intervals of length 8:	[1 8] ... [249 256]	35%	1.5
6	64 intervals of length 4:	[1 4] ... [253 256]	20%	1.8
7	128 intervals of length 2:	[1 2] ... [255 256]	15%	2.1
8	256 intervals of length 1:	[1] ... [256]	5%	2.4

Table 1: Signal structure: precision levels, signal characteristics, expected success rate of the cognitive tasks, and information flows of the signals.

the characteristics and information content of these signals.

As summarized in Table 1, subjects have access to 9 levels of signal precision. The levels start from a basic Level 0 signal, which is uninformative and simply replicates the prior of the income distribution, to a Level 8 signal, which is fully revealing about the income draw. Letting j indicate the precision level, a signal identifies one of the 2^j intervals containing 2^{8-j} prizes into which the full support of prizes is partitioned.

We can compute the information flow of the signals by finding their variation in entropy starting from the uniform prior $g(y)$. The information flows of the signals are reported in the last column of Table 1. Since the signal structure is such that signals pertain to uniform distributions and the supports of the signals are proportional, our signals have the property that the change in entropy is constant from one precision level to the next, and equal to 0.3 bits.

Subjects must successfully complete a cognitive task in order to acquire a signal. The precision of the signals is proportional to the cognitive difficulty of the task, and the subjects must exert a higher cognitive effort to obtain a better signal. We rely on the visual puzzles developed by Civelli and Deck (2018) to compose these tasks. These puzzles take on the form of a (3×3) graphical matrix in which eight images are provided and one is missing. Subjects must identify the missing image among a set of alternatives, after analyzing the patterns of attributes shared by the eight images in the matrix. The puzzles are generated in the spirit of the Raven's Progressive Matrices, and their level of difficulty is correlated with the reasoning abilities of an individual. An example puzzle is illustrated in Figure 2, while the corresponding solution set with eight images is shown in Figure 3.

Each image in the puzzle has six attributes that could change: shape, size, shade of the filling, orientation, border style, and corner marker style. These attributes are allowed to change following six schemes of patterns: orthogonal - along rows and columns; diagonal - along main or counter-diagonal; and corners - from NW to SE and from SW to NE. The difficulty level of a matrix is determined by the number of attributes allowed to change. Given the puzzle calibration exercises conducted by Civelli and Deck (2018), we are able to create tasks of any desired cognitive difficulty by requiring someone to solve a series of puzzles of various difficulty levels. A task is usually composed of 2 puzzles, with the exception of the highest precision level that requires solving 3 very difficult puzzles. The reference expected difficulty level of the task of each signal is reported in Table 1 as well.

In the example of Figure 2, two attributes change: shape and size, while shade, orientation, border style, and corner marker style remain unchanged. Notice that along the diagonal from upper left to lower right, the size is the same. Further, notice that the shapes vary along the counter-diagonal. Given these observations about the patterns of attributes in the puzzle matrix, one can deduce that option "A" in Figure 3 is the correct solution of the puzzle since it has the correct characteristics to complete the matrix.

3.3 Experimental Implementation

The experiment was written in Visual Basic, and it was deployed as a stand alone executable on the machines in the Behavioral Business Research Laboratory at the University of Arkansas. The experiment is a within-subject design, and it is broken into an un-paid practice round and 3 treatment rounds including the baseline. The practice round lasted 10 minutes and was otherwise identical to the baseline. Treatments are randomly ordered so as to mitigate order effects. Before beginning the practice round, the subjects worked through a series of instructions on their screen, and each subject was offered an opportunity to ask clarifying questions before beginning the actual experiment. These instructions have been included in Appendix B.

Subjects were recruited from the undergraduate student body at the University of Arkansas, and a total of 64 subjects participated in the experiment over the course of 6 sessions in the month of May 2017. Subjects earned an average total profit, including a $5 participation payment, of $25.30 USD. Each session lasted approximately 90 minutes, which included working through the instructions, the practice round, and participating in the actual experiment. Payouts in excess of the show-up fee were determined by the subjects total cumulative prizes across all three treatments divided by 500.

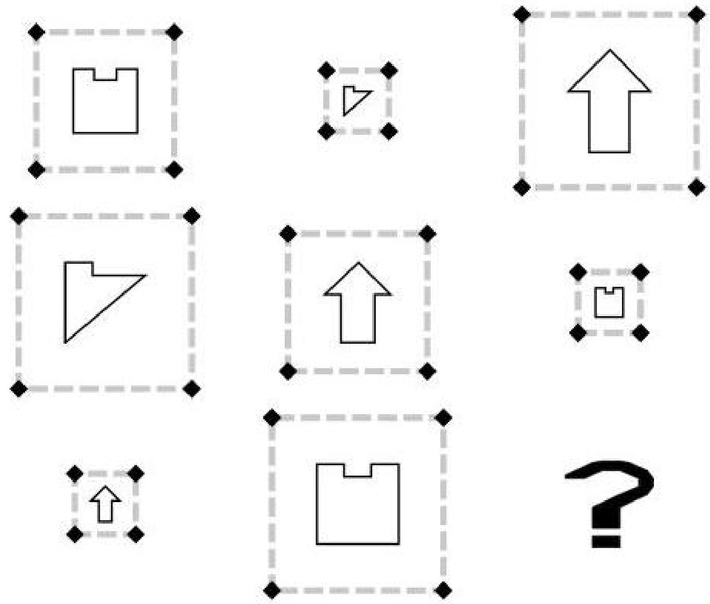

Figure 2: An example of a logic puzzle faced by the subjects.

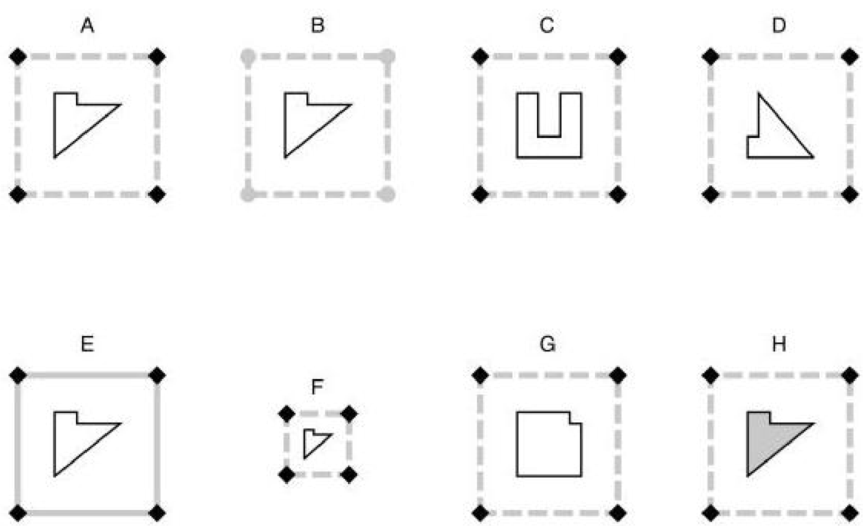

Figure 3: An example of a solution set for the the logic puzzle in Figure2.

In each treatment, subjects face an indefinite number of periods during which they can try to earn as much money as possible. For each treatment, subjects are told that the alloted time is random, with an average of about 20 minutes. The unspecified ending time is designed to mitigate end of game effects on behavior. At the beginning of a period, subjects receive a new an endowment draw. In each period, subjects are allowed to attempt to gain as much additional information about the endowment from the signals as desired. If a subject fails to successfully solve the task to obtain a more precise signal, she is free to attempt that same precision or another level of precision as many time as she wishes. However, once a precision level is successfully obtained, the subject is unable to reduce the precision of the signal.[6]

Subjects accumulate earnings by successfully claiming prizes. Successfully claiming a prize requires the subject to claim a prize which costs less than or equal to the endowment drawn in that period. There are three possible outcomes when a subject attempts to claim a prize:

1. The subject claims a prize of value less than or equal to the drawn endowment. In this case the subject earns the value of the prize claimed and moves on to the next period.

2. The subject successfully obtains additional information about her endowment, but chooses a prize with a value that is greater than the endowment draw. In this case, the subject earns zero profit for the period and moves on to the next period.

3. The subject does not obtain additional information beyond what is known at the start of the period, simply that the endowment falls in the range of 1-256, and the subject picks a prize of greater value than the endowment draw. This results in zero earnings and a 60-second reduction in the length of the treatment that is enforced by adding 60 seconds to the elapsed timer.[7] The theoretical model assumes an utility cost that the consumer incurs for zero consumption when it attempts $c > w$. This cost is embedded in the assumed functional form of the utility, CRRA family and log in its limit, which delivers negative infinite utility in case $c = 0$. This assumption is made to prevent the decision maker from adopting a random consumption strategy

[6]This assumption encoded in the experimental setting captures the well-known principle in information theory that information cannot be forgotten. In order to avoid subjects having to rely on cognitive effort to remember the more informative signals, we prevent them from choosing signals with lower precision than the ones successfully obtained.

[7]Incurring the penalty does not mean the subject has to sit idle for 60 seconds. It means the length of time over which the subject can try to earn money is reduced by 60 seconds.

without acquiring any information about income. In the experimental setting the 60-second penalty ensures that random consumption strategy is strictly dominated by an informed consumption strategy, preventing the subjects from utilizing a mechanical strategy in which they meaninglessly guess prizes a large number of times hoping to rapidly accumulate prizes.[8]

Choosing a prize concludes one period; as mentioned above, subjects encounter as many periods as they can within the timeframe of the treatment. At the beginning of each period, the subjects return to the state of only knowing that their income has been drawn from $g(y)$, and they are again able to choose whether or not to pursue more precise information regarding the value of endowment.

3.4 Experimental Treatments

Up to this point the description of the task has focused on the baseline (Treatment 1). In Treatment 2, the endowment is drawn randomly from a uniform distribution on the support [1 256] as in Treatment 1; however, in Treatment 2 the value of prizes are increased by an order of magnitude every (10^{th}) period and decreased by an order of magnitude in all other periods. That is, in periods 11, 22, 33, etc., the value of every prize shown on the screen is multiplied by 10 in comparison to the prizes shown in Treatment 1. For the other periods the displayed prizes are one tenth the amount displayed in Treatment 1. The ratio of high payoff to low payoff periods is such that the expected payoff is the same across all treatments. In Treatment 3, the mapping of prize value to profit was the same as in Treatment 1, but the endowment is determined differently. For the first period, the endowment is drawn from the usual uniform distribution. For each subsequent period, there is an 80% probability that the endowment is the same as it was in the previous period and a 20% chance that the endowment is determined by a new draw from the uniform distribution.

4 Results

In this section we introduce five testable predictions we obtain from the comparative statics based on the solution of the static model in (7). We use the data from the laboratory to test these predictions in each of the sections below.

[8]In the experiment the penalty is levied to the participants who attempt $c > w$ either without trying any signal or unsuccessfully trying a signal. The two cases are observationally equivalent in terms of information effectively acquired to reduce uncertainty with respect to the initial prior.

4.1 Information and consumption

The first testable prediction stemming from the model is that the more information agents acquire, the higher their consumption is. We test this prediction by examining the relationship between precision of information chosen by the participants and their realized consumption outcomes. In our experiment, precision of information relates to the cognitive tasks participants have to solve to know where their income lies. The harder the task is, the smaller the width of the interval revealing where the actual income is. We test whether a participant solving more difficult cognitive tasks enjoy higher profits than a participant who selects easier tasks.

Figure 4 illustrates the relation between successful consumption choices and the signal precision acquired, for all subjects in Treatment 1 of the experiment. The positive relation is clearly shown by the estimated OLS fitted line; the estimate slope coefficient is 16.9 and it is significant at 1%. Individuals with better information are also able to obtain higher successful realizations of consumption.

The figure also summarizes a couple of interesting properties of the consumption and precision decisions across individuals. First, we can clearly observe how the consumption choices are clustered at the intervals determined by each signal level. The points favor the lower portion of the intervals, as we would expect even for risk neutral subjects, but they are actually spread out and not concentrated on the lower extreme of the interval. We will go back to this point in Section 4.2 too. Second, the majority of the signals chosen by the subjects are for the first levels of precision up to level 4. Precision 5, 6, and 7 get progressively more infrequent as the cognitive cost of acquiring the signal gets higher. Level 7 is virtually never chosen and Level 8 is never successfully completed indicating it correctly approximates an infinite information capacity situation.

Similarly, Figure 5 shows the scatter plot between average precision acquired by each subject and average successful consumption in the baseline, Treatment 1. The color of the dots indicates the average time in seconds spent by the subject to acquire information before making the consumption choice, based on the scale on the right hand side of the figure. The plot confirms a strong positive relation between consumption and precision too, with an estimated coefficient of 14.1 (significant at 1%). As expected, we can also observe a general increase in the period length as the precision of the signal goes up.

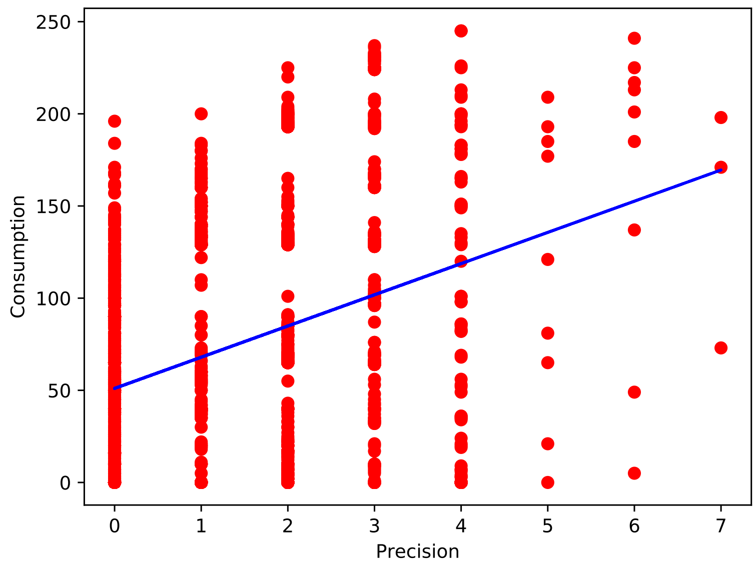

Figure 4: Scatter plot between consumption choices and the signal precision acquired. All subjects and all treatments.

