INNOCENZO GASPARINI INSTITUTE FOR ECONOMIC RESEARCH

Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

The Weight of Personal Experience: an Experimental Measurement

Joshua Miller and Zacharias Maniadis

Working Paper n. 452

This Version: August 31, 2012

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

The Weight of Personal Experience: an Experimental Measurement

Joshua Miller[a] and Zacharias Maniadis[b] *

August 31, 2012

Abstract

We present an experiment to address the question of whether a piece of information is more influential if it comes from experience, rather than from another source. We employ a novel experimental design which controls for the value of information and other potentially important confounding factors present in related studies. Overall, our results show that an event that is personally experienced has a stronger influence on subsequent behavior than an observed event with equally valuable information content. Importantly, in early rounds when information is more valuable from a rational viewpoint, this overweighting of personal experience is not statistically significant.

JEL Classification Numbers: C90; C91;
Keywords: Experiments; Learning; Observation; Reinforcement Learning; Belief-Based Learning

---

*a: Department of Decision Sciences, Università Bocconi, b: Department of Decision Sciences and Dondena Center, Università Bocconi. This research was supported by Miller's grant from the IGIER research institute. We are grateful to María José Aragón for her Spanish translations with an experimenter's eye and Giovanni Ponti and the LATEX experimental laboratory at the University of Alicante for their hospitality. We also are grateful to Adam Sanjurjo as well as the helpful Alicante graduate students who assisted with the experiments. We are also indebted to Graham Loomes, John Hey, Faruk Gul, Bart Lipman, Gary Charness, and Larry Samuelson for their time, encouragement and helpful comments. We would like to thank participants at the Bocconi University experimental group meetings. We would like to thank Marco Casari, Giulio Zanella and seminar participants at the University of Bologna. We would like to thank participants at the 2011 BEELab conference in Florence, the 2011 IMEBE conference in Barcelona, the 2011 M-BEES Symposium in Maastricht, and the Max Planck ESI Spring Workshop in Jena. Finally, we are grateful to Dimitris Christelis, Botund Kőszegi, Andrew F. Newman for their helpful comments as well as participants at the 2011 Csef-Igier Symposium in Capri. Corresponding author: Joshua Miller. Address: Department Of Decision Sciences, Università Bocconi, Via Roentgen 1, Milano, 20136, Italy, telephone: +39 02 5836.3411, e-mail: *joshua.miller@unibocconi.it*.

"There is nothing like first hand evidence. As a matter of fact, my mind is entirely made up upon the case, but still we may as well learn all that is to be learned." (Sherlock Holmes)

1 Introduction

A basic principle of rational learning is for people to fully use available information. This requires that, all else equal, the value placed on information should not depend on its source. While this is compelling as a normative benchmark, in many settings people appear to under-value information that is not derived from personal experience. People routinely ignore expert advice, such as doctor recommendations, and fail to be sufficiently influenced by information campaigns that are supported by solid statistical evidence, such as efforts to increase seat belt use (Sjőberg 2003; Robertson et al. 1974). An important question in these situations is whether people are truly under-valuing information and doing so because it is not coming from personal experience, or if there are more mundane explanations.

A failure to be influenced by valuable information may simply reflect a rational consideration of the costs and benefits of using the information to alter one's behavior.[1] More importantly, even when people are made worse off by not altering their behavior, the proposition that this is driven by the absence of personal experience is unlikely to be shown using observational data alone. Alternative explanations for under-valuing information, such as limited attention, forgetting, or suspicion of a particular information source, typically cannot be ruled out in natural settings. Moreover, without observing what an agent would have done had information of *equivalent value* been the product of personal experience, one cannot know if and to what degree a piece of information is under-valued because it is not a product of personal experience.

We introduce an experimental environment that is tractable enough to identify whether experimental subjects overweight personal experience relative to equally valuable information

[1] In many cases we may also be mistaken in assuming that information from an outside source is objectively valuable for someone. Second and third hand information is often simplified and does not speak to particular circumstances. For example, receiving information that a large percentage of car accidents are caused by cell phone use may only reveal the consequences of the worst cell phone user practices and may not represent the hazards faced by a more careful user. Note that people may also appeal to these reasons for the purpose of rationalization, i.e. they may be driven by motivated reasoning (Kunda 1990).

sources, and to control for the types of alternative explanations typically encountered in natural settings and in previous literature. In the experimental design, an informative event was operationalized as the draw of a colored ball with replacement from an urn whose contents were unknown, but fixed. A personally experienced event involved having the subjects' payoffs depend on their guess of the drawn ball's color. Events that were not personally experienced, but whose information content was equally valuable, involved each subject simultaneously observing the draws, but not the choices, of two other subjects whose urns had identical compositions to their own. In each period of the experiment the subjects were paid based on correctly guessing the color of the ball drawn from their own urn, before simultaneously witnessing the draws from all urns. In addition, the subjects were required to sort the balls according to color, a novel control designed to assure that subjects attended to all the colors drawn.[2]

The experimental task was designed so that the data analysis would be amenable to an empirical strategy that had been shown, via simulation, to identify the over-weighting of personal experience and, additionally, was robust to various data generating processes. The analysis involved estimating a parameter of personal experience in a discrete choice model whose latent index depended only on the subject's beliefs regarding the frequency of drawn colors, in the spirit of stochastic fictitious play (Fudenberg and Levine 1995). This *weight of experience* parameter measured if subjects behaved *as if* they counted a color drawn from their own urn as more frequent than the same color drawn from another urn.

There was good reason to expect the impact of personal experience with payoffs to lead subjects to place more weight on draws from their own urn in the experiment. In both the psychology and neuroscience literature learning is viewed primarily as product of an individual's interaction with the environment, via their experience with punishments and rewards (Schultz 2004). In the context of choice, the actual experience of choice influences subsequent choices via the rewards and punishments associated with chosen actions, while the counterfactual experience of a choice influences subsequent choices via the rewards and punishments associated with actions not chosen (Berridge 2001; Byrne 2002). In the economics literature, the actual and counterfactual components of personal experience have

[2] Additional design elements were incorporated to control for forgetting as well as the ability to count.

been successfully modeled separately and jointly and have been shown to be factors that influence learning in games (Cheung and Friedman 1997; Camerer and Ho 1999; Fudenberg and Levine 1995; Feltovich 2000; Roth and Erev 1995). Therefore, if learning is primarily driven by personally interacting with the environment, we should expect people to place more weight on information derived from events that relate to personal experience—actual or counterfactual—than on equally valuable information that arises from other sources.

In our experiment, overall, subjects did place more weight on information coming from personal experience relative to equally valuable information obtained by observing the experiences of others. More importantly, however, this effect was pronounced and significant only at times when information was least valuable. In earlier periods, while estimates indicate that subjects behaved as if personal experience had nearly twice the influence of observation on subsequent behavior, this premium was not significantly different than one. An explanation for this is that learning has two components: a rational component and an experience-related component which leads to an underweighting of information from sources not related to personal experience. The rational component disengages after beliefs regarding the distributions have presumably stabilized, and the information content of a draw is low. Thus, the experience-based component of learning dominates in later periods, which explains what we see in our data.

The remainder of the paper is as follows. In Section 2 we review the experimental literature and explain the limitations of previous studies. In Section 3 we introduce our experimental design. In Section 4 we present our empirical strategy. In Section 5 we present our results, and Section 6 concludes.

2 Related Literature

There are many experimental studies that have investigated how people treat sources of information that do not involve direct personal experience.[3] Only a few studies have quantitatively compared the relative importance of experience to other information sources such

[3] There are also field studies that consider the same question, but we feel that discussing them would bring us outside the scope of this paper.

as description and observation. More importantly, no prior study has completely and simultaneously controlled for important factors that we control for—such as attention and memory—while holding constant the value of information. We review first the literature involving individual decision making, which our study is an instance of, and then we review the work in economic games.

In the context of individual decision making Merlo and Schotter (2003) compared the performance of subjects who had repeatedly executed a given task to the performance of others who had only observed another subject execute it. They found that observers outperformed doers for this task, indicating that information that arises from personal experience is treated differently than information that arises from observation, suggesting that they may not be weighted equally when both are available. In other individual decision making studies, researchers have compared decision from description, where objective probabilities are provided by the experimenter, with decision from experience, where subjects have experienced a sample from the same probability distribution. Barron and Erev (2003) and Hertwig et al. (2004) find that in decisions from experience subjects weight rare events significantly less than in decisions from description. In these experiments, however, the information content is not kept constant: experience concerns only a small sample, whereas description informs about the entire distribution.

