

Are Biased Beliefs Fit to Survive?

An Experimental Test of the Market Selection Hypothesis

Chad Kendall

Ryan Oprea*

February 12, 2017

Abstract

We experimentally study the market selection hypothesis, the classical claim that competitive markets bankrupt traders with biased beliefs, allowing unbiased competitors to survive. Prior theoretical work suggests the hypothesis can fail if biased traders over-invest in the market relative to their less biased competitors. Subjects in our experiment divide wealth between consumption and a pair of securities whose values are linked to a difficult reasoning problem. While most subjects in our main treatment form severely biased beliefs and systematically over-consume, the minority who form unbiased beliefs consume at near-optimal levels – an association that strongly supports the market selection hypothesis.

Keywords: market selection hypothesis, survival of the fittest, efficient markets, Bayesian errors, Monty Hall problem, experimental economics

JEL codes: C9, D03 G1

1 Introduction

We know from decades of economic and psychological research that human decision-makers are prone to a range of inferential biases. However, we also know from this same body of research that humans are heterogeneous: though the average subject forms biased beliefs in a number of settings, most studies report at least *some* subjects forming unbiased beliefs. This heterogeneity is

*Kendall: Marshall School of Business, University of Southern California, Los Angeles, CA, 90089, chadkend@marshall.usc.edu; Oprea: Economics Department, University of California, Santa Barbara, Santa Barbara, CA, 95064, roprea@gmail.com.

crucial to economists because of a long standing hypothesis – sometimes called the “market selection hypothesis” – that biased agents in competitive markets will bankrupt themselves by repeatedly making poor bets relative to their unbiased competitors, leaving only unbiased decision-makers to influence market outcomes in the long run.

This idea – an economic analogue to natural selection – has a long and influential history in economics (stretching back at least to Alchian (1950) and Friedman (1953)), but theoretical research in the past few decades has emphasized that the validity of this hypothesis is ultimately an empirical question. If agents universally save and consume optimally, market selection will indeed tend to hold in competitive markets (e.g. Sandroni (2000); Blume and Easley (2006)), but heterogeneity in consumption/savings behavior (of a sort commonly observed in both laboratory and field studies) can generate environments in which the hypothesis fails (e.g. De Long et al. (1991); Blume and Easley (1992)). In particular, if biased agents systematically hoard wealth relative to unbiased agents, biased agents may survive, driving their unbiased competitors from the market. Wealth dynamics can thus either improve or worsen market outcomes in the long run, depending upon how biased agents consume relative to their unbiased competitors. These effects are difficult to evaluate using evidence from naturally occurring markets, where essential variables like beliefs, consumption behavior, and wealth are typically unobservable, and confounding forces such as learning, self-selection and liquidity concerns abound. For this reason, the question is well suited to laboratory experiments where all of the relevant variables can be cleanly measured and potential confounds carefully controlled.

In this paper, we examine the market selection hypothesis in the laboratory, focusing on the survival of one of the most severe (and persistent) types of biases ever observed in the lab: conditional reasoning failures of the sort arising in settings like the Monty Hall problem, the hot hand fallacy (Miller and Sanjurjo (2015)), and the winner’s curse. Our experiment closely mirrors the setting of Blume and Easley (1992): subjects are endowed with wealth and repeatedly decide how much to take in immediate payouts and how much of the remainder to allocate across a pair of Arrow securities whose values are determined by a Monty Hall-like process. Investment in the Arrow securities determines wealth in future rounds, which in turn allows for future consumption and investment. Importantly, subjects can accumulate wealth relative to the rest of the market not only by allocating wealth across securities in a less biased manner, but also by consuming less wealth (and investing more in the market) than one’s competitors. It is thus possible, as Blume and Easley (1992) emphasize, for subjects with highly biased beliefs to out-survive competitors simply by consuming a sufficiently small amount of wealth each round.