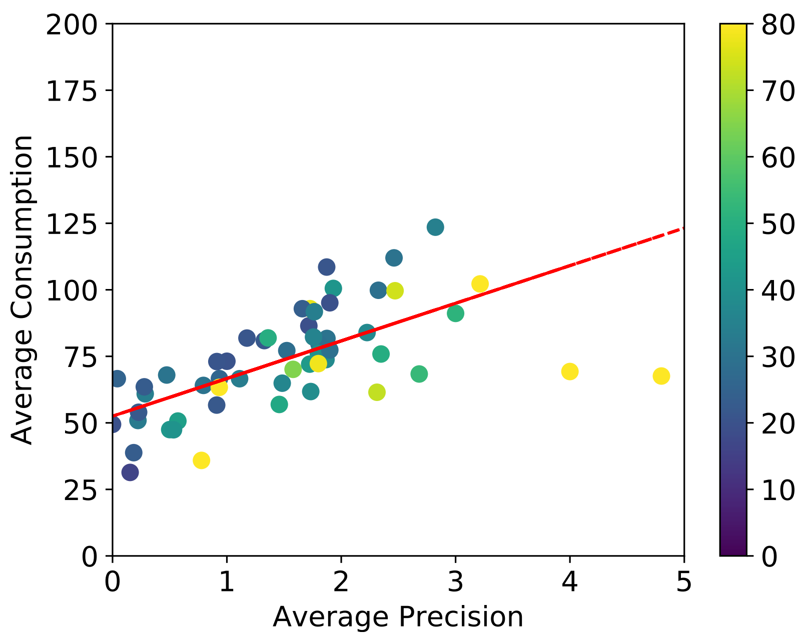

Figure 5: Treatment 1: Scatter plot between average successful consumption and average signal precision acquired by each subject. The color-bar indicates the average time in seconds spent for information acquisition by the subject before a consumption choice is made.

4.2 Restricted consumption support

The second testable prediction of the model is that the optimal probability distribution $p^*(c, w)$ assigns positive probability to only an handful of values out of the full the support of c. A rationally inattentive consumer focuses only on a restricted subset of c, placing higher probability of selecting within that restricted support than everywhere else.

This prediction from the rational inattention model is observationally distinct from other theories of values where either the mean of the support is chosen with probability one or all the values in the support of c are equally likely to be chosen. Moreover, unlike habit formation, where an ad-hoc shape of utility is needed to elicit the behavior predicted by the theory, rational inattention derives this prediction from people's optimizing behavior of balancing cost and benefit of information gathering.

In the experimental setup, we test for this prediction by looking at the empirical distribution of consumption selections conditional on the signal received by the subjects. If the prediction holds, choices of consumption would display a restricted support; within that support, the consumption distribution would not take on the form of a uniform distribution since some outcomes will be more likely to be selected than others and some will not be selected at all, nor would it be a degenerate distribution with all probability concentrated on one particular value of the support.

We next discuss some evidence that corroborates this prediction. We begin with the simple histograms of the aggregate distributions of consumption choices made by all subjects in all treatments, reported in Figure 6. Each panel of the Figure refers to a precision level in which the different intervals corresponding to that precision level – disjoint, but of the same length – are superimposed. We focus on precision 1 to 6, since we do not have sufficient observations to generate reliable plots for higher precision levels.

The distributions typically exhibit higher densities on the lower portion of the intervals; however, we also find substantial distribution density on the rest of the segment. This evidence suggests that the empirical distribution of the subjects' consumption choices is left-skewed – reflecting moderate risk aversion. However, the distribution is not spurious and it is not concentrated at a single point common to all subjects. Similarly, the optimal distributions of consumption are not uniform either, a difference from the prior.[9]

[9] The distribution of Precision 3 exhibits a density peak at position negative one. This position indicates a consumption choice outside the interval revealed by the signal, corresponding to the upper extreme of the interval preceding the revealed one. This concentration of the density reflects the behavior of an individual who acquired precision 3 signals multiple times and systematically chose one unit of consumption less than the minimum allotted by the revealed signal. There are three other subjects for which this behavior occurs

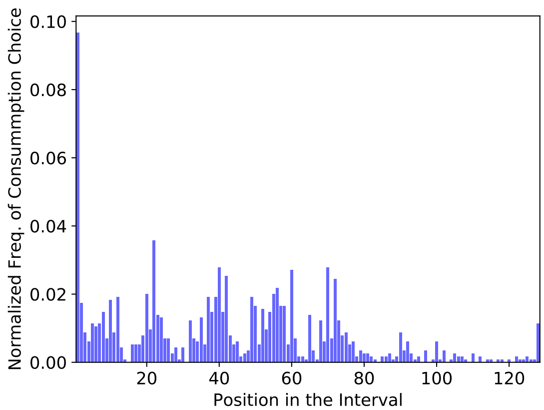

(a) Precision 1

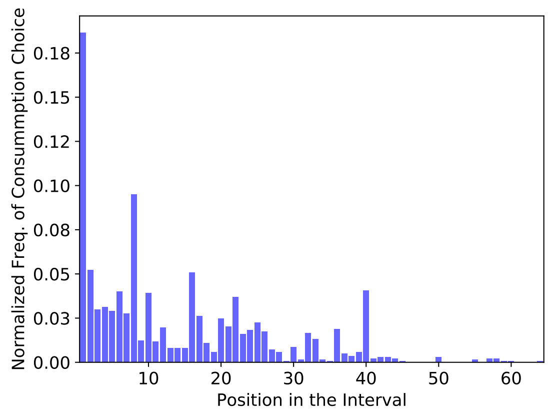

(b) Precision 2

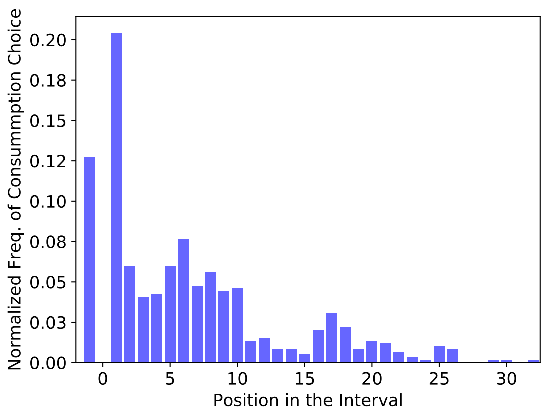

(c) Precision 3

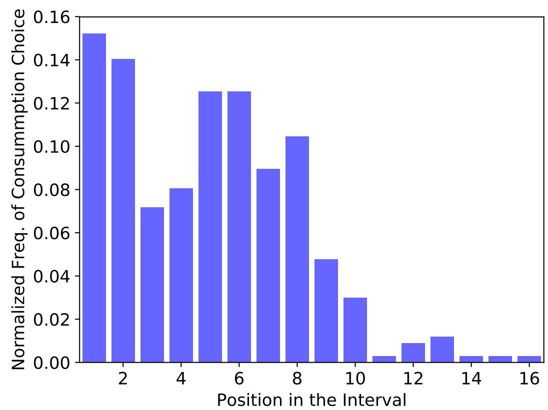

(d) Precision 4

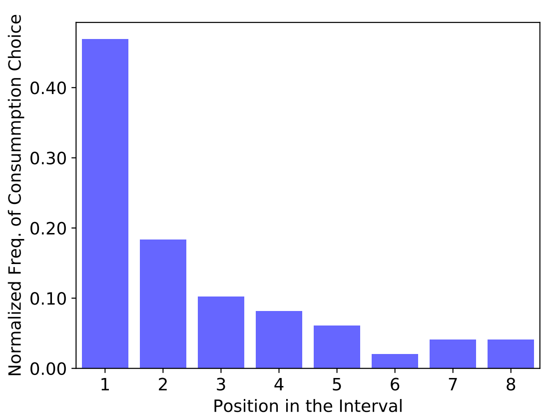

(e) Precision 5

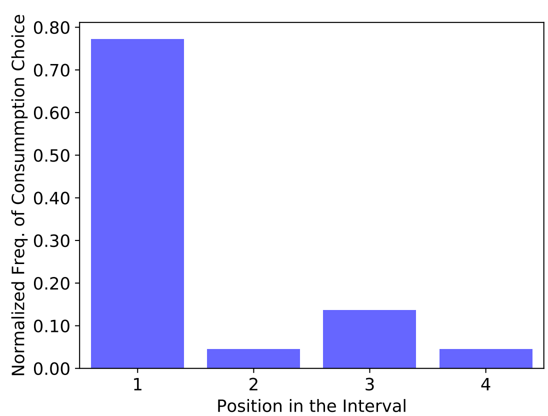

(f) Precision 6

Figure 6: Aggregate distributions of consumption choices for all subjects by signal precision. The intervals corresponding to a signal that spans different portions of the consumption space are superimposed to form a common distribution support.

These histograms help us show some properties of the observed distribution of consumption, but they are based on aggregate data for all subjects and hide individual choices. We consider distributions at the individual level, then, in order to fully assess the restricted support prediction. In fact, individual subjects might select only an handful of consumption points, but if these points are distinct from one subject to the next then aggregation would cause overall distribution to look more dispersed. We assess this possibility by reporting in Figure 7 the Herfindahl-Hirschman Index (HHI) of the individual consumption distributions for each subject for signal precision 1 to 4.

The HHI is typically used to measure market concentration. It is calculated by summing the square of the market shares of competing firms in an industry, and it has a maximum value of 10,000 when the market is full concentrated in the hands of one single firm. In our situation, we first compute the frequencies of the consumption choices made by a subject over the support defined by the overlapped intervals corresponding to a certain signal precision. We then calculate the HHI of the consumption choices of an individual summing up the square of the frequencies of selected consumption levels. Finally, we normalize the HHI's by 10,000 and order the normalized HHI values in decreasing order to report them in Figure 7.

The HHI is an intuitive and rapid way to provide information about the dispersion of consumption levels chosen by the subject although it is not always possible to uniquely map an HHI value into the exact number of points chosen. To this end, the horizontal red lines in Figure 7 provide some reference levels. For instance, when an HHI bar is equal to 1, the subject chose the same point every time he received a signal with that precision. The red line at a height of 0.5 denotes the normalized HHI for a subject who chooses each of two points with equal frequency. The following lines follow the same principle: 0.25 for four choices each with a $\frac{1}{4}$ share, .1 for 10 choices each with a $\frac{1}{10}$ share. The figure omits subjects that make only one consumption choice for which the HHI is trivially 1.[10] We include subjects that make two choices to check whether their HHI differs from the trivial result of 0.5. These cases, HHI=1 with two choices, are displayed in red bars.

The subjects are ordered on the horizontal axis. Since the number of signals is endogenous the number of subjects varies by precision level. The histograms show that the majority of subjects, do not focus exclusively on a single consumption level for a given precision. Further, almost half of the subjects spread their choices among four or more consumption choices within the revealed interval (HHIs fall below .25).

episodically. We judge those instances as mistakes as opposite to deliberate choice and do not include them in Figure 6.

[10]There are a total of 23 such cases out of 122 subject-precision observations in the sample of Figure 7.

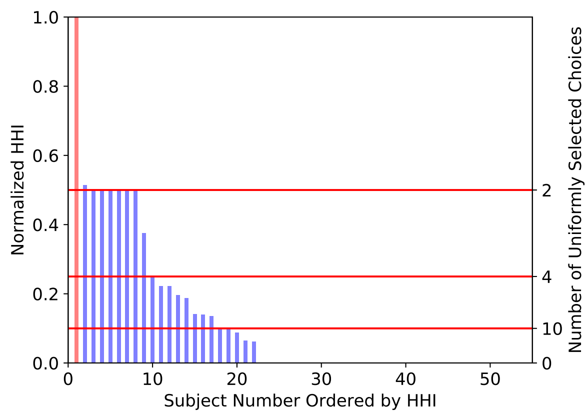

(a) Precision 1

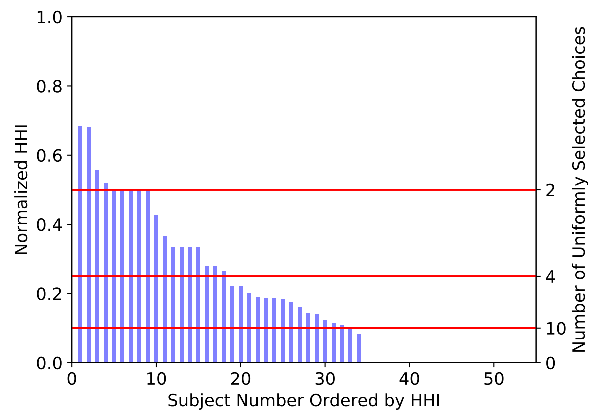

(b) Precision 2

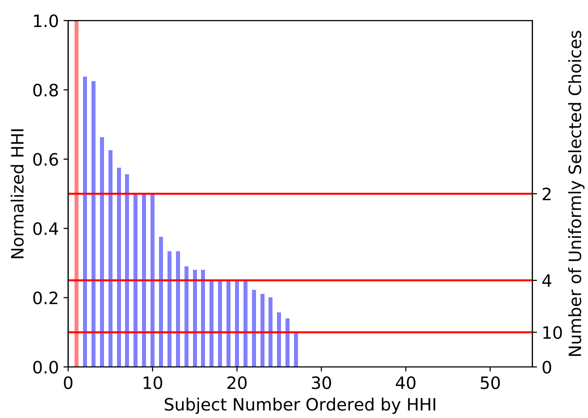

(c) Precision 3

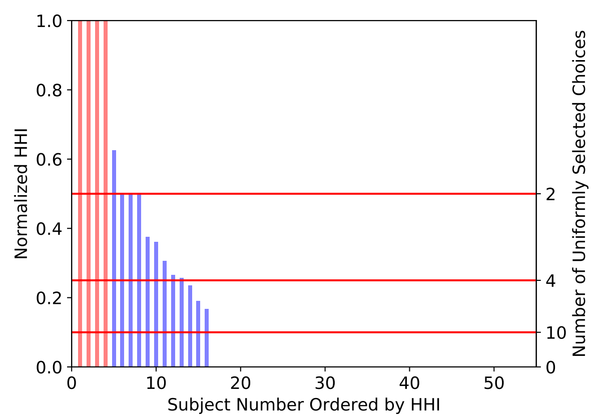

(d) Precision 4

Figure 7: Herfindahl-Hirschman Index (HHI) of the individual consumption distributions for all subjects for signal precision 1 to 4. The indexes are normalized by 10,000 and sorted in decreasing order. The horizontal lines indicate the HHI value that would occur if the subject were to choose the number of points indicated on the right vertical axis with the same frequency. Subjects who make a single consumption choice are omitted. Red bars indicate subjects with two consumption choices and HHI=1.