Many related experimental studies have involved economic games. Duffy and Feltovich (1999) found that the importance of observed information depends on the game. When each player observed the play of another matched dyad, observation mattered for the ultimatum game, but not for the best-shot game. The authors did not directly compare the importance of experienced and observed information. Armantier (2004) performed an experiment of a first-price common value sealed-bid auction. In that experiment observation affected learning significantly, and observational learning was of comparable magnitude with experiential learning. In this study, however, information from experience was limited to whether one has won the auction or not, while observation concerned others' bids, signals and payoffs.

In an information cascade experiment, Goeree et al. (2007) found that private information had two to four times more influence on behavior than public information had, although information acquisition did not relate to an event that involved payoffs. Simonsohn et al.

(2008) used a finitely repeated Prisoner's Dilemma game without re-matching to investigate whether or not subjects were influenced by observing other matched dyads play. While controlling for attention, they found that information from other dyads did not significantly influence behavior while a subject's own past experience did. While the result suggests learning from observation may be irrelevant when personal experience is available, the game form of the Prisoner's Dilemma involves dominated strategies, which means that both observed and experienced information on past play have no value unless ad hoc assumptions are made with regard to the players' preference profile.

In our view, an individual decision making experiment, not an economic game, is a more suitable environment for addressing our research question, which was at most of ancillary interest in many of the aforementioned studies involving games. For our purposes, the first problem with games is that when a small number of subjects per session is used, or when subjects' play is observed by others, repeated-game effects are relevant. If subjects do not simply learn from the past, but also "invest" in affecting the future, their behavior may allow for multiple interpretations. Second, in repeated games, it is not clear if and how people might find past behavior relevant for future play without knowing their beliefs or the strategy they are implementing. Moreover, the lack of stationarity in the distribution of other players' actions may confound participants' attempts to learn from past behavior.[4] As we shall explain, our experimental design, which does not involve participants observing or being affected by the decisions of others, is not exposed to any of these limitations.

3 The experiment

Six experimental sessions were conducted at the "LATEX" laboratory of the University of Alicante on two consecutive days in October of 2010. In total 99 subjects participated, and all of them were undergraduate students of the university.[5] The experiments were conducted

[4] An additional problem is that if subjects are involved in a game, observing another subject's experience often also involves observing another subject's choice. This makes it unclear what construct is being investigated, the influence of information from observing the experience of others, or the social influence from observing the choices of others.

[5] We immediately discarded the sessions conducted on the first day (with a total of 42 subjects). In these sessions there was an error in the z-Tree code causing the colors displayed on the monitors in the lab to be different than intended, and two colors to appear identical for a substantial number of subjects.

using z-Tree experimental software (Fischbacher 2007). Subjects were recruited via email and posted flyers on campus. Participants earned an average of 20.2 Euros in one hour, including a guaranteed 5-Euro participation fee. The average monetary earnings was high for Spain (four times the minimum wage), and subjects appeared motivated and carefully engaged in the task.

3.1 Design and protocol

At the beginning of the experimental session each subject was assigned randomly and permanently to a fixed group with two other participants. Each of the three persons in the group had an urn containing 30 balls with each ball being of one of 3 possible colors: brown, purple or green. All three urns had identical contents and this did not change, i.e. they had the same number of balls of each color in every period. The actual contents of the urns were unknown to the subjects before the 50-period decision task began. Subjects began the experiment with 8 euros and in each period they gained 1 euro if they correctly guessed the color of the ball drawn from their urn and lost 0.50 euro if their guess did not match the color.

In each period the choice task went as follows: first, each subject chose one of three possible colors and after this each of the identical urns was shaken and a single ball was drawn from each urn. Immediately after the draws, the payoff was displayed on the screen (+1.00 euro for a match and −0.50 euro for a failed match) and then each subject was prompted to sort the three drawn balls into bins of the appropriate color.[6] Finally, subjects observed the balls being automatically replaced to the urns they were drawn from, and then they were shown their current cumulative earnings.[7] Table 1 shows a brief description of the characteristics of the six experimental sessions. The number of subjects, the distribution and the dominant color differed across sessions.[8]

We chose three colors instead of two to prevent subjects from explicitly counting the

[6]Subjects had to sort with overall 90% accuracy or forfeit their winnings. The actual accuracy achieved by subjects was 98%.

[7]If their monetary balance reached zero, subjects could continue to play without compensation. This happened in only one case.

[8]Screenshots are available in Appendix B. A video of the experimental task can be viewed by entering or clicking on the following shortened url link: http://goo.gl/khwe6.

number of times that each color was drawn. This allowed us to assure that behavior could not be attributed to the greater ease of counting draws from one's own urn. In order to control for the possibility that subjects may still preferentially remember the colors from their own urn, we designed an experimental manipulation where half of the subjects in each session were randomly assigned to a condition called "full recall". In this condition, instructions and everything else were the same, except for the fact that subjects in the "full recall" treatment were presented with an additional screen before they began. This screen informed them that in each round they would see the number of times that each of the three colors has been drawn in past rounds. In particular, one of the frequencies they observed corresponded to past draws from their own urn, and the other frequency pertained to past draws from the two other urns together. The records with the number of past draws were continually presented at the lower part of the experimental screen (see Figure 6 in Appendix A).

Session	**S1**	**S0**	**Actual Contents**	**Dominant color**
1	11	10	13/30, 9/30, 8/30	Brown
2	11	10	13/30, 9/30, 8/30	Brown
3	10	11	13/30, 9/30, 8/30	Brown
4	6	6	14/30, 9/30, 7/30	Purple
5	6	6	14/30, 9/30, 7/30	Green
6	6	6	14/30, 9/30, 7/30	Green

Table 1: *Characteristics of each experimental session.*

Note: S0 and S1 denote the number of subjects who were assigned to the low-recall and full-recall treatments, respectively.

The protocol of the experimental sessions was conducted so that subjects were introduced to the task in an accessible and simple way, and that their understanding was assured. In each experimental session, subjects were seated in front of a computer terminal and were presented with the instructions on their screens. The instructions were read to them orally by the same person in each session to ensure uniformity and common knowledge. Midway through the instructions subjects were asked to look at an overhead projector and watch a one-minute video of someone performing the task, in order to make them more familiar with the task and the interface.[9] After introducing the rules of the experiment, we presented a few slides, which provided examples of the three urns and their contents. Our objective was to make it clear that the three distributions were the same, and also fixed across rounds. After this, subjects watched once more the video with the presentation of the sample rounds. This time, the experimenter explained in detail what subjects saw on the screen. After the instruction phase, they answered a short quiz with 3 questions, which was aimed at ensuring that they understood the basic structure of the decision task. Feedback with an explanation of the correct answers to the questions was given when subjects answered the quiz incorrectly.

3.2 Discussion of design

Our experimental design achieves several objectives. We keep the value of information constant across information sources, and at the same time we address the aforementioned problems of previous studies. Using our decision task, we resolve the problem of repeated-game effects, namely the incentive to act in order to affect the future beliefs of other subjects. Furthermore, we make the unknown distribution stationary, and make this known to subjects. In such a setting, it is in the interest of agents to attend to the information and learn, as past observations are relevant for predicting the future. In addition to addressing these limitations, our experiment controls for two important factors: attention and memory. It is possible that subjects simply pay little attention to information that does not have direct consequences for them. The novel sorting task, which was described above, ensures that subjects pay similar attention to the color of each of the three urns. Finally, the full recall condition examines whether differences are driven by the possibility that people tend to forget observed draws easier than experienced draws.[10]

[9]The task displayed different ball colors than the actual experiment.

[10]Additionally, learning about the hedonic implications of a given outcome is generally a non-trivial factor in economic games, since players might not know how it feels to be treated in a certain way in a game. Note that only experience can provide this type of learning. For our environment, the affective information of experience should be less important, since Nature, rather than a human being, is the opponent. This makes us more justified in claiming that the "same" piece of information is being observed and experienced. Of course, even in our environment there may be some initial learning with regard to the hedonic value of the payoffs, which can only occur with personal experience.

4 Model, Simulations and Empirical Strategy

Our multinomial logit model-based empirical strategy was customized for the experiment and tested via simulation before the design was finalized. This approach contributed to the validity and credibility of our results in two ways: (1) the success of the statistical model in identifying the true parameters of the underlying simulated data-generating process lent validity to the model itself, and (2) a prior commitment to our empirical strategy imposed discipline on how we analyzed our data, lending credibility to our conclusions.

Let N_B^{own}, N_P^{own} and N_G^{own} be the total number of times that the colors *brown*, *purple* and *green* respectively, have been drawn from a subject's own urn (experienced information). Let N_B^{other}, N_P^{other} and N_G^{other} be the total number of times that the colors *brown*, *purple* and *green* respectively, have been drawn from the other two urns (observed information).[11] A rational observer of the draws, who has a uniform Bayesian prior, would keep track of the total number of balls drawn of each color and choose the color most frequently observed. This is the rational benchmark.