Prior research on consumption/savings behavior has documented a widespread tendency to make sub-optimal decisions in experimental tasks that closely resemble the savings component of our experiment (e.g. Hey and Dardanoni (1988); Ballinger et al. (2003); Carbone and Hey (2004); Ballinger et al. (2011)). Likewise, a number of studies on belief formation have observed a pattern of highly biased choices in reasoning tasks closely related to the investment component of our experiment (e.g. Nalebuff (1987); Friedman (1998); Slembeck and Tyran (2004); Kluger and Wyatt (2004)). Based on this prior research, we have strong reasons to expect both the consumption decisions and the revealed beliefs in our experiment to be, on average, suboptimal (though, in each case, heterogeneous). Market selection, however, fundamentally depends not on the characteristics of the *marginal* distribution of either consumption decisions or beliefs, but rather on characteristics of their *joint* distribution in the population. Our contribution is to examine how failures of standard economic rationality are co-distributed in the subject population, and to use this empirical distribution to form conclusions regarding the effectiveness of market selection for correcting biases like the one we study. If hoarding tends to be higher for biased subjects, market selection can fail or even reverse; if biased subjects instead tend to over-consume relative to unbiased subjects, market selection can occur with greater speed than benchmark models that assume optimal consumption (e.g. Sandroni (2000) and Blume and Easley (2006)) would suggest.

We find strong support for the market selection hypothesis in our data. In our main condition (called CORE), subjects face exogenous prices (perfectly implementing the competitive setting of the model). As expected, we observe pervasive bias in asset allocations, with over 80% of subjects systematically betting on the less valuable state and a small sliver of subjects investing in a Bayesian manner. The average CORE subject also substantially over-consumes (by over 50% at the median) but, crucially, relatively Bayesian CORE subjects tend to consume at a rate much closer to optimal. We show (using theoretical tools from Blume and Easley (1992) and simulations) that the resulting joint distribution between beliefs and consumption rules suggests that biased beliefs in our sample have low survival value (given associated consumption choices) and will tend to be driven out of markets by less biased beliefs. Indeed, the results imply that market selection will operate at a substantially *faster* rate than had subjects uniformly consumed optimally. Counterfactual exercises show that these results are fundamentally driven by the empirical joint distribution between beliefs and consumption in our data: under alternative joint distributions, wealth dynamics would have instead driven *unbiased* subjects systematically from the market, reinforcing rather than eliminating the bias observed in the average subject.

In two additional treatments, we assess the robustness of our main findings. In the endogenous

price treatment (END), we group subjects into small, imperfectly competitive markets and determine prices based on subjects' betting decisions via a market mechanism. This treatment studies how our results extend to *imperfectly competitive* markets. In our LEARN treatment, we study environments in which subjects have sufficient information to form unbiased beliefs simply by counting. Here we study how our results extend to settings in which unbiased beliefs are considerably less cognitively demanding to form. In both robustness checks, we continue to find evidence supporting market selection, but the results are noisier and, unlike in CORE, do not imply a speed-up of market selection relative to the benchmark of optimal intertemporal allocation of wealth.

Our results imply that the average belief that survives highly competitive wealth dynamics in the long run will tend to be considerably different from the sample averages measured in the lab.¹ As a result, care must be taken in exporting (informally or formally) biases measured in the average subject in the lab onto counterfactual market settings. However, as we emphasize in our concluding discussion, this observation carries with it important caveats. First, in non-market settings (for instance in large, insulated bureaucratic hierarchies) agents are to varying degrees shielded from wealth dynamics so that counterfactuals using average behavior in laboratory studies may be considerably more reasonable. Second, our experiment studies one particularly important type of bias with a wide range of applications, but in principle market selection might work better, worse, or not at all with other types of biased beliefs. Research examining the survivability of other non-neoclassical behaviors using the methods we introduce can provide a sharper picture of the corrective scope of wealth dynamics.

The remainder of the paper is organized as follows. In Section 2 we introduce the model and provide theoretical foundations. In Section 3, we describe and motivate our experimental design. In Section 4 we describe the results which we discuss further in Section 5.

¹We are unaware of previous, direct experimental evidence evaluating how well biases measured in the laboratory are likely to survive the effects of wealth dynamics. Perhaps the closest related paper is that of Asparouhova et al. (2015) who show that ambiguity aversion can cause biased agents to voluntarily *self* select out of trade (a related, but distinct mechanism to market selection). Casari et al. (2007) study the role of learning and liquidity constraints in a repeated common value auction environment. Oprea (2014) studies biases towards survival in a cash management task that is related to a version of the market selection hypothesis. A small experimental literature exists on how arbitrage mechanisms – which are quite distinct from the wealth dynamics behind market selection – mitigate the impact of biases on market prices in speculative markets (e.g. Camerer (1987); Ganguly et al. (2000); Kluger and Wyatt (2004); Slembeck and Tyran (2004); Fehr:2005). Finally Oberlechner and Osler (2012) use survey data from currency markets to argue that experienced traders are no less likely to express overconfident beliefs than are inexperienced traders, though they emphasize that the mechanism for this is unclear.