Overall, the evidence of this section strongly suggests both the transformation of the prior distribution into a non-uniform and non-degenerate optimal consumption distribution, and the concentration of consumption decisions on relatively small sets of choices, but not on a single choice.

4.3 Convex Payoffs and Information Processing

The third testable prediction from the model concerns the effects of incentives on subjects' information acquisition and information processing. The prediction states higher payoffs are associated with more informed decisions. Intuitively, the prediction states that individuals pay more attention to the economic environment the higher the stakes are. We verify this prediction in Treatment 2, in which we vary the concavity of the payoffs by varying their magnitude across periods within the Treatment. As we explained in Section 3, the payoff variations occur at a pre-determined and known frequency: every eleventh period.

We assess the implications of payoff fluctuations in Figures 8 and 9. In these Figures, the red bars refer to signals pursued on the "eleventh" periods, when we impose higher payoffs in Treatment 2, while the blue bars correspond to the other periods, when we have lower payoffs in Treatment 2. It is worth keeping in mind that the payoffs differ by a factor of 100 between the high and low stakes periods of Treatment 2. Figure 8 illustrates the distribution of the maximum signal precision that subjects attempted to acquire before each consumption decision, regardless of the successful attainment of the signal. All subjects and all information acquisition instances are pooled together by treatment, and the distribution in Treatments 2 is compared to the baseline distribution in Treatment 1.

Two observations are noteworthy. First, as expected, there should not be differences between the eleventh and off-eleventh period distributions in Treatment 1, since the subjects do not observe any change in the payoff structure in this Treatment. A KS test to compare the two distributions cannot reject the null that empirical distributions are the same (p-value of .88). In Treatment 2, the two distribution noticeably move apart from each other. The high payoff distribution shifts to the right, and reaches its peak for precision level 3. This shows how subjects seek better information in the high payoff periods in Treatment 2. The regular payoff distribution, on the contrary, shifts to the left, relative to the Treatment 1, with the highest mass concentration at precision 0. This reflects the optimal strategy of reducing attention in the low stakes periods, while increasing it in the high stakes moments. The difference in the two distributions is confirmed by the KS test too, which rejects the null at very high confidence level (p-value of .00).

Comparing the distribution across treatments, we find that the mean of the distributions in Treatment 1 is between precision 1 and 2; in Treatment 2 the mean of the high payoff distribution is between precision 2 and 3, while for low payoffs it is between 0 and 1. We formally test whether these distribution are the same across Treatments using the KS test again. The tests reject the null of same empirical distribution for both high and low payoffs states at extremely high level of confidence (p-value of .00).

As a robustness check, Figure 9 illustrates the distribution of the average maximum precision of the signals attempted by each individual by treatment. As before, the eleventh and off-eleventh period distributions are very similar in Treatment 1. There is no reason for subjects to change their behavior in this case since there is no material change in the economic environment from one period to the next in the baseline treatment. By contrast, in Treatment 2, the distribution of the high payoff periods shifts to the right, away from the least informative signal that dominates in the low payoff periods. Thus, this second set of figures corroborates the result that participants trade-off effort when the payoff is small for higher precision when the payoff is more substantial.

As a consequence of the decision of allocating more attention to the economic environment in the high payoff state, consumption should be higher in the "eleventh" periods.[11] We illustrate this point in Figure 10, where we estimate the non-parametric aggregate density distribution of consumption in the high and low payoff periods in Treatment 1 and 2 by using a normal kernel estimator.

In panel (a) of the Figure we see the two densities for Treatment 1 closely overlap. The density for the non-eleventh periods is less smooth, though, since the larger number of observations allows for a more precise estimation of the density. On the contrary, the densities in panel (b) for Treatment 2 exhibit two distinguishable density peaks. For the high payoff distribution the density maximum is achieved around 130 unit of consumption, while for the low payoff density this peak is shifted far to the left indicating a high density concentration at poor consumption levels.

A similar result is conveyed by Figure 11 where we plot the aggregate density distributions of unused income, as a percentage of the income draw – a measure of the "error" in consumption made by the subjects. Again, the two densities are the same for Treatment 1, while they have a completely different shapes in Treatment 2. In particular, we find a mass concentration at very low and very high levels of unused income respectively for the high

[11] This is obviously the case if the attempted signal precision reported in Figure 8 lead to successful information acquisition too. In our data, we find that about 85% of the signal attempts are successful.

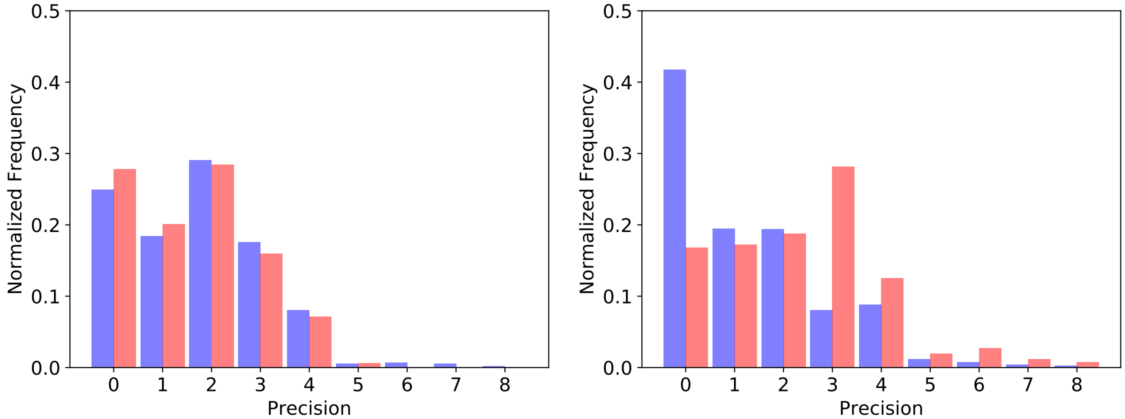

(a) Treatment 1 (b) Treatment 2

Figure 8: Distributions of the maximum precision of attempted signals – Treatment 1 Vs. Treatment 2. Period numbers that are a multiple of 11 in red while other periods are in blue. All subjects, all attempts.

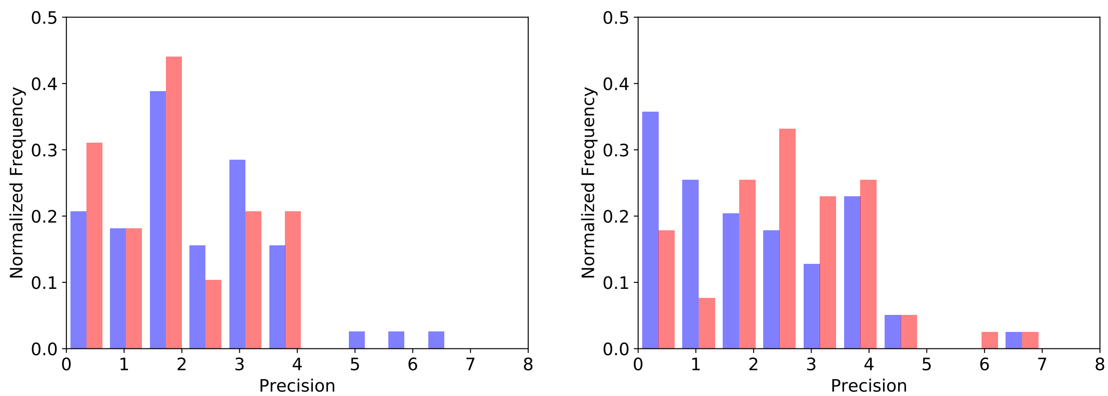

(a) Treatment 1 (b) Treatment 2

Figure 9: Distributions of the average maximum precision of the signals attempted by each subject – Treatment 1 Vs. Treatment 2. Period numbers that are a multiple of 11 in red while other periods are in blue. All subjects, averages by subject.

and low payoff periods.

In summary, the results of this Section imply that participants with limited ability to process information choose more information about options that provide them with higher utility and, as a consequence, make better consumption decisions. This result is intuitive since a higher reward implies an higher benefit to process more precise information offsetting its cost. Consistent with these findings, the optimal probability distribution $p^* (c, w)$ is more informative the higher the payoff is.

4.4 Predictability of economic environment

The fourth prediction posits that a rationally inattentive person responds to a more predictable environment – higher persistence of $g (w)$ – by processing less information. That is, the optimal probability distribution of consumption conditional on the signal chosen by the participant, $p^* (c|w)$, is more dispersed in the case in which the environment is less uncertain, reflecting either a less precise signal acquired in the more predictable environment or more "mistakes" in consumption when learning about the environment turns out to be more costly than forgoing consumption units. We rely on Treatment 3 to test this prediction.

This prediction can have two different outcomes on the behavior of the subjects in the experiment. One the one hand, they could periodically increase signal precision, when they realize an income shock has occurred, to reset information about their income level. After acquiring more information, they would maintain a low precision signal relying on the high persistence of the process. On the other hand, if the cost of increasing precision level is sufficiently high, they could opt for a low precision level by default and use consumption choices to explore the income space. In this case, they would gropingly increase consumption after a positive income shock, trying to figure out the new higher income. Similarly, they would progressively reduce consumption after a negative shock, trying to detect the new lower income.

As we saw in Section 4.3, the cognitive cost of the signals necessary to acquire better information is relatively high in our experimental setup. In Treatment 2, for instance, the ten fold increase in the magnitude of payoffs was sufficient to move up precision of the signal by roughly only one notch. This tendency is confirmed by Treatment 3 as well, and the second type of behavior described above prevails. We illustrate this point in the two panels of Figure 12, where the average income, consumption, and precision time profiles in the neighborhood of an income shock are reported for all subjects in Treatment 3. The figures are re-centered on the period in which the shock occurs, time 0 on the horizontal axis. The

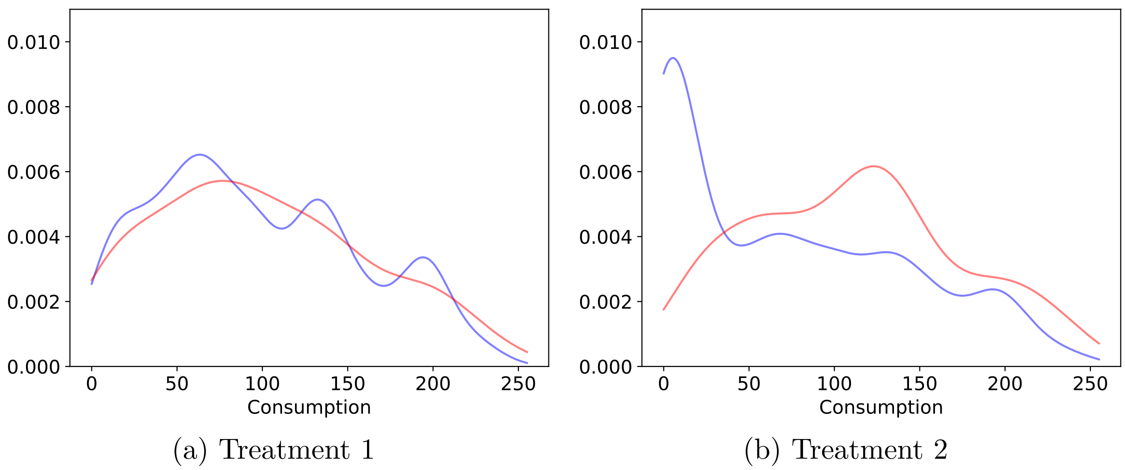

Figure 10: Aggregate density distribution of successful consumption choices – Treatment 1 Vs. Treatment 2. Estimation by normal kernel function. High payoff periods in red; low payoff periods in blue. All subjects.

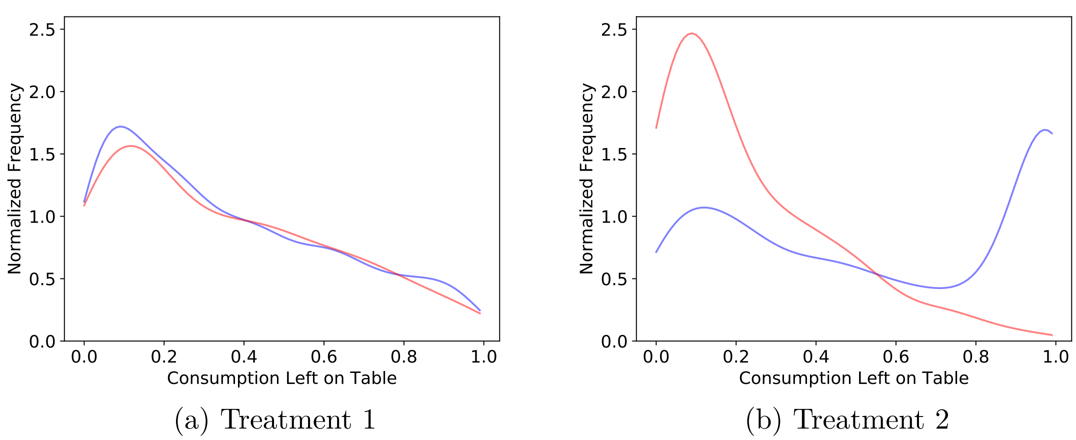

Figure 11: Aggregate density distribution of unused income, as a percentage of the income draw – Treatment 1 Vs. Treatment 2. Estimation by normal kernel function. High payoff periods in red; low payoff periods in blue. All subjects.