This decision maker can be modeled as forming an index for each color I_B, I_P, and I_G where each index is the total number of draws of that color, e.g. for *brown* $I_B = N_B^{own} + N_B^{other}$. The rational decision-maker then chooses the color with the highest index and randomizes for ties, e.g. choosing *brown* if $I_B > \max\{I_P, I_G\}$. An agent that overweights information derived from personal experience may instead form an index $I_B = \theta N_B^{own} + N_B^{other}$, where the θ can be thought of as the weight of experienced information (so that there is a biased counting procedure). If $\theta > 1$, experienced information is overweighted, or, equivalently, observed information is discounted. If $\theta = 1$, this becomes the same index as the rational benchmark. If $\theta < 1$, the agent discounts information from experience. The index described in the paragraph above is the *one-color* latent index.

The one-color latent index specification was for illustrative purposes. In our estimations, we also allow for the possibility that the latent index for a particular color is influenced by experience and observation of other colors, as well as the possibility that colors are weighted

[11] Note that, for ease of exposition, we suppress the index i for denoting a given individual.

differently.[12] In particular, the index for color $c \in \{B, P, G\}$ is the following:

$$I_c = \beta_B^c(\theta N_B^{own} + N_B^{other}) + \beta_P^c(\theta N_P^{own} + N_P^{other}) + \beta_G^c(\theta N_G^{own} + N_G^{other}) \tag{1}$$

This is the *all-colors* latent index.

We model agents as having random utility for each color. This utility consists of their potentially biased latent index plus a random component. The utility of color $c \in \{B, P, G\}$ is given by $U_c := I_c + \epsilon$. If U_c is thought of as the true utility of color c at the time of decision, ϵ can be thought of as the random variation in taste. If I_c is the true utility of color c, then ϵ can be thought as an error term that captures unmodeled heterogeneity. Agents choose color c when it has the highest utility. Due to the error term, the choice will be probabilistic, so for example, $\mathbb{P}(\text{Choosing } \textit{brown}\ ) = \mathbb{P}(U_B \geq \max\{U_P, U_G\})$. We assume that the errors are independently and identically distributed according to the type-I extreme value distribution, therefore the choice probabilities will be multinomial logit with an index that is linear in the cumulative color counts and θ.[13] In particular, the choice probabilities are:

$$\mathbb{P}(\textit{color } c) = \frac{e^{I_c}}{\sum_{k \in \{B,G,P\}} e^{I_k}} \qquad \text{for } c \in \{B, P, G\} \tag{2}$$

This discrete choice model is a belief-based learning model, a version of stochastic fictitious play in our environment (Fudenberg and Levine 1995).[14]

Our primary goal is to estimate the hidden weighting parameter θ. To this end, prior to finalizing the design of the experiment we simulated an experimental data set consisting of

[12] Both models performed equally well in identifying theta in the simulations.

[13] Independence of Irrelevant Alternatives is implicitly assumed here with the i.i.d. assumption in errors. This is reasonable as all the colors are perfect substitutes.

[14] For measurement purposes the model assumes a (potentially) experience-biased counting process that influences choices, and thus appears to make separate assumptions of how players update their beliefs and how they use their beliefs. Nevertheless, the one-color latent index model can be directly related to a special form of fictitious play in games, called κ-exponential fictitious play, and described in Fudenberg and Levine (1995). The relationship is as follows: (1) in our model fictitious play is extended from actual and counterfactual choice to weighted hypothetical choice among the observed draws of others, (2) in our experimental environment the discrete utility comparison of the κ-exponential fictitious play model—namely the difference between the historical average utility of always choosing color a and the historical average utility of always choosing color b—becomes equivalent to a simpler comparison. In particular, this comparison concerns the difference between the (experience-biased) frequency count of the number of times color a was drawn and the (experience-biased) frequency count of the number of times color b was drawn. (This equivalence holds since payoffs for matching or failing to match the drawn color do not vary over colors and periods).

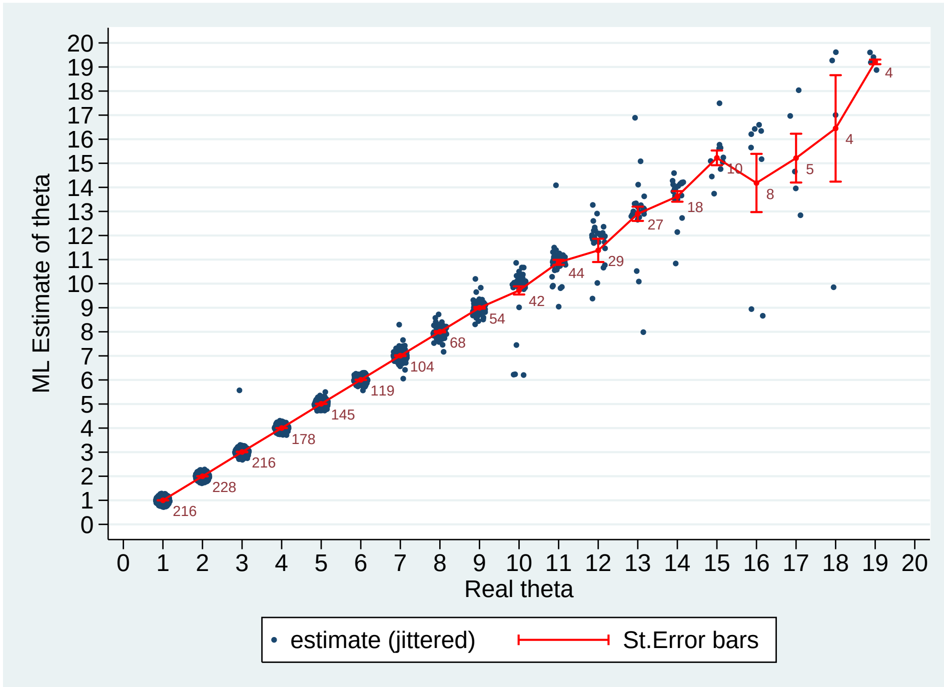

Figure 1: *Simulation and Estimation of θ (numbers indicate the frequency of simulations where the MLE converged within* 50 *iterations).*

agents who chose according to a biased counting procedure.[15] The data-generating processes consisted of random draws from the urns, and simulated subjects choosing deterministically the most frequent color according to their biased counting procedure.[16] In the simulations the hidden θ had a range $\{0, 0.5, 1, 1.5, \ldots, 20\} \cup \{\infty\}$ and we allowed for heterogeneous θs in the population, with up to four different uniformly distributed values. For each parameter value (or population mean) we simulated data and estimated θ for 250 experimental runs, each run consisting of 5 sessions of 60 periods with 20 subjects per session.[17] Our discrete choice model reliably estimated the true θ of the data-generating process. With 250 estimations, both the mean estimate of θ and the error bars were to within 1% of the the true value of θ

[15] While the properties of the multinomial logit model are well known, the simulations were necessary as there are identification issues related to the range of parameter values we expected to encounter as well as the number of periods and variety of urn configurations present in our particular data set.

[16] We also ran simulations where in the data generating process subjects chose stochastically, either choosing in proportion to the biased relative frequency of a color's count, or according to multinomial logit transformations of the biased frequency of a color's count. The results were the same except for improved identification for higher θ.

[17] The urn contents (distribution) were $[14, 10, 6]$. As in the eventual experiment, in each session, the draws were the same across all subjects who were assigned to the same urn (a given subject was assigned either to the left, middle or right urn).

when all subjects had the same $\theta < 8$, as can be seen in Figure 1.[18]

5 Experimental Results

As can be seen in Figure 2, in aggregate, subjects do not generally select the most common color, and by around round 15 (45 total observations from all urns) subjects are choosing colors according to their relative frequency.[19] After round 30, subjects appear to be moving towards choosing the best color, and in the final 5 periods subjects converge to choosing that color with 95% frequency.[20] In the following section these final 5 periods will be removed from analysis.

5.1 Estimation

Using the all-colors latent index discrete choice model discussed in Section 4 we estimated the own-information weighting parameter θ and the coefficients on the total cumulative counts of each color by maximum likelihood. With *brown* as our reference category there are six coefficients to estimate, three representing the impact of each color on the log-odds of the probability of choosing *purple* over *brown*, and three representing the impact of each color on the log-odds of the probability of choosing *green* over *brown*. The overweighting hypothesis predicts that we will find θ to be significantly larger than 1. We anticipated that the treatment effect on θ would lower its value in the full-recall treatment, as a possible memory bias on experienced draws should play a smaller role.