2 Theory and Hypotheses

2.1 Model

We implement a simplified version of the model analyzed in Blume and Easley (1992). Time is discrete with an infinite horizon, $t = 1 \dots \infty$. At each date, one of two possible states of the world, $s_t \in \{0, 1\}$, distributed i.i.d. with $Pr(s_t = s) = q_s \in (0, 1)$, is realized. n agents (labeled $i = 1 \dots n$) are initially endowed with wealth $w_0^i = \frac{Y}{n}$ and at each time, t , consume, and then invest unconsumed wealth across two Arrow-Debreu securities corresponding to the two states of the world.

At each t , each agent i chooses a fraction of her wealth, c_t^i , to consume at each date and a fraction of her wealth remaining after consumption to invest in each asset, α_{st}^i . After agents have made their decisions, the state, s_t , is revealed. Asset $s \in \{0, 1\}$ corresponding to state s_t pays Y and the other asset pays zero. Prices are determined in a competitive market with a single share of each asset. Given constant aggregate wealth of Y , we define a normalized price which is pinned down through market clearing:

$$p_{st} = \frac{\sum_{i=1}^n \alpha_{st}^i (1 - c_t^i) w_t^i}{Y} \quad (1)$$

where w_t^i is the wealth of agent i at time t .² Given the normalized price, each agent's wealth evolves according to

$$w_{t+1}^i = w_t^i \frac{\alpha_{st}^i (1 - c_t^i)}{p_{st}} \quad (2)$$

when state s is realized.

Agents maximize the objective function $\pi^i = E \sum_{t=1}^{\infty} \delta^{t-1} \log(c_t^i)$, where δ is a stationary discount rate. As Blume and Easley (1992) (Proposition 5.1) show, this reward structure has three important properties. First, consumption and investment decisions are *stationary*, meaning we can denote consumption rules by c^i and investment rules by α_s^i . Second, agents employ simple and easily interpretable consumption and investment rules: they “bet their beliefs” by setting α_s^i equal to the probability with which they believe state s will occur, and choose $c^i = 1 - \delta$. Finally, decisions are independent of prices (prices drop out of the first-order conditions) meaning agents make the

² Assuming at least one agent places a positive fraction of wealth on the asset corresponding to the state that is realized. This assumption is never violated in our experimental data.

same decisions at any price. In our experiment we provide subjects with a log reward function to match the model (see Section 3, below, for motivation), allowing us to simplify the experiment and analysis considerably along multiple dimensions. In Section 5, we motivate and state a proposition drawing from Blume and Easley (2006) showing that, under minimal assumptions, the type of conclusions we draw from our experiment under log utility can also be drawn for a much broader range of preferences.

2.2 Selection

Our primary question is whether agents who make relatively biased investments (by setting α_s very different from the true state probabilities q_s) tend to run out of wealth so that those that make less biased investments tend to accumulate wealth. In order to answer this questions we study how an agent's *wealth share*

$$r_t^i = \frac{w_t^i}{\sum_{i=1}^n w_t^i}$$

evolves over time, and how this evolution is impacted by bias in investment rules.³ Because we are interested in the growth rate of the wealth share, we focus on it's natural log at time T :

$$\log r_T^i = \log r_0^i + \sum_{t=1}^T \sum_{s=0}^1 1_{st} \left(\log(1 - c^i) + \log \frac{\alpha_s^i}{p_{st}} \right) \quad (3)$$

where 1_{st} is the indicator function that takes the value 1 at time t if state s is realized, and is 0 otherwise. Following Blume and Easley (1992), we study the relative wealth share across agents i and j , $\log r_T^i - \log r_T^j$. Assuming that each agent starts with the same wealth (as in our experiment), in the limit as $T \rightarrow \infty$ we have

$$\lim_{T \rightarrow \infty} \frac{\log r_T^i - \log r_T^j}{T} = \log \frac{1 - c^i}{1 - c^j} + \sum_{s=0}^1 q_s (\log \alpha_s^i - \log \alpha_s^j). \quad (4)$$