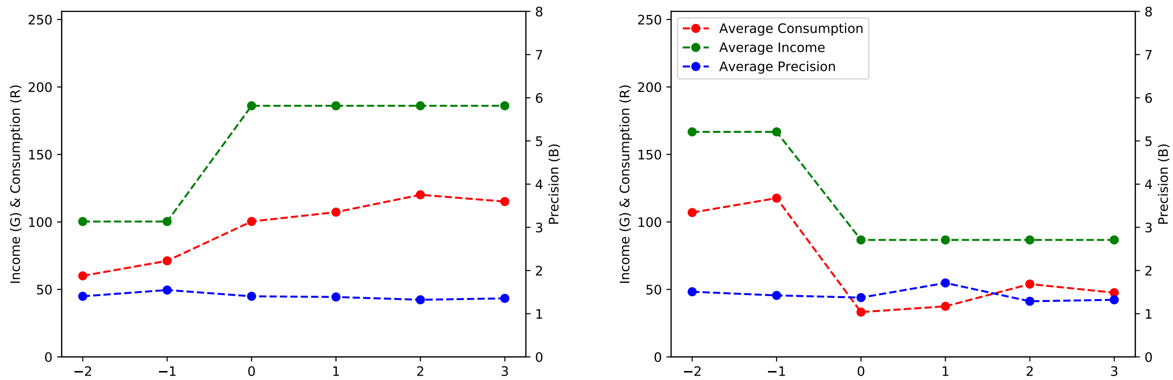

(a) Positive shocks

(b) Negative shocks

Figure 12: Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

time series show the two periods before the shock and the three periods after it. We consider all episodes followed by at least three periods of constant income level. The left panel refers to the positive shocks, while the right panel to the negative ones.

While income jumps up with a positive shock (green line), the precision level minimally varies (blue line). However, consumption (red line) follows the direction of the shock and it is used by the subjects in place of a change in precision. In case of a negative shock, though, the response is more hybrid. Consumption drops at time 0, following income. Precision is increased in the next period, for just a period to provide one with a better sense of the new income level. After that, subjects start slowly moving up their consumption choices to learn more about income, as done for a positive shock.

We also find, however, the behavior of the subjects can be quite idiosyncratic. Some of them, probably those who have higher cognitive capacity, prefer to manipulate the precision level after a shock, as we document in Appendix C.

4.5 Asymmetry in delays

The fifth proposition deals with the differences in subjects' responses to positive and negative income shocks. The theory predicts an asymmetric response to shocks of opposite sign. Although subjects naturally exhibit a delay in the response regardless of the type of shock, the response to bad income shocks should be faster and sharper than to good income shocks. Intuitively, this is because bad shocks elicit the allocation of more attention from the subjects.

This prediction, which is unique to RI with respect to other mainstream theories, can be tested with our results comparing subjects' reaction to a period in which income falls with the reaction to a period in which income increases. Basically, a negative shock to income should be detected faster since previous consumption choices may no longer be available, whereas if the participant is sticky in its consumption choices an increase in income may go undetected for several periods.

Figure 13 replicates Figure 12 for two levels of shock intensity. The first row refers to shocks larger than two thirds of the income support (i.e. 162 units of income); the second row refers to shocks smaller than one third of the support (smaller than 94 units of income). Considering different sizes of the shocks helps us highlight differences in the speed of the responses.

We find that the responses of consumption and precision to income shock closely follow the predicted asymmetry. In the first row, a large negative shock causes an immediate correction of consumption that rapidly drops and remain low. On the contrary, subjects respond only gradually to a large positive shock. The adjustment process is indeed quite slow, and a faster correction would require subjects to increase the precision of the acquired signal. For small shocks, illustrated by the last row of the same Figure, the subjects' response is less sharp and milder for the positive shock, while it is still quite rapid and deep for the negative shock. Very interestingly, the negative shock induces a clear increase in the precision level the first period after the shock, which is totally missing for positive shocks. The negative shocks are perceived as more worrying, so that they also require an adjustment in the signal precision.

Overall, the size and timing of the reactions to income shocks is strongly supportive of this prediction of the model.

5 Concluding Remarks

The paper shows that subjects in a controlled laboratory experiment largely behave in accordance with the predictions of rationally inattentive consumption theory. Specifically, we find that subjects

- consume stochastically, with their consumption choices reflecting a draw from their optimal posterior distribution of consumption conditional on signals about income.

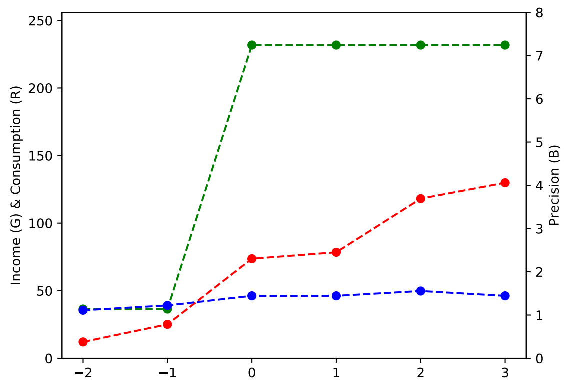

(a) Large positive shocks

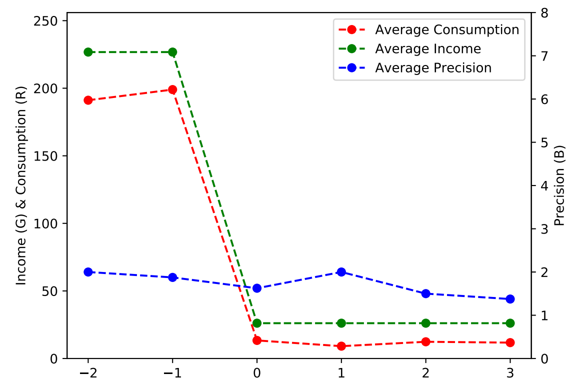

(b) Large negative shocks

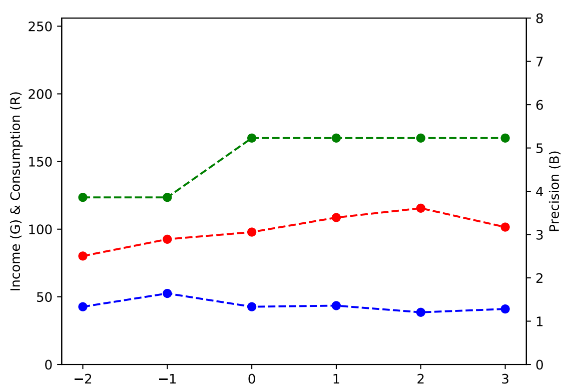

(c) Small positive shocks

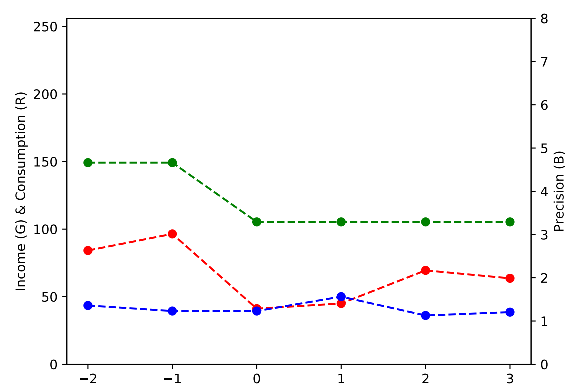

(d) Small negative shocks

Figure 13: Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. Large shocks are larger than 162 units of income; medium shocks are between 94 and 162 units of income; small shocks are smaller than 94. All subjects in Treatment 3.

- respond to incentives by varying the informativeness of their optimal posterior toward higher consumption values, consistent with the postulates of elastic information processing capacity.

- react asymmetrically to shocks of income, with stronger and faster adjustment of consumption to negative shocks than positive ones.

The last of these findings is particularly noteworthy because, to the best of our knowledge, it is the first direct evidence of a pattern that is predicted by rational inattention, but not by other models.

We map the mutual information between prior and posterior beliefs by rigorously measuring attention and information processing using cognitive tasks and consumption choices. As research on behavioral modeling strategies focusing on information processing and attention advances, this paper constitutes a step towards bridging the gap between theoretical models and applied measurements.

References

Abaluck, J., and A. Adams (2017): "What do consumers consider before they choose? Identification from asymmetric demand responses," *NBER Working Paper No. 23566*.

Broda, C., and J. A. Parker (2014): "The Economic Stimulus Payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics, Elsevier*, 68(5), S20–S36.

Caplin, A., and M. Dean (2015): "Revealed preference, rational inattention, and costly information acquisition," *The American Economic Review*, 105(7), 2183–2203.

Cheremukhin, A., A. Popova, and A. Tutino (2015): "A theory of discrete choice with information costs," *Journal of Economic Behavior and Organization*, 113, 34–50.

Chetty, R., A. Looney, and K. Kroft (2009): "Salience and taxation: Theory and evidence," *The American Economic Review*, 99(4), 1145–1177.

Civelli, A., and C. Deck (2018): "A Flexible and Customizable Method for Assessing Cognitive Abilities," *Review of Behavioral Economics, Forthcoming*.

de los Santos, B., A. Hortasu, and M. Wildenbeest (2012): "Testing models of consumer search using data on web browsing and purchasing behavior," *The American Economic Review*, 102(6), 2955–2980.

Dean, M., and N. Neligh (2017): "Experimental Test of Rational Inattention," *Mimeo, Columbia University*.

DellaVigna, S., and J. Pollet (2007): "Demographics and industry returns," *The American Economic Review*, 97(5), 1667–1702.

Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006): "Costly information acquisition: Experimental analysis of a boundedly rational model," *The American Economic Review*, 96(4), 1043–1068.

Hellwig, C., S. Kohls, and L. Veldkamp (2012): "Information choice technologies," *The American Economic Review*, 102(3), 35–40.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006): "Household Expenditure and the Income Tax Rebates of 2001," *The American Economic Review*, 96(5), 1589–1610.

Khaw, M. W., L. Stevens, and M. Woodford (2016): "Discrete adjustment to a changing environment: Experimental evidence," *Technical report, National Bureau of Economic Research.*

Mackowiak, Bartosz, and M. Wiederhold (2009): "Optimal Sticky Prices under Rational Inattention," *The American Economic Review*, 99, 769–803.

Manzini, P., and M. Mariotti (2015): "Stochastic choice and consideration sets," *Econometrica*, 82(3), 1153–1176.

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

Mondria, J. (2010): "Portfolio choice, attention allocation, and price comovement," *Journal of Economic Theory*, 145(5), 1837–1864.

PInkofskiy, M. (2009): "Rational inattention and choice under risk: Explaining violations of expected utility through a Shannon entropy formulation of the costs of rationality," *Atlantic Economic Journal*, 37(1), 99–112.

Shea, J. (2009): "Myopia, Liquidity Constraints, and Aggregate Consumption: A Simple Test," *Journal of Money, Credit, and Banking*, 1(1), 229–257.

Shue, K., and E. Luttmer (2009): "Who misvotes? the effect of differential cognition costs on election outcomes," *American Economic Journal: Journal Policy*, 1(1), 229–257.

Sims, C. A. (1998): "Stickiness," *Carnegie-Rochester Conference Series on Public Policy*, 49(1), 317–356.

——— (2003): "Implications of Rational Inattention," *Journal of Monetary Economcis*, 50(3), 665–690.

Tutino, A. (2013): "Rationally inattentive consumption choices," *Review of Economic Dynamics*, 16(3), 421–439.

Appendix

A The theoretical Model

The theoretical model borrows from Tutino (2013). An agent faces a consumption decision under uncertainty about her wealth and acquires information about the distribution of her income under an information processing capacity constraint. To understand the implications of limits to information processing, we start with the standard, full information version of the problem.

Let $(\Omega, \mathcal{B})$ be the measurable space, where Ω represents the sample set and $\mathcal{B}$ the event set. States and actions are defined on $(\Omega, \mathcal{B})$. Let c_t, c_t, and w_t respectively indicate consumption, income, and wealth at time t. Let $\mathcal{I}_t$ be the $\sigma-$algebra generated by $\{c_t, w_t\}$ up to time t, i.e., $\mathcal{I}_t = \sigma(c_t, w_t; c_{t-1}, w_{t-1}; ...; c_0, w_0)$. Then, the collection $\{\mathcal{I}_t\}_{t=0}^{\infty}$ such that $\mathcal{I}_t \subset \mathcal{I}_s$ $\forall s \geq t$ is a filtration.

Let $u(c)$ be the utility of the household defined over a consumption good c.[12] The consumer's problem is:

$$\max_{\{c_t\}_{t=0}^{\infty}} \quad E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t) \middle| \mathcal{I}_0 \right\} \tag{8}$$

$$\text{s.t.} \quad w_{t+1} = R(w_t - c_t) + y_{t+1} \tag{9}$$

$$w_0 \text{ given} \tag{10}$$

where $\beta \in [0, 1)$ is the discount factor and $R = 1/\beta$ is the interest on savings, $(w_t - c_t)$. We assume that $y_t \in Y \equiv \left\{y^1, y^2, .., y^N\right\}$ follows a stationary Markov process with constant mean $E_t(y_{t+1}|\mathcal{I}_t) = \bar{y}$.

Consider now a consumer who cannot process all the information available in the economy to track precisely her wealth. This not only adds a constraint to the decision problem, but fundamentally affects each constraint (9)-(10) in four important ways.

1. Since the consumer does not know her wealth, (10) no longer holds. Her uncertainty about wealth is given by the prior $g(w_0)$.