Table 2 presents the estimation results.[21] The first column summarizes the estimation with the interaction of the treatment dummy *full-recall* with θ. Surprisingly, as can be seen,

[18] For large finite values of θ ($9 < \theta < \infty$), the MLE routine rarely converged. This should be expected. Since the distribution of colors was set to be roughly 50%, 30%, 20%, this made the types of observations that could identify and discriminate such large values of θ nearly impossible. Despite this fact, the simulations converged reliably for a data generating process with $\theta = \infty$: in a model with the latent index modified to be $I_c = N_c^{own} + \alpha * N_c^{other}$, we estimated $\alpha \approx 0$ on average (implying $\theta = 1/\alpha \approx \infty$) with similarly tight standard errors, as found in Figure 1.

[19] This finding is consistent with the "probability matching" literature (see Vulkan (2000)).

[20] The behavior after period 45 appears consistent with the results of Gneezy and Potters (1997), who find that as evaluation periods become longer, subjects become less risk-averse, which in this case may translate to less diversification in choices when the end is near, i.e. subjects become less averse to choosing the same color repeatedly.

[21] The results were identical for the one-color latent index model. See Appendix C, Table 4.

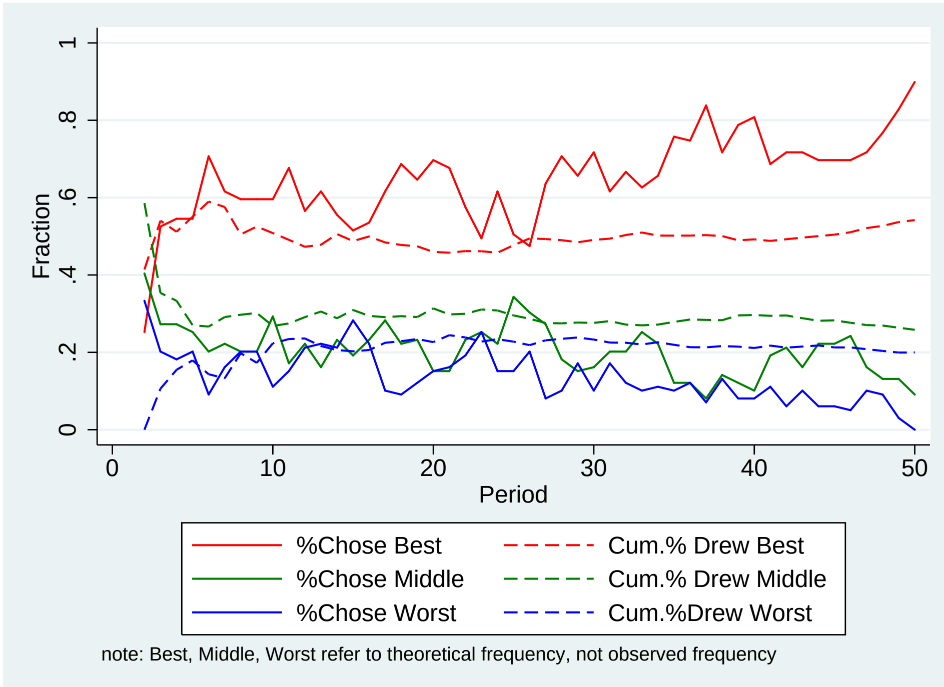

Figure 2: *Draws and Choices for* 50 *periods for the theoretically Best, Middle, and Worst color.*

providing information on all the previous draws had no significant effect on the weighting parameter.[22] Since we could not detect a difference between treatments, we pooled the data from both treatments together. The second column is the estimation of the full model in this case, and as can be seen, we estimate $\hat{\theta} = 3.483$. This means that for subjects who behave as if they choose (with error) the color with the highest relative frequency in the past, draws from their own urn are weighted over three times as much as draws from other urns. The estimate for θ is significantly greater than 1 with $p = 0.019$.

While the coefficients on the cumulative counts of each color were not of particular interest in the estimation, their signs are sensible. For the latent index related to choosing *purple*, the three coefficients represent the change in the log-odds of choosing *purple* over *brown* (the base color), when the cumulative count of a given color increases. For the total cumulative count of *brown* there is a significant negative coefficient (-0.046) while for *purple* there is a significant positive coefficient (0.036), both of which are sensible. There is no a priori reason to believe that the coefficient for *green* should have a particular sign, as that

[22] Although with limited memory capacity it would be irrational to ignore the table presenting the records of previous draws, there was no experimental task to guarantee that subjects attended to the table.

Table 2: *Biased-Counting Multinomial Logit Model*

		Pooled			
	All Periods	All Periods	First 15	Middle 15	Last 15
Counting Weight (θ)					
constant (θ_0)	3.350***	3.483***	1.921***	1.932***	5.102***
	(1.009)	(0.989)	(0.666)	(0.664)	(1.690)
Treatment (θ_1)	0.377				
	(0.803)				
Purple Index					
Brown(β_B^P)	−0.046***	−0.046***	−0.106***	−0.068***	−0.032***
	(0.009)	(0.009)	(0.021)	(0.012)	(0.008)
Purple(β_P^P)	0.036***	0.036***	0.090***	0.047***	0.029***
	(0.009)	(0.008)	(0.031)	(0.015)	(0.007)
Green(β_G^P)	0.009**	0.009**	0.012	0.016*	0.003
	(0.004)	(0.004)	(0.019)	(0.009)	(0.003)
Green Index					
Brown(β_B^G)	−0.044***	−0.044***	−0.130***	−0.062***	−0.029***
	(0.008)	(0.008)	(0.022)	(0.012)	(0.007)
Purple(β_P^G)	0.004	0.003	0.003	−0.006	0.004*
	(0.004)	(0.004)	(0.022)	(0.011)	(0.002)
Green(β_G^G)	0.039***	0.039***	0.109***	0.067***	0.024***
	(0.008)	(0.008)	(0.028)	(0.013)	(0.006)

Latent Indices: $I_c = \beta_B^c(\theta N_B^{own} + N_B^{other}) + \beta_P^c(\theta N_P^{own} + N_P^{other}) + \beta_G^c(\theta N_G^{own} + N_G^{other})$ for $c \in \{P, G\}$.
Base Outcome: Brown Choice. Standard errors in parentheses. 4812 observations, 99 subjects. Periods 2-45.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

should not influence a choice between *purple* over *brown*.[23] The coefficients corresponding to the cumulative counts influencing the log-odds of choosing *green* over *brown* have a similar sensible pattern in terms of their sign.

We further hypothesized that parameter θ might change over time. To this end, we divided the pooled data into 3 epochs of 15 periods, as a visual inspection of Figure 2 reveals that the cumulative draws do not settle down until period 15 and there appears to be a regime change in choices around period 30.[24] As can be seen in Table 2, the values of θ get progressively larger in each epoch. The final epoch is significantly different than the earlier two, as can been seen in the regressions presented in Appendix C, Table 3, which include the epochs as controls.[25] Note, in particular, that $\hat{\theta} = 1.921$ in the first 15 periods, which is not significantly different from 1. It appears that the overweighting of personal experience is driven primarily by behavior in later rounds.

This result does not depend on how the periods are divided into epochs. An additional analysis was performed which involved dividing the 45 periods into 2 epochs and allowing the cutoff defining the epochs to vary between 15 and 30. We used an epoch dummy control on each coefficient, e.g. for the θ coefficient, $\theta = \theta_0 + \theta_1 * \mathbb{1}_{[Period>T]}$. This led to 16 estimations, and Figure 3 presents the graph of the estimates for the θ coefficients only. In the graph, the horizontal axis corresponds to the cutoff defining the epoch division while the vertical axis corresponds to the values of the estimated θ_0, θ_1 and their 95% confidence intervals.[26] As can be seen, the weighting parameter in the earlier periods (θ_0) is larger than 1 regardless of the cutoff and it is not significantly larger than 1 for any cutoff. Only in the later periods is the weighting parameter ($\theta_0 + \theta_1$) significantly greater than 1 overall. As can be seen in Figure 7, the later periods have more of an influence on the value of θ_1 as its estimate increases with the cutoff.[27]

[23] The marginally significant positive value (0.009) likely reflects how frequently in the data an observation of *green* occurred in periods where *purple* was more frequent than *brown*.

[24] We ignored the final 5 periods due to the previously discussed clear shift in behavior noted in Figure 2. Continuing to pool the data was justified, as in each 15-period epoch there is no significant treatment effect and the estimates for θ do not change markedly.

[25] These results also hold up for the one-color latent index model, as can been seen in Appendix C, Table 5.

[26] An analogous graph of the one-color discrete choice model is presented in Appendix C, Figure 7.

[27] The influence appears to be non-linear. When a linear control for each period is used for all coefficients, there is no significant period effect.