(4) can be rewritten in terms of a measure of bias in the investment rule: the *relative entropy* of an investment rule with respect to the true probabilities, defined as

³We assume each agent uses rules, $c^i < 1$ and $\alpha_s^i \in (0, 1)$. An agent that instead consumes her entire wealth or bets all of her remaining wealth on a single asset ends up with zero wealth almost surely and therefore has no impact on wealth distributions or prices.

$$\begin{aligned}
I_q(\alpha) &\equiv \sum_{s=0}^1 q_s \log \left(\frac{q_s}{\alpha_s} \right) \\
&= q \log \frac{q}{\alpha} + (1-q) \log \frac{1-q}{1-\alpha}
\end{aligned}$$

where we define α to be the investment in asset 0 and q to be the corresponding probability. $I_q(\alpha)$ is weakly positive and is zero only when $q = \alpha$. It serves as a measure of distance between the investment rule and the true probabilities. Applying this definition, (4) becomes

$$\begin{aligned}
\lim_{T \rightarrow \infty} \frac{\log r_T^i - \log r_T^j}{T} &= \log \frac{1 - c^i}{1 - c^j} + I_q(\alpha^j) - I_q(\alpha^i) \\
&= SI^i - SI^j
\end{aligned} \tag{5}$$

where the *survival index* of agent i is defined by

$$SI^i \equiv \log(1 - c^i) - I_q(\alpha^i) \tag{6}$$

Blume and Easley (1992) (Proposition 3.2) show that the agent with the highest survival index accumulates all wealth as $T \rightarrow \infty$ almost surely (and those with lower survival indices lose all wealth).⁴ We restate their result here:

Proposition 1. *Suppose all agents consume constant fractions of wealth and make constant investment decisions. If there is a unique investor i such that for all agents, $j \neq i$, $SI^i > SI^j$, then $\lim_{T \rightarrow \infty} r_T^i \rightarrow 1$ a.s. and $\lim_{T \rightarrow \infty} r_T^j \rightarrow 0$ a.s. for all $j \neq i$. Furthermore, in the limit p_T is determined solely by the investment rule of i , α^i .*

Proposition 1 shows that knowing each agent's survival index is sufficient for knowing which agent comes to own all of the wealth, and therefore whether market selection works (an unbiased agent survives) or fails (a biased agent survives). Clearly if all agents consume optimally (and have the same discount factor), then market selection must work because the difference in survival indices depends only on differences in beliefs.⁵ However, if agents do not consume optimally then two possibilities exist: (i) market selection may fail or (ii) it may work even more rapidly than in

⁴In our data, survival indexes are always unique so that the technicalities that arise with ties are avoided.

⁵This statement is also a special case of a more general result in Blume and Easley (2006).

the optimal consumption case.⁶ To quantify the speed of selection for (ii), we prove Proposition 2 in Appendix A, showing that survival indices not only govern who survives, but also how quickly they come to dominate the market: the speed depends on the distance between the index of the highest survival index agent and that of this agent's nearest competitor.

Proposition 2. *Suppose all agents start with the same wealth, consume constant fractions of wealth, and make constant investment decisions. If there is a unique agent i such that for all agents, $j \neq i$, $SI^i > SI^j$, then the expectation of the time, $\tilde{T}$, at which agent i 's relative wealth share is first expected to exceed that of all other agents by at least a factor of m is given by*

$$E[\tilde{T}] = \frac{\log m}{SI^i - SI^k}$$

where SI^k is the second largest survival index among all remaining agents, $j \neq i$.

To summarize, the theory provides us with two key observations: the agent with the highest SI is selected by the market (survives) and the speed of selection is determined by how much higher her SI is than that of her nearest competitor.

2.3 The Market Selection Hypothesis

According to the market selection hypothesis, agents with biased beliefs lose wealth relative to competitors with less biased beliefs in competitive markets. Proposition 1, on the other hand, says that the agent with the highest survival index, SI_i will be the eventual exclusive survivor in a competitive market. Are these two statements consistent?