[12] We can assume that the utility belongs to the CRRA family, $u(c) = c^{1-\gamma}/(1-\gamma)$ with γ the coefficient of risk aversion.

2. Before processing any information, consumption is also a random variable. This is because the uncertainty about wealth translates into a number of possible consumption profiles with various levels of affordability. It follows that to maximize lifetime utility, consumers need to jointly reduce uncertainty about wealth and choose consumption. The consumer chooses the joint distribution $p(c, w)$.

3. With respect to the program (8)-(10), there is a new constraint on the amount of information the consumer can processes. The reduction in uncertainty conveyed by the signal depends on the attention allocated by the consumer to track her wealth. We append the Shannon channel in equation (11) to the constraint sets. The information flow available to the consumer is a function of the signal, i.e. the joint distribution $p(\cdot_{c_t}, \cdot_{w_t})$. In formulae:

$$\kappa_t \geq I\left(p\left(\cdot_{c_t}, \cdot_{w_t}\right)\right) = \int p\left(c_t, w_t\right) \log\left(\frac{p\left(c_t, w_t\right)}{p\left(c_t\right) g\left(w_t\right)}\right) dc_t dw_t \tag{11}$$

4. The update of the prior replaces the law of motion of wealth using the budget constraint in (9). To describe the way individuals transit across states, define the operator $\mathbb{E}_{w_t}\left(\mathbb{E}_t\left(x_{t+1}\right)\middle| c_t\right) \equiv \hat{x}_{t+1}$, which combines the expectation in period t of a variable in period $t+1$ with the knowledge of consumption in period t, c_t, and the remaining uncertainty over wealth. Applying $\mathbb{E}_{w_t}\left(\mathbb{E}_t\left(\cdot\right)\middle| c_t\right)$ to equation (9) leads to:

$$\hat{w}_{t+1} = R\left(\hat{w}_t - c_t\right) + \widehat{\bar{y}} \tag{12}$$

where,

$$\begin{aligned}
\widehat{\bar{y}} &= \mathbb{E}_{w_t}\left(\mathbb{E}_t\left(y_{t+1}\right)\middle| c_t\right) \\
&\equiv \mathbb{E}_{w_t}\left(\mathbb{E}_t\left(y_{t+1}\middle|\mathcal{I}_t\right)\middle| c_t\right) + \left[\mathbb{E}_{w_t}\left(\mathbb{E}E_t\left(y_{t+1}\right)\middle| c_t\right) - \mathbb{E}_{w_t}\left(\mathbb{E}_t\left(y_{t+1}\middle|\mathcal{I}_t\right)\middle| c_t\right)\right] \\
&\overset{LIE}{=} \bar{y} + \mathbb{E}_{w_t}\left[\left(\mathbb{E}_t\left(y_{t+1}\right)\middle| c_t\right) - \left(\mathbb{E}_t\left(y_{t+1}\right)\middle| c_t\right)\right] \\
&= \bar{y}.
\end{aligned}$$

Given the initial prior state $g(w_0)$, the next belief state $g'_{c_t}(w_{t+1})$ is determined by

Bayesian conditioning as

$$g'\left(w_{t+1}\right|c_t) = \int \tilde{T}\left(w_{t+1}; w_t, c_t\right) p\left(w_t|c_t\right) dw_t, \tag{13}$$

which is the transition equation we have in the main text too. In (13), the function $\tilde{T}$ is the transition function representing (12) and the belief state itself is completely observable. Bayesian conditioning satisfies the Markov assumption by keeping a sufficient statistics that summarizes all information needed for optimal control. Thus, (13) replaces (9) in the limited processing world.

Combining all these points, the problem of the household under rational inattention can be expressed by the program in equations (2)-(6) in Section 2.

B Instructions for the Laboratory Experiment

We report below the instruction sets that were provided to the subjects during the experiment:

A. General Instructions

B. Instructions for Treatment 1

C. Instructions for Treatment 2

D. Instructions for Treatment 3

General Instructions

You will be paid in cash at the end of today's study based upon the decisions you make, so it is important that you understand the directions completely. If you have a question at any point, please raise your hand, but otherwise you should not talk or communicate with anyone.

Overview

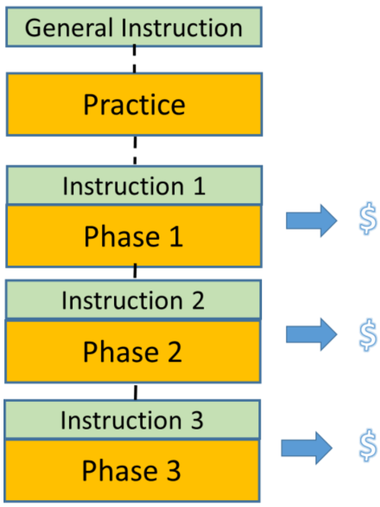

This study involves 3 paid phases that each last about 20 minutes (the exact time for a part is random). During that time you can complete as many periods as possible. Each period you can earn money and you will be paid based on the total amount of money you earn from all three paid phases. Your payment in $US will equal your cumulative earnings from all three phases divided by 500. Each of the paid phases is slightly different and you will be given specific instructions about each phase just before it begins. Before the paid phases you will go through a practice phase to familiarize you with the decision process.

How You Earn Your Payoff

Each period in a phase involves the following sequence of events.

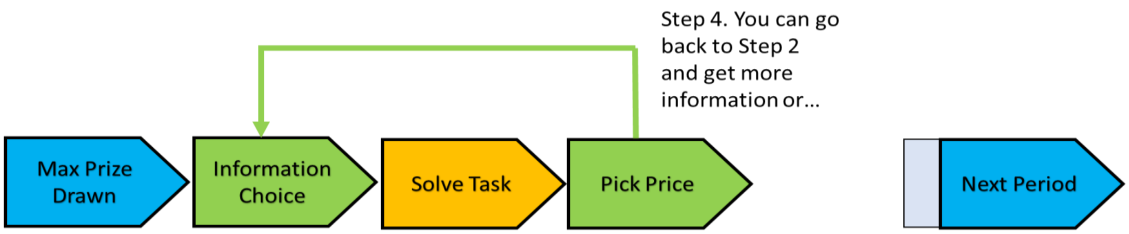

Step 1. The computer determines the maximum prize available that period, but this is not revealed to you.

Step 2. You decide what information you would like to receive about the maximum prize that period.

Step 3. You attempt to solve some logic puzzles to earn information.

Step 4. … Go ahead and pick a prize.
ATTENTION: If you pick a prize less than or equal to the max. prize you earn the amount of the prize you pick. But if the prize you pick is greater than the maximum prize, you earn 0 for the period. So picking a larger prize could increase your payment, but it is riskier. The maximum prize does not change no matter how many times Step 2 is repeated.

Step 5. The next period begins.

More Details About Each Step

You will use the interface (shown below) to interact with the computer. The interface tells you the current period, the elapsed time in seconds, and your accumulated prizes. Then, it shows a grid of possible prizes available in the current and a row of blue buttons to choose the information you want to receive about the prizes.

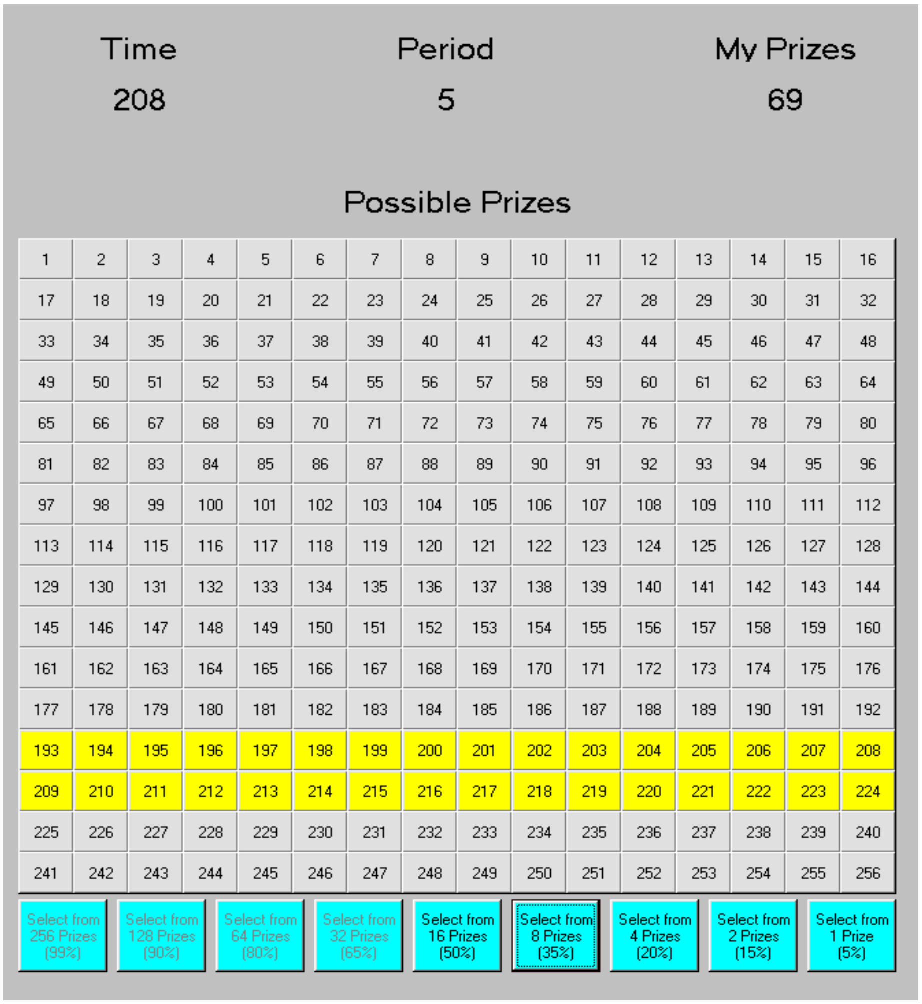

Step 1: Maximum Possible Prize

Each period there are 256 possible prizes. For instance, in the practice phase the prizes are 1, 2, …, 255, 256. The computer will randomly determine the maximum possible prize each period. You will not be told the amount of the maximum prize, but it is the maximum amount you can earn in that period.

Step2: Information about the Maximum Possible Prize

Gathering information allows you to find out more about the maximum prize in the current period. The blue buttons at the bottom of your screen let you pick the type of information you would like to know.

Select from 256 Prizes: You get no additional information and select from all 256 prizes.
Select from 128 Prizes: You learn if the maximum possible prize is in the first 128 prizes or the last 128 prizes.
Select from 64 Prizes: You learn if the maximum possible prize is in the first set of 64 prizes, the second set of 64 prizes, the third set of 64 prizes, or the fourth set of 64 prizes.
…
Select from 1 Prize: You learn the maximum possible prize.

In the screen shot above, we picked Select from 32 Prizes and we found out the maximum possible prize is in the 7th set of 32 prizes: somewhere between 193 and 224.

Step 3: Solve Logic Puzzles

To receive the information selected in Step 2, you have to solve some logic puzzles which work as follows. A 3x3 table of images will be shown to you. Each image has a particular shape, direction, size, color, border edge, and border corner. Some of these characteristics will change from row to row, column to column, diagonally, or from corner to corner. You have to identify the image that belongs in the lower right corner. The blue buttons include information on how frequently people have successfully answered the required puzzles in a previous study. Below is an example of a puzzle. In this example, the shape changes row to row while the size changes from corner to corner.

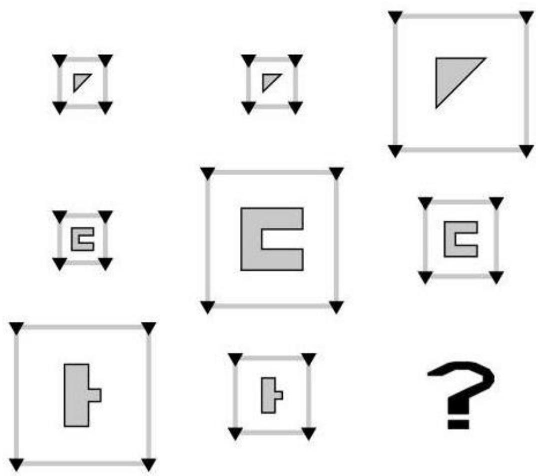

General Instructions

The better information you want to receive the harder the logic puzzles you have to solve. Typically, you have to correctly solve 2 puzzles. One exception is that you have to solve 3 very difficult puzzles to Select from 1 Prize. The other exception is that you only have to answer a single easy puzzle if you want to Select from 256 Prizes. *The blue buttons on your screen indicate how often people who have participated in this study before have been able to correctly solve the necessary puzzles.*

The computer will not tell you if you answer a logic puzzle correctly or not; all you will observe is the yellow range for the maximum possible prize if you solve the task correctly. If you do not solve the task correctly, the full set of 256 prizes will be yellow.

Step 4: Picking a Prize

Once you have obtained all of the information you want then you can pick a prize.

To pick a prize you simple click on it. If the prize you pick is less than or equal to the maximum possible prize then the prize you pick is added to your payoff. Otherwise you earn 0 for the period.

In the example above we know the maximum possible prize is between 193 and 224. Suppose you picked a prize of 199. If the maximum prize was 211 then you would earn 199. But if the maximum prize was 194 then you would earn 0. The maximum prize is not revealed. You will simply see your prize turn green if you receive the prize and red if you do not.

ATTENTION – TRADEOFF: You face a tradeoff in that better information helps you make better decisions about what prize to claim in a period but acquiring less accurate information enables you to claim more prizes during that phase of the study.

ATTENTION - TIME PENALTY: Your time will be reduced by 60 seconds (technically 60 seconds are added to your elapsed time) if you pick from the full set of 256 prizes and are unsuccessful. This can occur either when you decide to Select from 256 Prizes because you did not try to get better information or when you tried for better information but did not correctly solve the required puzzles. When all of the prizes are yellow when you pick a prize the time penalty is in play.