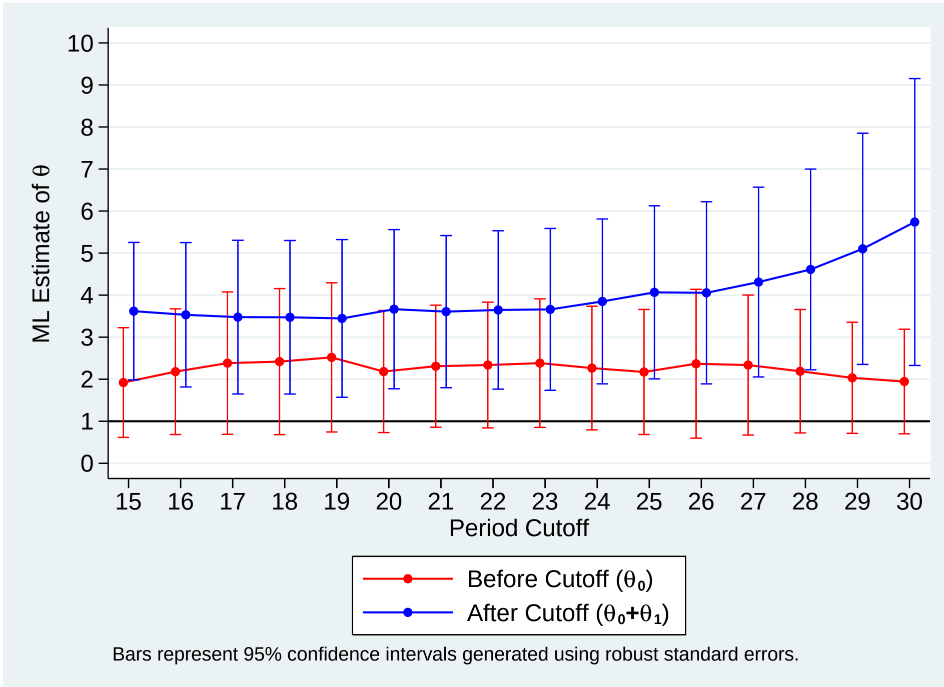

Figure 3: *Estimates for θ with various 2-epoch partitions of periods $2-45$. All-colors discrete choice model: $I_c = \beta_B^c(\theta N_B^{own} + N_B^{other}) + \beta_P^c(\theta N_P^{own} + N_P^{other}) + \beta_G^c(\theta N_G^{own} + N_G^{other})$ for $c \in \{P, G\}$. Each coefficient is of the form $\alpha = \alpha_0 + \alpha_1 * \mathbb{1}_{[Period>T]}$ where T is the cutoff that defines the partition.*

A plausible explanation for these estimates is that subjects incorporated all the information available in the early periods, a time when they were minimally informed. After participating in a sufficient number of periods, the total number of observations, both from one's own urn and that of others, was large. At this point, since the dominant color was represented $40\% - 50\%$ of the time in the true distribution, beliefs may have been nearly stationary. This behavior is consistent with an unbiased balanced weighting of experienced and observed information in the early rounds, i.e. a rational belief-based learning process. In later rounds, the fact that θ is significantly greater than 1 may be driven by the influence of realized own payoffs.[28] This type of behavior is more consistent with a reinforcement-based learning process. This may be reasonable as there is not much to learn in later rounds from the perspective of an agent who has rational beliefs, understands the stationary nature of the experimental environment, and has participated in every round.[29]

[28] It has been suggested that this influence leads to probability matching in environments when probabilities are known, as in past experiments (see Vulkan (2000)).

[29] An agent may have rational beliefs but not best-respond to those beliefs, either due to a misunderstanding

6 Conclusion

This paper presents the results of an experiment designed to test whether information is more influential if it is derived from personal experience, rather than from another source. Overall, we found that subjects behaved as if personally experienced events had over three times the influence on subsequent behavior relative to events of equivalent value that were merely observed. This means that events which are personally experienced influenced behavior in ways not attributable to the objective value of the information they revealed, or its ease-of-use. Nevertheless, experimental subjects clearly incorporated valuable information from events that are not related to personal experience. More importantly, in early rounds of the experiment, when information was more valuable, we cannot rule out the possibility that subjects gave this information equal weight. This means that there may be no premium placed on personal experience when information is sufficiently valuable and controlled for.

In this study, the experimental design and analysis have important features not present in previous work. In the design, experienced and observed events were constructed to yield information that is equally valuable, assuring that a rational agent that fully uses available information should weight both sources of information equally. In addition, a novel control in the form of a sorting task was employed to keep the level of attention between the two sources of information constant. Finally, the possibility that differential demands on memory may be driving the result was ruled out with a randomized treatment design. In the analysis, we employed an empirical strategy whose ability to identify the weight placed on experienced information was investigated by numerically simulating experimental data sets prior to the study. The commitment to use the strategy from the simulations on the subsequent experimental data disciplined our analysis while the performance of the estimators in the simulations provided confidence in our measurements.

While the present paper does address the question of how much weight people place on information derived from personal experience relative to other sources, other questions remain. The experimental design operationalizes information that is not derived from personal experience as the first hand observation of the experiences of others. How would these

of chance processes or because of an aversion to always choosing the same color.

results generalize to second-hand and third-hand sources of information? How can the value of information be controlled for in those settings? Moreover, this study does not investigate what mechanisms led to the bias towards personal experience when the value of information was low in later rounds of the experiment. A plausible driver of this behavior is the unconscious reinforcement learning and fictive (counterfactual) learning that relates to personal experience (see, e.g. Camerer and Ho (1999)). An appropriately designed neuroeconomic study may be able to shed light on this issue.

In our view, in terms of policy implications, the early-rounds results presented here contain a positive message: when people are exposed to information that is sufficiently valuable and representative of their own experience, they will weight it appropriately, regardless of its source.

References

Armantier, O. (2004): "Does observation influence learning?" *Games and Economic Behavior*, 46, 221–239.

Barron, G. and I. Erev (2003): "Small feedback-based decisions and their limited correspondence to description-based decisions," *Journal of Behavioral Decision Making*, 16, 215–233.

Berridge, K. C. (2001): "Reward Learning: Reinforcement, Inventives, and Expectations," *The Psychology of Learning and Motivation*, 40, 223–278.

Byrne, R. M. (2002): "Mental models and counterfactual thoughts about what might have been," *Trends in Cognitive Sciences*, 6, 426–431.

Camerer, C. and T.-H. Ho (1999): "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67, 827–874.

Cheung, Y. and D. Friedman (1997): "Individual Leraning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46–76.

Duffy, J. and N. Feltovich (1999): "Does observation of others affect learning in strategic environments? An experimental study," *International Journal of Game Theory*, 28, 131–152.

Feltovich, N. (2000): "Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games," *Econometrica*, 68, 605–641.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10, 171–178.

Fudenberg, D. and D. K. Levine (1995): "Consistency and cautious fictitious play," *Journal of Economic Dynamics and Control*, 19, 1065–1089.

Gneezy, U. and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 112, 631–645.

Goeree, J. K., T. R. Palfrey, B. W. Rogers, and R. D. McKelvey (2007): "Self-Correcting Information Cascades," *Review of Economic Studies*, 74, 733–762.

Hertwig, R., G. Barron, E. U. Weber, and I. Erev (2004): "Decisions from Experience and the Effect of Rare Events in Risky Choice," *Psychological Science*, 15, 534–539.

Kunda, Z. (1990): "The case for motivated reasoning," *Psychological Bulletin*, 108, 480–498.

Merlo, A. and A. Schotter (2003): "Learning by not doing: an experimental investigation of observational learning," *Games and Economic Behavior*, 42, 116–136.

Robertson, L. S., A. B. Kelley, B. O'Neill, C. W. Wixom, R. S. Eiswirth, and W. Haddon (1974): "A controlled study of the effect of television messages on safety belt use." *American Journal of Public Health*, 64, 1071–1080.

Roth, A. and I. Erev (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164–212.

Schultz, W. (2004): "Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology," *Current Opinion in Neurobiology*, 14, 139–147.

Simonsohn, U., N. Karlsson, G. Loewenstein, and D. Ariely (2008): "The tree of experience in the forest of information: Overweighing experienced relative to observed information," *Games and Economic Behavior*, 62, 263–286.

Sjőberg, L. (2003): "Neglecting the Risks: The Irrationality of Health Behavior and the Quest for La Docle Vita," *European Psychologist*, 8, 266–278.

Vulkan, N. (2000): "An Economists' Perspective on Probability Matching," *Journal of Economic Surveys*, 14, 101–118.

A Appendix: Instructions

PRELIMINARIES

Welcome. Please turn off pagers and cell phones now. It is important that you do not talk, or in any other way try to communicate during the study. This is a decision making study.

Please follow the instructions carefully. You may earn a considerable amount of cash, which will be paid to you as you leave.