The answer depends crucially upon the relationship between agents' investment rules, α^i , and their consumption rules, c^i . Intuitively there are two ways to acquire a high survival index. One is, indeed, to make less biased investment choices by choosing α^i close to the true state probabilities. The alternative is simply to under-consume by setting a low c^i . Because these two factors act as substitutes, the validity of the market selection hypothesis depends crucially upon the joint distribution of α^i and c^i in the population of traders.

⁶The literature provides examples of alternative settings in which the market selection hypothesis fails for reasons other than the ones emphasized in Blume and Easley (1992). For instance, Kogan et al. (2006) and Kogan et al. (2017) show that irrational traders can have persistent influence on prices even though they are eventually driven out of the market. Likewise Blume and Easley (2006) point out that market selection can fail altogether simply due to markets being incomplete.

If agents that make biased α^i choices save (and invest) more than their less biased competitors, the market selection hypothesis can fail, with relatively unbiased investors systematically running out of wealth while relatively biased investors are enriched. On the other hand, if biased subjects save systematically less than unbiased ones, survival indices between the biased and unbiased will be diffuse and, by Proposition 2, market selection will not only occur, but will unfold at a significantly faster rate than if consumption were uniform (for example, if agents uniformly chose the same c^i).

The market selection hypothesis, and the speed with which it operates, is therefore an empirical question, hinging on the relationship between the way subjects make inferences about states of the world (to make investment decisions) and reason their way through intertemporal tradeoffs (to make consumption decisions). In the next section, we describe the experimental design we use to pose this empirical question in the context of a severe type of bias with a wide range of applications.

3 Experimental Design

Each session of our experiment is divided into several periods, each of which is a full run of the model described in the previous section. Each period is, in turn, divided into some unknown (to subjects) number of rounds. Because of the obvious impracticalities of studying infinite horizon settings in the lab, we transform the problem into a theoretically isomorphic indefinite horizon setting: a period ends in each round with probability δ (0.1 in our experiment), creating a shadow of the future and a natural source of discounting (Roth and Murnighan (1978)).⁷ At the beginning of each period, subjects are given $w_0 = 40$ units of wealth and, using the interface shown in Figure 1, are asked to choose:

- a withdrawal rule, c – the fraction of wealth, w_t , to be withdrawn and consumed each round of the period – and
- an investment rule, α – the fraction of un-consumed wealth, $w_t(1 - c)$, to invest in asset 0 (described to subjects as the “green” stock) each round of the period. Remaining wealth is then automatically invested in asset 1 (the “blue” stock).

Subjects input their choices of c and α in boxes in the center of the screen (not pictured in Figure 1). The stationary decision rules we impose are optimal under the log payoff scheme we implement (described below) and allow us both to run standard 2-hour sessions and to avoid serious censoring

⁷We pre-drew period lengths and used the same draws in each session of the experiment.

issues that would arise in non-stationary designs. Subjects’ earnings are determined entirely by their withdrawals (their consumption, $w_t c$) using a log reward function, $\log(w_t c)$. In order to control for risk aversion, we pay subjects in binary lottery tickets, and in order to prevent flat payoff problems,⁸ we, in the background, apply subjects’ decisions to hundreds of distinct period-length draws and pay them the average over all of these payments (In each period we show subjects the evolution of their wealth over one realized period length (see Figure 1), but also report their actual earnings generated by this repetition procedure.) Collectively we call this payoff scheme “LOG-LP” (we show in Section 5 that results measured under this payoff scheme are relevant for a much broader set of preferences). Subjects can view the log transformation on the left side of the screen (see Figure 1) to assist in visualizing earnings.

Each round, subjects are presented with a five-door version of the classic Monty Hall problem and are informed that the winning asset for the round will be determined by the outcome of the problem. Under this procedure for determining the state, the green asset (asset 1) pays out (is the “winning” asset) with probability 0.4, leading an unbiased, optimizing subject to choose $\alpha = 0.4$. However, the naive heuristic employed by most decision-makers in previous experiments with the Monty Hall problem will lead to a choice of $\alpha = 0.67$ (see footnote 21). Thus, we expect unbiased subjects to put most of their wealth on the blue asset but for most subjects to express a strong bias by putting more of their wealth on the green asset. (Details of the decision problem and its framing to subjects are provided in Appendix C and in the instructions in Appendix D.) Each round, any wealth the subject has invested in the winning asset is multiplied by a multiplier $m_{st} > 1$ (see the next subsection for details) to determine w_{t+1} , the wealth the subject begins the next round with. The process is visualized on the right side of Figure 1 in the “Francs Graph”. The fraction of wealth allocated to the winning stock last round is shown in green or blue (depending upon the color of the winning stock). An arrow points to the new wealth with the multiplier listed beside it (in this example, “2.1x”).