Summary

1. Each period the computer will determine the maximum possible prize.
2. You can solve puzzles to get information about the maximum possible prize. Better information requires you to solve more and harder puzzles.
3. If you pick a prize less than or equal to the maximum possible prize then the prize you pick is added to your earnings. Otherwise you earn 0 for the period.

4. The faster you pick a prize the more periods you will have (but there is a 60 second penalty if you unsuccessfully pick from all 256 prizes).
5. You will complete three paid phases, each of which lasts about 20 minutes. You will be paid based upon your cumulative earnings in each of those phases.

If you, have a question raise your hand. Otherwise, you can press the start button to begin the practice phase.

Instructions for Treatment 1

The next phase of the study will count towards your payment.

This phase is the <u>same as the practice</u> phase you did initially.

1. The 256 prizes are just amounts 1, 2, 3, …., 255, 256 (as in the practice phase).

2. Each period the maximum possible prize that you could claim is randomly determined and equally likely to be any of the show prize amounts (as in the practice phase).

If you, have a question raise your hand. Otherwise, you can press the start button to begin this paid phase of the study. One you press start your time will begin and you cannot pause it.

Instructions for Treatment 2

The next phase of the study will count towards your payment.

This phase is <u>different from the practice</u> phase you did initially.

1. In most periods the prizes are 0.1, 0.2, 0.3, …., 25.5, 25.6, which is one-tenth of what they were in the practice phase. However, every 11th period (that is in period 11, period 22, period 33, and so on) the prizes are 10, 20, 30, …., 2550, 2560, which is ten times what they were in the practice phrase.

2. Each period the maximum possible prize that you could claim is randomly determined and equally likely to be any of the show prize amounts (as in the practice phase).

If you, have a question raise your hand. Otherwise, you can press the start button to begin this paid phase of the study. One you press start your time will begin and you cannot pause it.

Instructions for Treatment 3

The next phase of the study will count towards your payment.

This phase is <u>different from the practice</u> phase you did initially.

1. The 256 prizes are just amounts 1, 2, 3, …., 255, 256 (as in the practice).

2. In the first period, the maximum possible prize that you could claim is randomly determined and equally likely to be any of the show prize amounts. After that, in any period there is an 80% chance that the maximum possible prize does not change from one period to the next and only a 20% chance that a new maximum prize is randomly drawn. Below are two examples of this process.

Example 1

Period	Maximum Prize
1	235
2	235
3	235
4	235
5	66
6	66
7	66
8	66
9	66
10	66
11	66
12	151
…	…

Example 2

Instructions for Treatment 3

Period	Maximum Prize
1	117
2	117
3	117
4	117
5	117
6	117
7	117
8	189
9	189
10	189
11	189
12	189
…	…

If you, have a question raise your hand. Otherwise, you can press the start button to begin this paid phase of the study. One you press start your time will begin and you cannot pause it.

C Individual Results by Subject

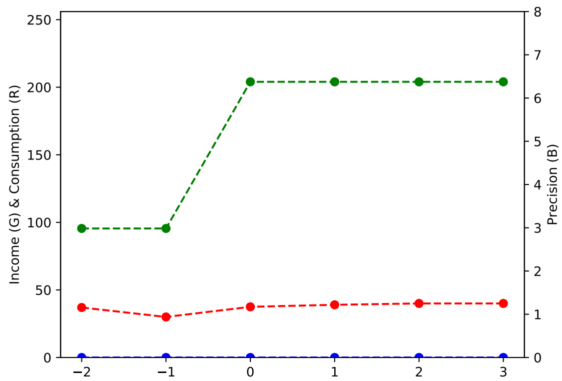

(a) Positive shocks - Sbj 1

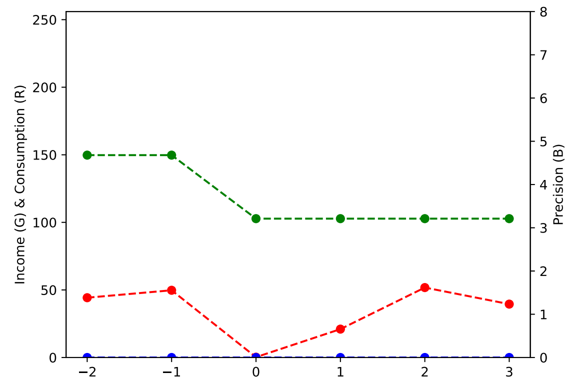

(b) Negative shocks - Sbj 1

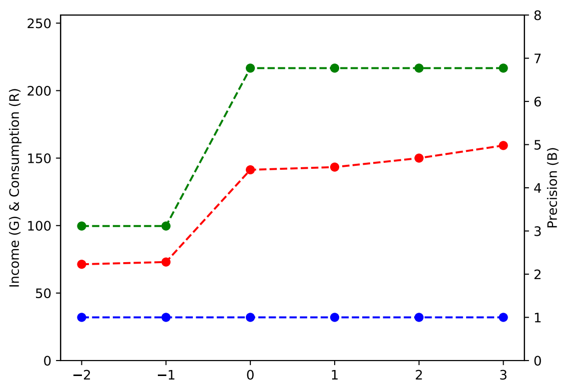

(c) Positive shocks - Sbj 3

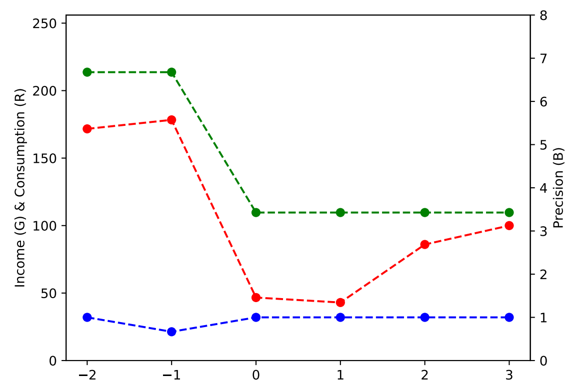

(d) Negative shocks - Sbj 3

Figure A1: INDIVIDUAL RESPONSES TO INCOME SHOCKS 1 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

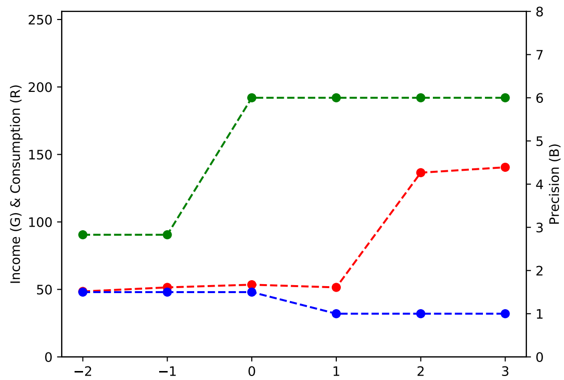

(a) Positive shocks - Sbj 4

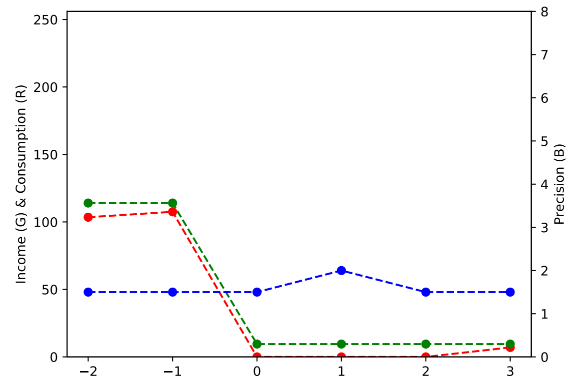

(b) Negative shocks - Sbj 4

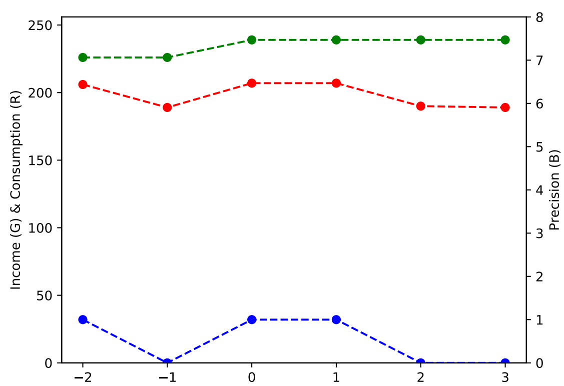

(c) Positive shocks - Sbj 5

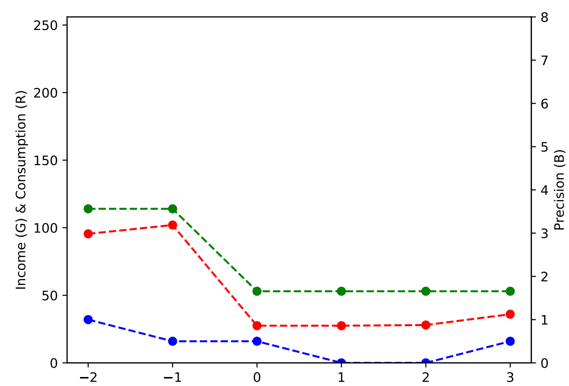

(d) Negative shocks - Sbj 5

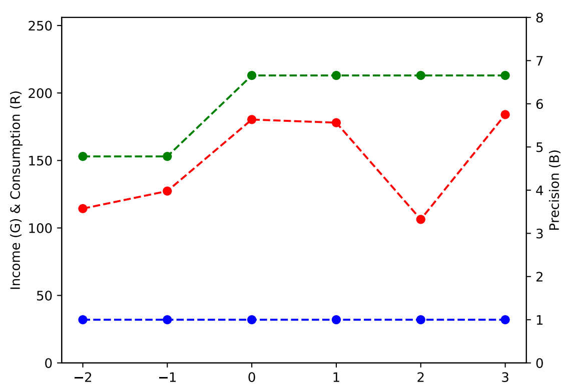

(e) Positive shocks - Sbj 9

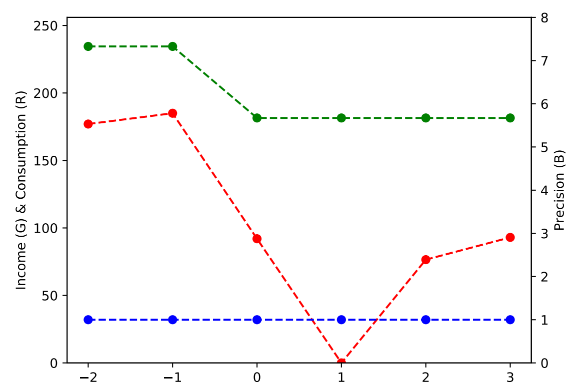

(f) Negative shocks - Sbj 9

Figure A2: INDIVIDUAL RESPONSES TO INCOME SHOCKS 2 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

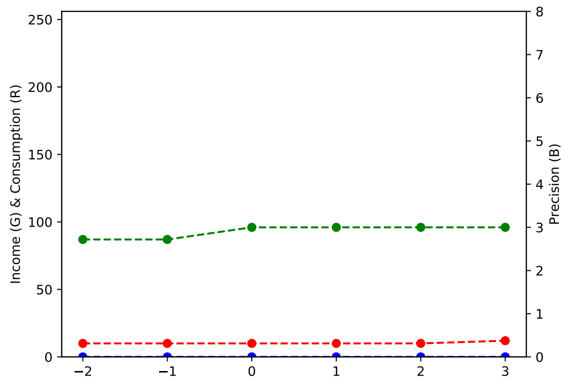

(a) Positive shocks - Sbj 12

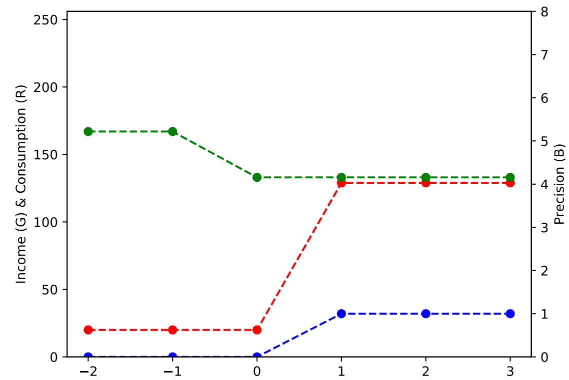

(b) Negative shocks - Sbj 12

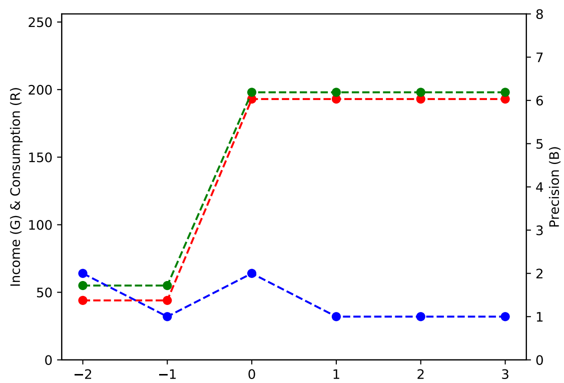

(c) Positive shocks - Sbj 13

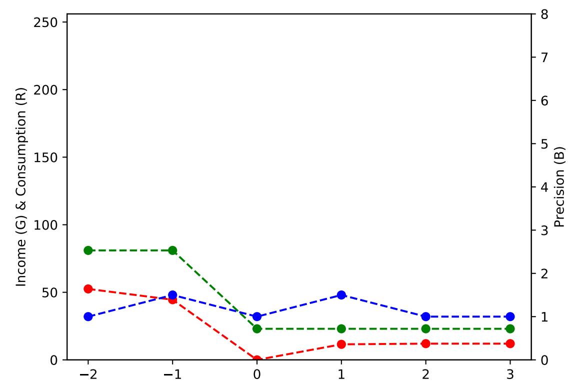

(d) Negative shocks - Sbj 13

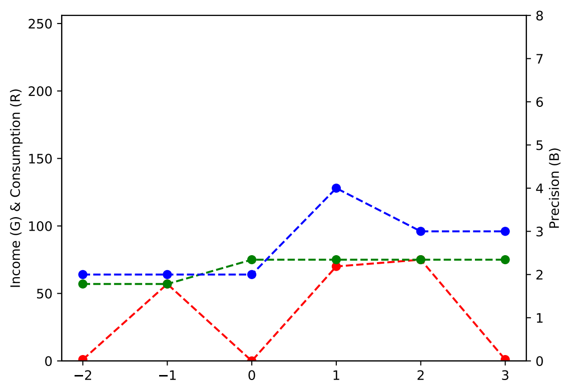

(e) Positive shocks - Sbj 16

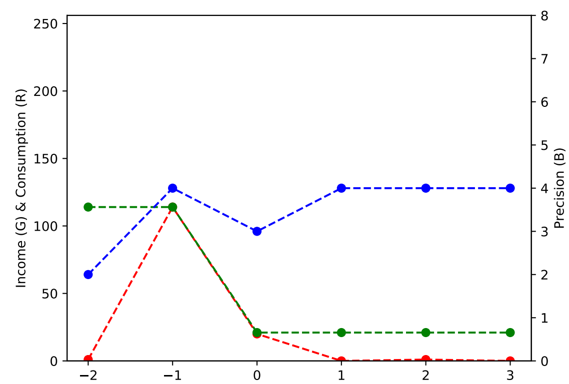

(f) Negative shocks - Sbj 16

Figure A3: INDIVIDUAL RESPONSES TO INCOME SHOCKS 3 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