What to Expect:

- 15 Minutes of Instructions
- 5 Minutes of Quiz
- 20 minutes of Decision Making (50 periods)
- Payment

How you are paid:

- 15 Minutes of Instructions
- In the decision task period, you start with 8 Euros and in each round you either earn 1,00 Euro or lose 0,50 Euro.
- In addition to your earnings from the decision task, you will get a fixed participation fee of 5,00 Euro.
- The possible amount you earn during the whole study ranges from 5,00 Euros to 63,00 Euros.

INSTRUCTIONS

In this initial instruction phase, we will now introduce you to the decision task. After the reading of the instructions, if you have questions please raise your hand (do not yell out) and an assistant will be there to help you individually. Following the instruction phase, the 50 periods (20 min) of decision tasks will begin. After the decision tasks start you will not be able to ask any further questions, so please be sure you understand the rules now.

The Decision problem. *Please watch the following 1 minute video on the screen. We will show it once again after we complete the instructions.*

[**VIDEO SCRIPT-NOT IN WRITTEN INSTRUCTIONS**]

This video presents what you will see in a typical round of the study. We shall explain this video in much more detail later, just watch and notice that there are balls and there are boxes. Now please listen to these instructions carefully and read along.

You will participate in up to 50 periods of a decision task. You will start the first period with 8,00 Euros. You may continue to participate as long as you have a positive amount of money. We describe the basic features of the decision task below.

You will be given one box. Your box will contain 30 balls, and each ball has one of 3 possible colors. There are two other participants in this room today that have a 30-ball box with identical contents as your box, i.e. they have the same number of balls of each color.

Each period will be as follows:

1. You choose one of the three colors.
2. After this, each box is shaken and one ball is randomly drawn from each box. You see the color of each ball.
3. If the color you chose matches the color of the ball drawn from your box, you gain 1,00 Euro, if your color doesn't match that ball you lose 0,50 Euro.
4. Next you use your mouse to drag each of the 3 balls to the square of the same color. If you fail to do this with at least 90% accuracy, you will not receive your winnings from

the task and instead will be paid only your 5,00-Euro show-up fee. (it is very easy to do this accurately)

5. After you do this, each ball is replaced in the box that it came from, so that in each box the number of balls of each color stays fixed.

6. You move to the next period, which will proceed in the same way.

Now please watch the video as we describe it to you. When it is over we will continue on the next page.

[VIDEO SCRIPT-NOT IN WRITTEN INSTRUCTIONS]

We will begin a video showing you 3 example periods of the task. FIRST PERIOD: Your screen first presents three bars, which represent the three different colors. In this example it is orange, grey and purple. When the text reading "Get Ready" disappears you may choose a color. In order to choose a color, you simply click at the appropriate bar. Make sure that you choose quickly, because if you don't choose before time runs out you will lose 0,50 Euros for sure. In this example orange is chosen.

After you choose, you see that your chosen orange bar remains and the others disappear. Next the boxes are shaken and then the balls are drawn. You observe the color of the balls drawn from each box. In this example period the color of the ball drawn from your box is orange and matches the color you chose. You are notified that you win 1,00 Euro. Next, the three balls remain, and you need to place each ball in the square of the same color. You move the balls by simply dragging them with your mouse. You have a limited amount of time for doing this, but it is very easy to do. Finally, you see the balls replaced to the boxes they came from and you move to the next period.

SECOND PERIOD: In this example period you fail to respond in time and lose 0,50 Euro. Please watch. (experimenter stays silent after and waits for the the period to end)

THIRD PERIOD: In this example period your chosen color *purple* does not match the color of the ball drawn from your box, which is grey. In this case you lose 0,50 Euro. Please watch. (experimenter stays silent after and waits for the the period to end)

Any questions? Please raise your hand and someone will come around to answer your question. Please make sure that everything is clear, because there will be no further opportunities for questions

[**NOW THE EXPERIMENTER SHOWS OVERHEAD SLIDE WITH THE BOXES**]

<u>So what is happening in the boxes?</u> *Please look at the overhead projector as we describe it to you:*

1. On the slide there are 3 boxes. In each period, before the balls are drawn from each box, the boxes are shaken to make sure that for each box, all balls have an equal chance of being chosen.

2. Notice there are 30 balls in each box colored either Orange, Pink, or Yellow (these are just example colors)

   - In this example there are 26 Orange balls, 2 Pink balls, and 2 Yellow balls in each box (these are just example numbers). Boxes are always identical for you and the two others both here and in the real decision task, therefore all boxes will have the same number of balls of each color.
   - Each ball has the same chance to be drawn. Therefore, in the above example the chance Orange is drawn is 26 out of 30, the chance that Pink is drawn is 2 out of 30 and the chance that Yellow is drawn is 2 out of 30. This is true for each box. Notice, the contents of the boxes are the same but different colors may be drawn from each box in a given period.

3. At the end of the period drawn balls are replaced. Before the draw of the next period the boxes are "SHAKEN" to make sure that each ball has an equal chance of being

chosen.

4. On the slide representing the next period, you see that the chance Orange is drawn is still 26 out of 30, the chance that Pink is drawn is still 2 out of 30 and the chance that Yellow is drawn is still 2 out of 30.

   Note: The number of balls of each color used in this example was just for illustration. In addition, the 3 colors used in the 50 periods of the actual task today will not be the colors from this example.

5. Please look at the final slide. We remind you here what is true of the actual decision task:

   - All three boxes will have the same contents, but you will not know the actual contents. You will only see draws from your box and the two other boxes in your group.

You may be curious how we make sure the box is well "shaken" so each ball has an equal chance of being chosen? We have a paid subscription to the random number service www.random.org which generates true random numbers. It is a service used in many countries for lotteries, draws and sweepstakes in order to obtain truly fair and unbiased random numbers that can be publicly verified. For each draw from a box, www.random.org generates a random number between 1 and 30 in order to pick one of the 30 balls at random from a given box. Since these draws are made publicly with a timestamp on www.random.org's website you may personally verify the accuracy of your draws. If you would like to do this please inquire with us when we are finished for details.

B Appendix: Video and ScreenShots

A video of the task can be viewed by entering or clicking on the following google-shortened url link: http://goo.gl/khwe6. (A permanent link can be found here : http://youtu.be/QHEncPC-xYk?hd=1)

Below are several screen shots

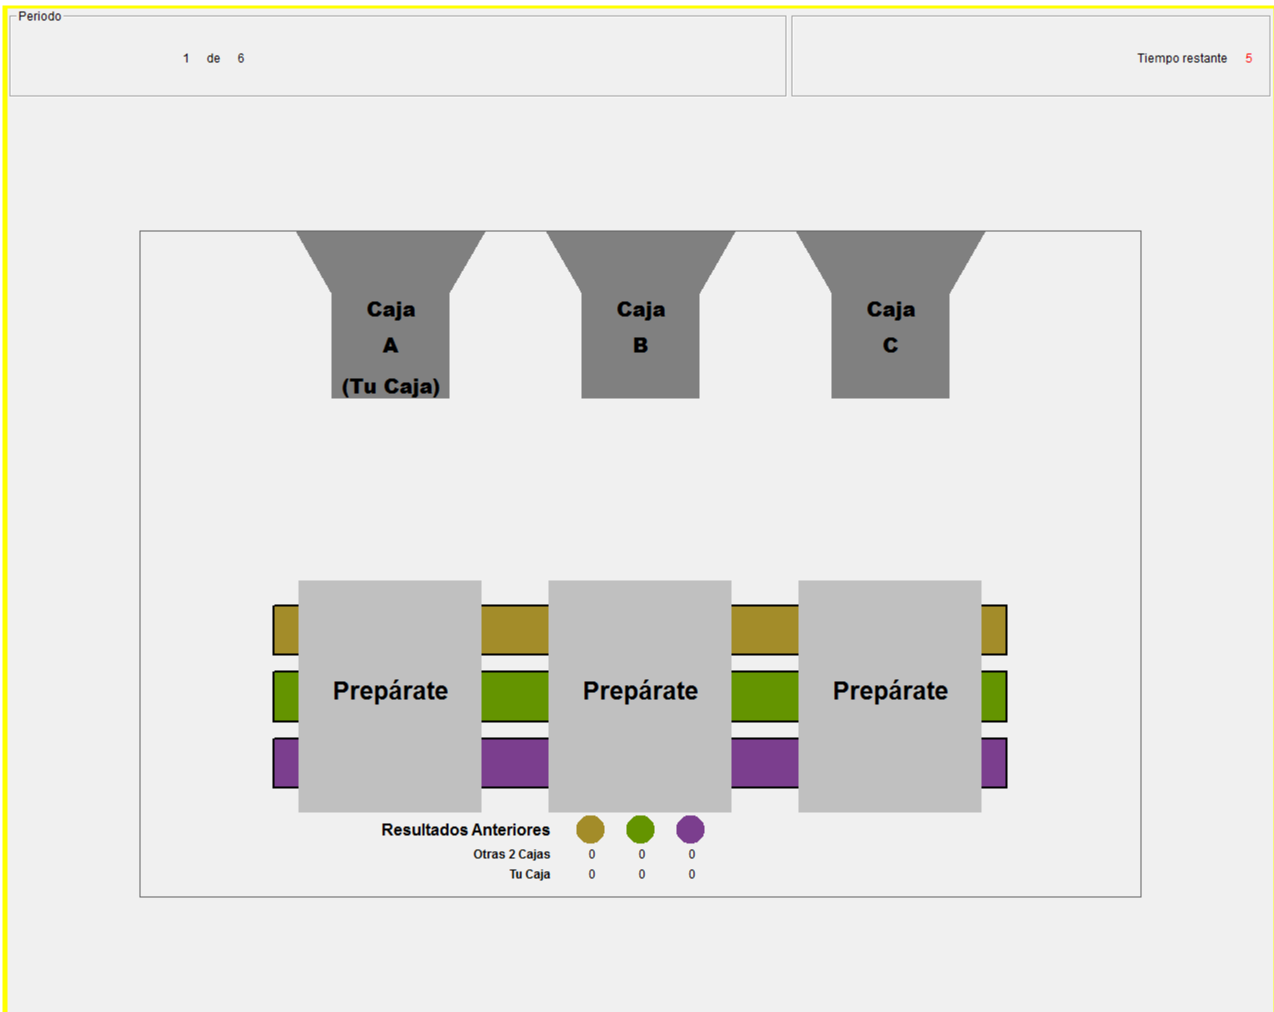

Figure 4: *Experimental screen showing the full-recall treatment, just before subjects are allowed to choose. (In the low-recall treatment the records below were omitted).*

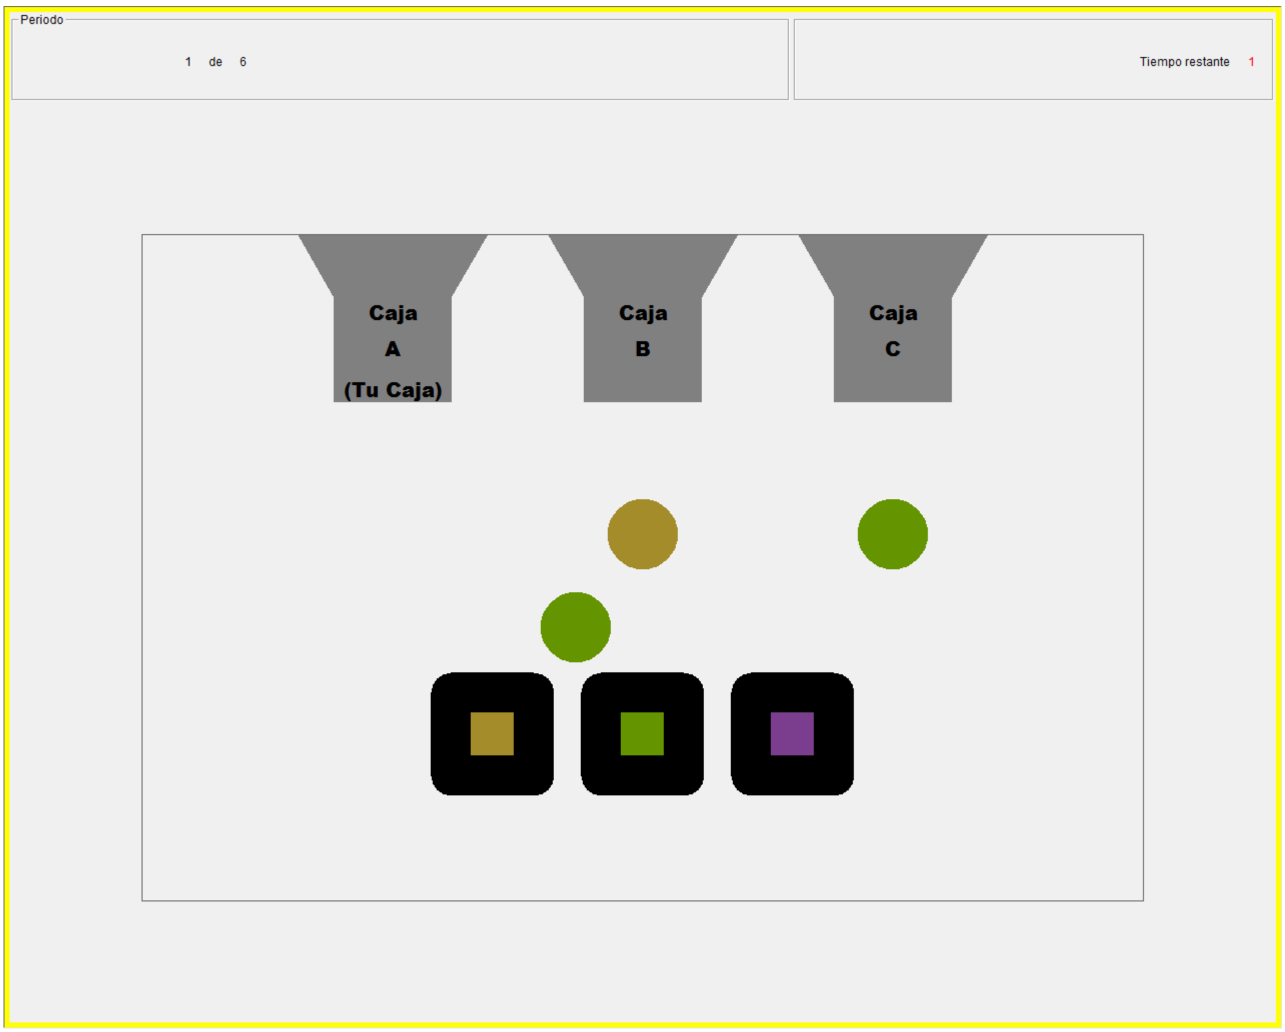

Figure 5: *The Attention Task: Experimental screen showing subjects sorting balls in the low-recall treatment. (Subjects also sorted the balls in the full-recall treatment).*