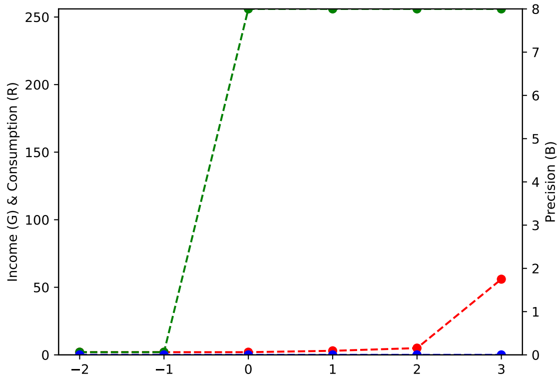

(a) Positive shocks - Sbj 18

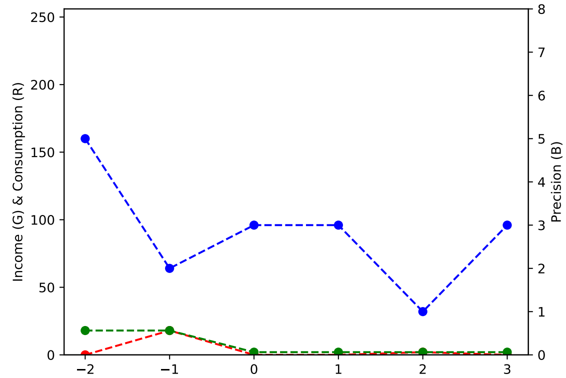

(b) Negative shocks - Sbj 18

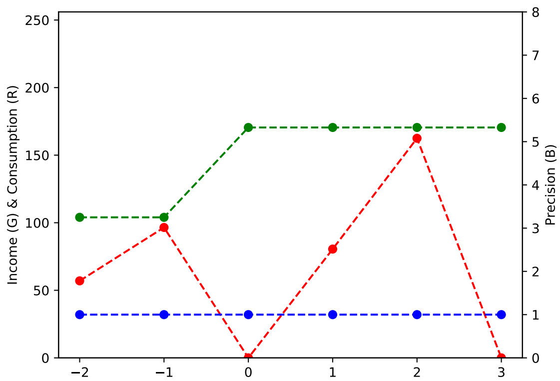

(c) Positive shocks - Sbj 19

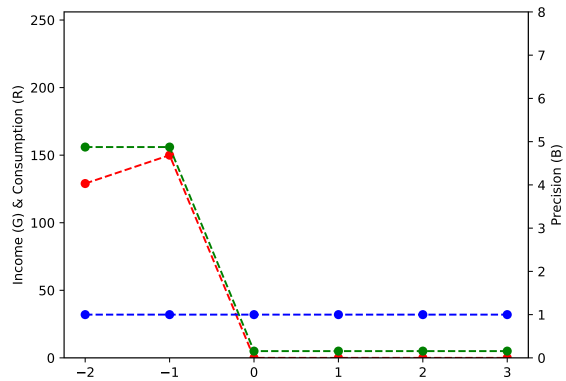

(d) Negative shocks - Sbj 19

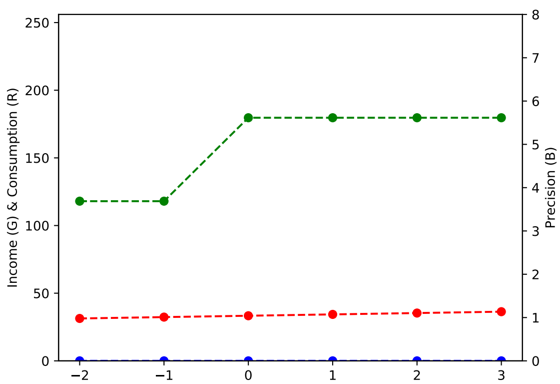

(e) Positive shocks - Sbj 20

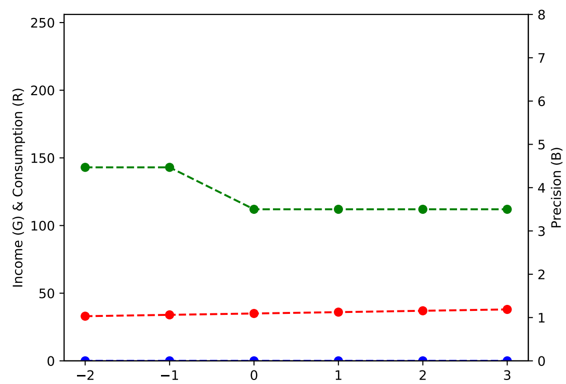

(f) Negative shocks - Sbj 20

Figure A4: INDIVIDUAL RESPONSES TO INCOME SHOCKS 4 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

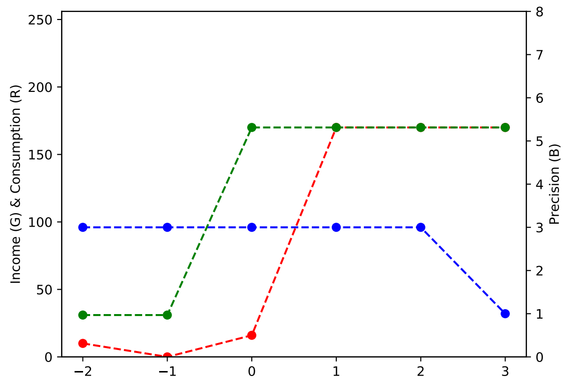

(a) Positive shocks - Sbj 21

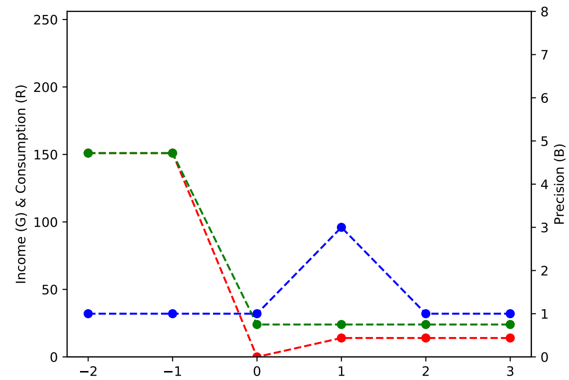

(b) Negative shocks - Sbj 21

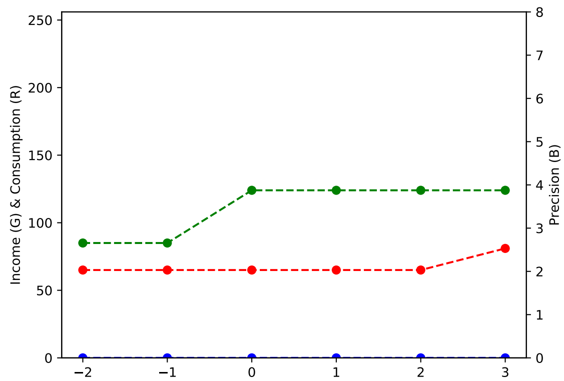

(c) Positive shocks - Sbj 24

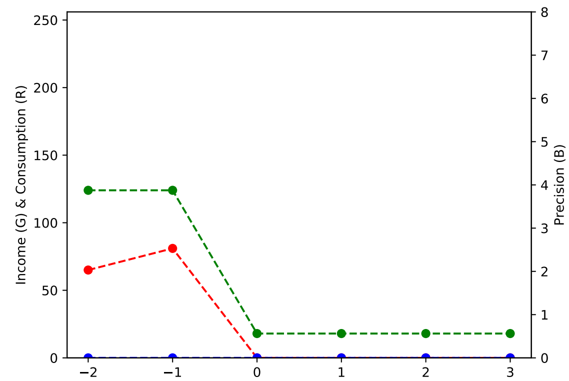

(d) Negative shocks - Sbj 24

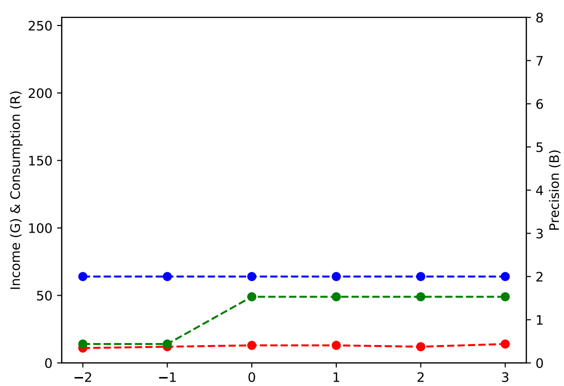

(e) Positive shocks - Sbj 25

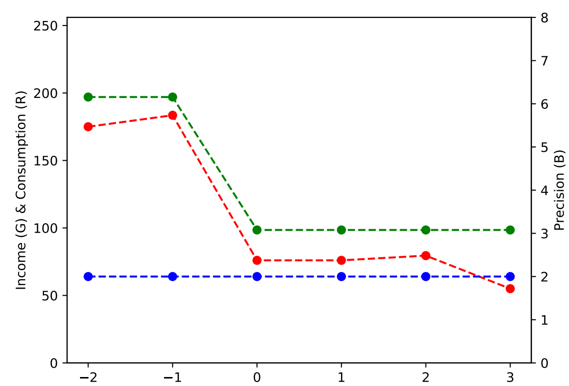

(f) Negative shocks - Sbj 25

Figure A5: INDIVIDUAL RESPONSES TO INCOME SHOCKS 5 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

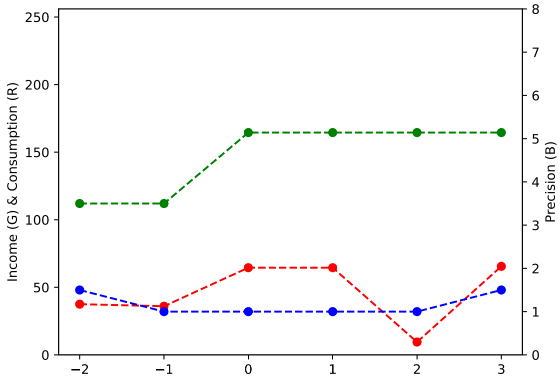

(a) Positive shocks - Sbj 26

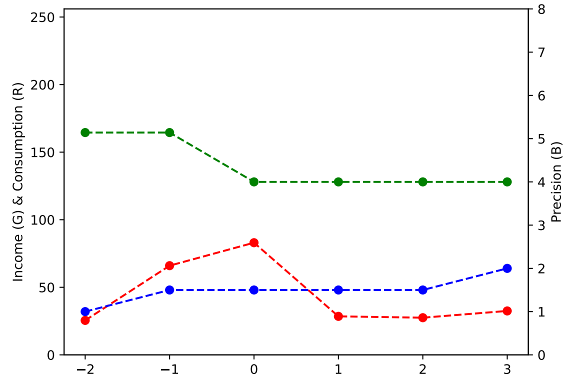

(b) Negative shocks - Sbj 26

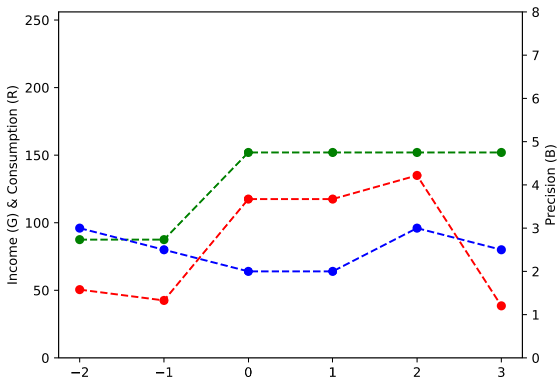

(c) Positive shocks - Sbj 28

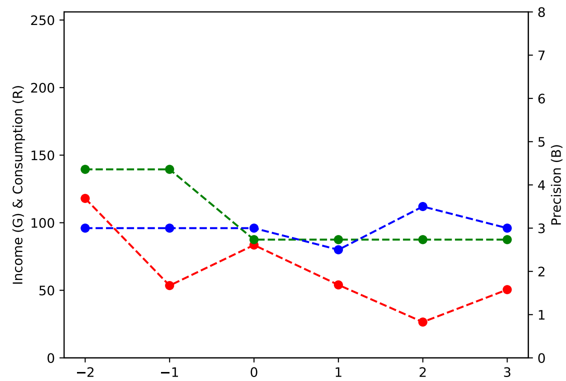

(d) Negative shocks - Sbj 28

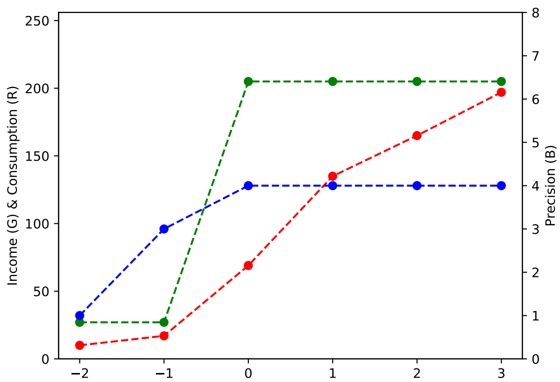

(e) Positive shocks - Sbj 29

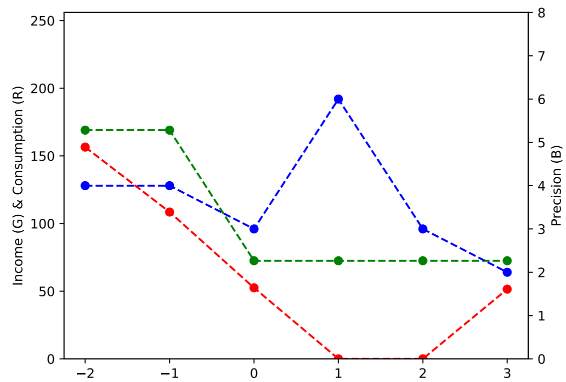

(f) Negative shocks - Sbj 29

Figure A6: INDIVIDUAL RESPONSES TO INCOME SHOCKS 6 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

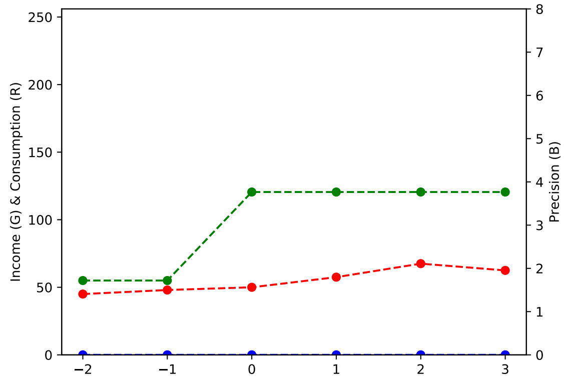

(a) Positive shocks - Sbj 34

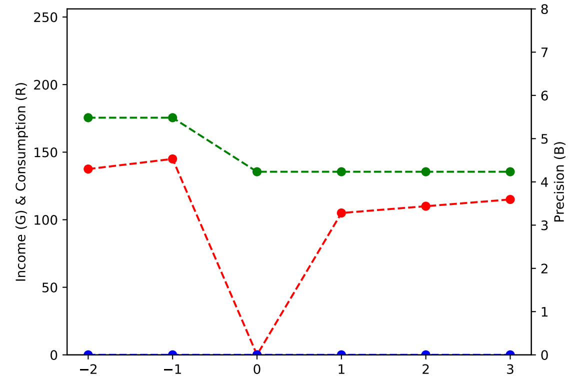

(b) Negative shocks - Sbj 34

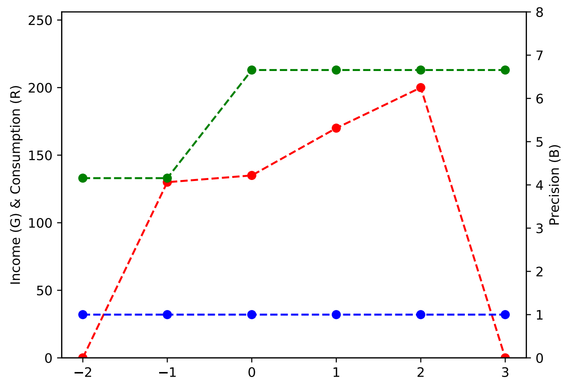

(c) Positive shocks - Sbj 35

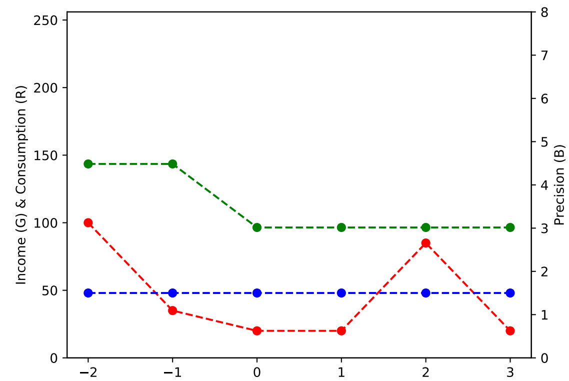

(d) Negative shocks - Sbj 35

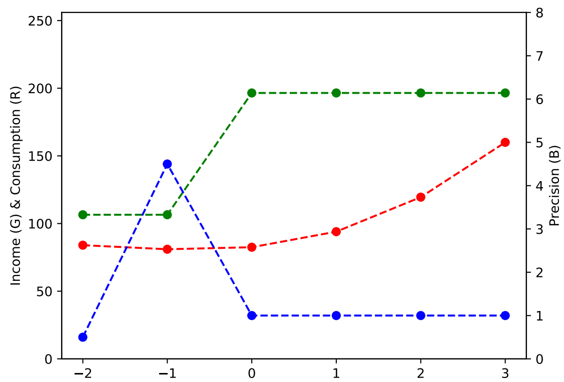

(e) Positive shocks - Sbj 36

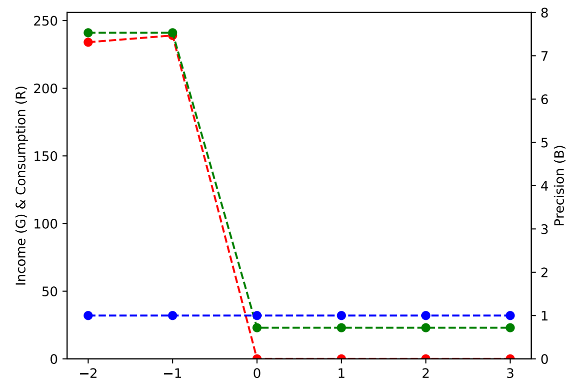

(f) Negative shocks - Sbj 36

Figure A7: INDIVIDUAL RESPONSES TO INCOME SHOCKS 7 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

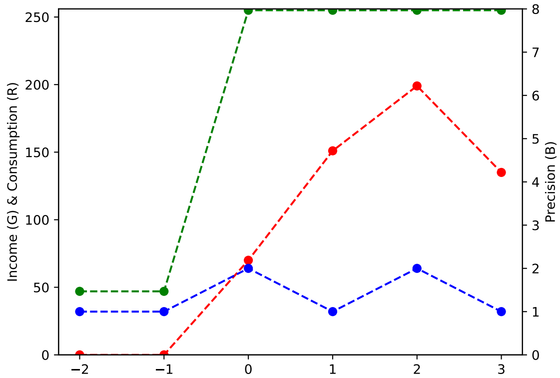

(a) Positive shocks - Sbj 38

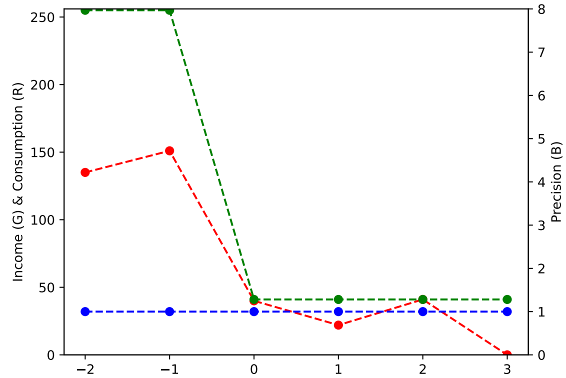

(b) Negative shocks - Sbj 38

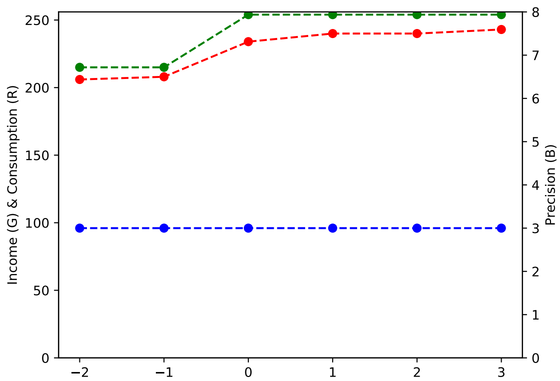

(c) Positive shocks - Sbj 39

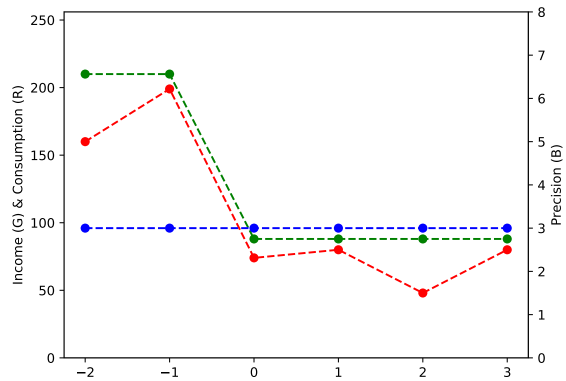

(d) Negative shocks - Sbj 39

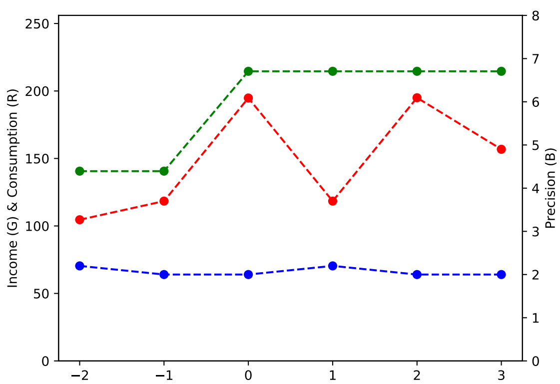

(e) Positive shocks - Sbj 40

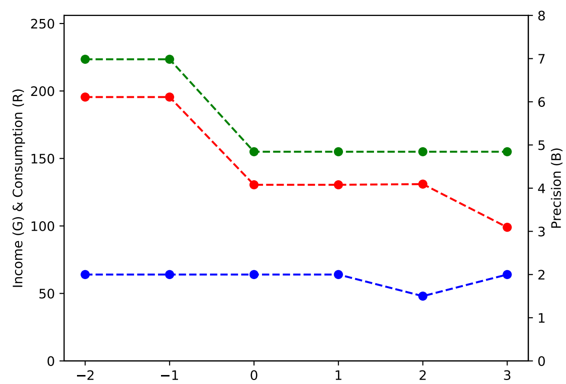

(f) Negative shocks - Sbj 40

Figure A8: INDIVIDUAL RESPONSES TO INCOME SHOCKS 8 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.

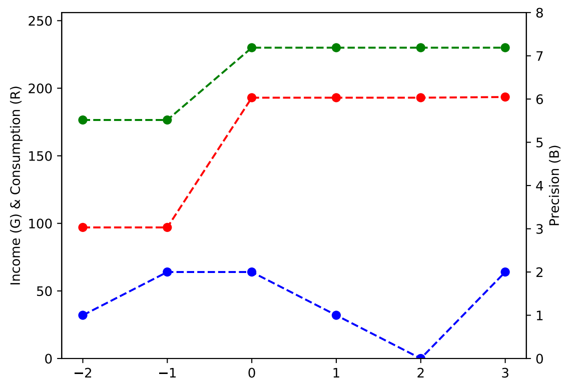

(a) Positive shocks - Sbj 42

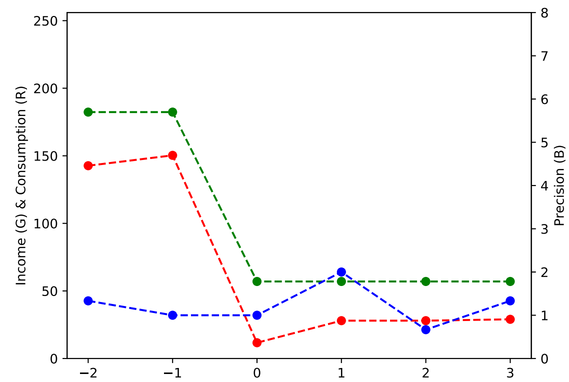

(b) Negative shocks - Sbj 42

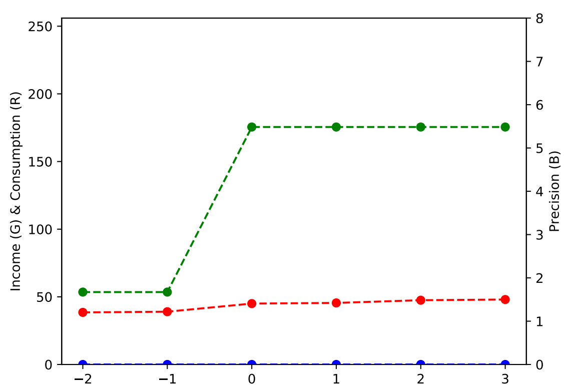

(c) Positive shocks - Sbj 60

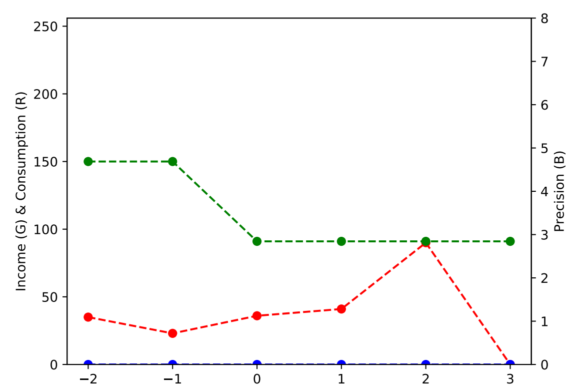

(d) Negative shocks - Sbj 60

Figure A9: INDIVIDUAL RESPONSES TO INCOME SHOCKS 9 - Average income, consumption, and precision in the neighborhood of an income shock, centered at time 0. Two periods before the shock and the three periods after the shock are showed. All subjects in Treatment 3.