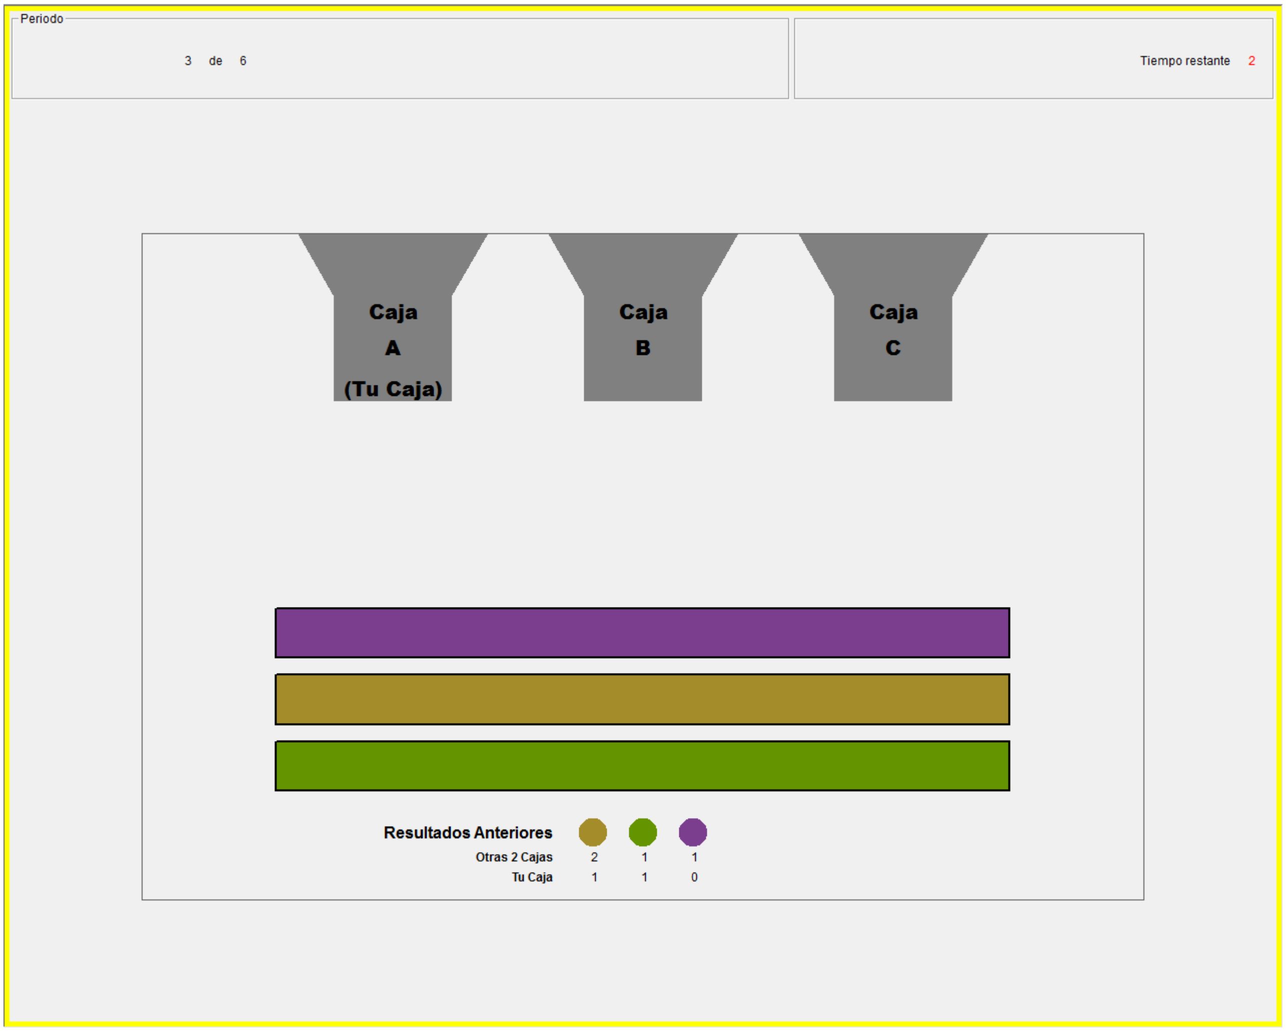

Figure 6: *Experimental screen showing full-recall treatment. Here are the three bars, one of which subjects choose.*

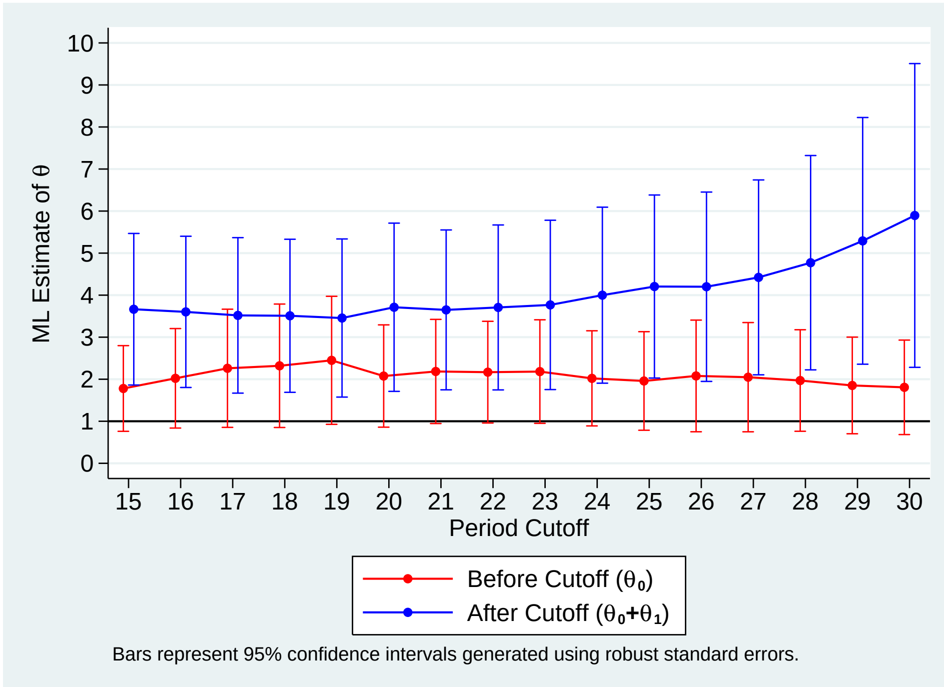

Figure 7: *Estimates for θ with various 2-epoch partitions of periods $2-45$. One-color discrete choice model: $I_c = \beta_c(\theta N_c^{own} + N_c^{other})$ $c \in \{B, P, G\}$. Each coefficient is of the form $\alpha = \alpha_0 + \alpha_1 * \mathbb{1}_{[Period>T]}$ where T is the cutoff that defines the partition.*

C Appendix: Additional Tables & Figures

Table 3: *Biased-Counting Multinomial Logit Model with Epoch Controls*

	All Periods	Time Control	Epoch Controls
Counting Weight (θ)			
constant	3.483***	3.287***	1.921***
Period		−0.013	
Last 15 Periods			3.181**
Second 15 Periods			0.011
Purple Index			
Brown(β_B^P)			
constant	−0.046***	−0.079***	−0.106***
Period		0.001**	
Last 15 Periods			0.074***
Second 15 Periods			0.038**
Purple(β_P^P)			
constant	0.036***	0.055**	0.090***
Period		−0.000	
Last 15 Periods			−0.061**
Second 15 Periods			−0.043*
Green(β_G^P)			
constant	0.009**	0.018	0.012
Period		−0.000	
Last 15 Periods			−0.009
Second 15 Periods			0.004
Green Index			
Brown(β_B^G)			
constant	−0.044***	−0.089***	−0.130***
Period		0.001***	
Last 15 Periods			0.100***
Second 15 Periods			0.067***
Purple(β_P^G)			
constant	0.003	−0.001	0.003
Period		0.000	
Last 15 Periods			0.001
Second 15 Periods			−0.010
Green(β_G^G)			
constant	0.039***	0.086***	0.109***
Period		−0.001***	
Last 15 Periods			−0.085***
Second 15 Periods			−0.042*

All-Colors discrete choice model with latent index:
$I_c = \beta_B^c(\theta N_B^{own} + N_B^{other}) + \beta_P^c(\theta N_P^{own} + N_P^{other}) + \beta_G^c(\theta N_G^{own} + N_G^{other})$ for $c \in \{P, G\}$.
Base Outcome: Brown Choice. Standard errors suppressed due to space restrictions.
4812 observations, 99 subjects. Periods 2-45.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: *Biased-Counting Multinomial Logit Model with 1-color latent index*

	All Periods	All Periods	First 15	Middle 15	Last 15
Counting Weight (θ)					
constant (θ_0)	3.411***	3.453***	1.780***	1.750***	5.291***
	(1.070)	(1.040)	(0.520)	(0.638)	(1.816)
Treatment (θ_1)	0.101				
	(0.806)				
Color Weight					
Brown (β_B)	0.041***	0.041***	0.120***	0.063***	0.027***
	(0.008)	(0.008)	(0.021)	(0.011)	(0.007)
Purple (β_P)	0.038***	0.038***	0.112***	0.057***	0.026***
	(0.008)	(0.008)	(0.025)	(0.011)	(0.007)
Green (β_G)	0.036***	0.036***	0.102***	0.060***	0.024***
	(0.008)	(0.008)	(0.021)	(0.011)	(0.006)

Latent Indices: $I_B = \beta_B(\theta N_B^{own} + N_B^{other})$, $I_P = \beta_P(\theta N_P^{own} + N_P^{other})$, $I_G = \beta_G(\theta N_G^{own} + N_G^{other})$.
Base Outcome: Brown Choice. Standard errors in parentheses. 4812 observations, 99 subjects. Periods 2-45.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: *Biased-Counting Multinomial Logit Model with 1-color latent index with controls*

	All Periods	Time Control	Epoch Controls
Counting Weight (θ)			
constant	3.453***	3.060***	1.780***
	(1.040)	(1.039)	(0.520)
Period		−0.009	
		(0.016)	
Last 15 Periods			3.512**
			(1.597)
Second 15 Periods			−0.030
			(0.599)
Color Weight			
Brown (β_B)			
constant	0.041***	0.082***	0.120***
	(0.008)	(0.018)	(0.021)
Period		−0.001***	
		(0.000)	
Last 15 Periods			−0.093***
			(0.018)
Second 15 Periods			−0.056***
			(0.017)
Purple (β_P)			
constant	0.038***	0.073***	0.112***
	(0.008)	(0.021)	(0.025)
Period		−0.001**	
		(0.000)	
Last 15 Periods			−0.086***
			(0.022)
Second 15 Periods			−0.055***
			(0.020)
Green (β_G)			
constant	0.036***	0.078***	0.102***
	(0.008)	(0.020)	(0.021)
Period		−0.001***	
		(0.000)	
Last 15 Periods			−0.079***
			(0.018)
Second 15 Periods			−0.042**
			(0.017)

One-color discrete choice model with latent indices:
$I_B = \beta_B(\theta N_B^{own} + N_B^{other})$, $I_P = \beta_P(\theta N_P^{own} + N_P^{other})$, $I_G = \beta_G(\theta N_G^{own} + N_G^{other})$.
Base Outcome: Brown Choice. Standard errors in parentheses. 4812 observations, 99 subjects. Periods 2-45.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$