In Appendix C, we discuss these design choices – including the stationary setting, the payoff scheme and the use of the Monty Hall task to determine asset values – in more detail.

⁸In particular, paying the expectation over period lengths allows us to calibrate the experiment to avoid subject payoffs that vary little with their decisions. Inducing a reward function that pays in lottery tickets requires a lower bound of zero tickets, excluding the (infinitely) negative payoffs that occur as consumption goes to zero. Truncating earnings at zero therefore distorts optimal consumption decisions upwards. One can endow subjects with tickets prior to each period to lessen the chance of the lower bound binding, but this weakens incentives by making the number of tickets earned less sensitive to one’s decisions. By paying subjects their expected earnings over period lengths, we need only ensure the lower bound doesn’t bind too often on average, rather than in the extreme case of a long period. Therefore, we can reduce the endowment of tickets, substantially strengthening payoff salience.




Figure 1: Screenshot of the experimental software.

3.1 Treatments

Our experimental design consists of a main treatment we call CORE and two robustness treatments, END and LEARN.

CORE Treatment. In our main CORE treatment, prices are exogenous to subjects' decisions, allowing us to perfectly implement the price-taking assumption in Blume and Easley's (1992) general equilibrium model. To do this we simply fix $m_{st} = 2.1$ exogenously; importantly, under the LOG-LP payoff scheme used in the design, prices drop out of first order conditions, allowing us to make this transformation without changing subjects' optimal decisions from those described in Section 2, above. By implementing the price-taking assumption of the model, the CORE treatment gives us our most direct implementation of the Blume and Easley's (1992) and thus serves as our cleanest environment for posing our main questions.

END Treatment. Relaxing the price-taking feature of the CORE treatment allows us to evaluate the importance of the competitiveness assumption of the theory to the process of market selection. For this reason, we ran a second cohort of subjects through ten independent, 8-person⁹ markets (with fixed matching across periods) in which everything is identical to the CORE treatment except that prices are endogenously determined by subjects' decisions, so that subjects are not, formally, price-takers. In this treatment, the multiplier, m_{st} , is determined directly by subjects' decisions via

⁹One of the markets consisted of only 7 traders due to low subject show-ups to the session.

prices generated through a parimutuel betting market, $m_{st} = \frac{1}{p_{st}}$, where p_{st} is directly determined by the wealth subjects allocate to each asset as described by (1). Moreover, we inform subjects of the fraction of wealth allocated to each asset in each round, giving them the ability to update their beliefs based on overall market investments. This Endogenous Price (END) treatment allows for possible strategic decision-making on the part of subjects, a condition under which theory provides us less guidance. Our question is whether the relationship between c and α measured in CORE changes as markets become less competitive and whether imperfect competition therefore has implications for market selection.

LEARN Treatment. We made an *ex ante* design decision to (i) run eight periods total in CORE sessions but (ii) to end the CORE treatment once subjects had seen a sufficient number of draws to learn away biases by sheer statistical learning (this “learnability” threshold occurs after completing period 4).¹⁰ This procedure ensures that the CORE treatment studies how market selection operates on belief distributions generated by subjects’ *ex ante reasoning* about the Monty Hall process (without the benefit of learning). After the close of the CORE treatment, we continue the session for an additional four periods. In these additional periods, beliefs are shaped not only by deductive reasoning about the state-generating process, but also by statistical learning about the state drawn from having observed dozens of state realizations: by the time subjects enter this additional treatment – which we call LEARN – subjects have observed 56 realizations of the Monty Hall problem and have enough information to reject the naive heuristic of 0.67 at the five percent level. Importantly, while avoiding bias in the CORE is quite cognitively demanding, requiring sophisticated Bayesian reasoning, doing so in LEARN requires little more than an ability to count. The LEARN treatment thus allows us to study whether selection operates differently on biases that require less cognitive ability and effort to avoid.

Summarizing, the CORE treatment serves as our main treatment. The END treatment allows us to ask whether these results extend to less competitive environments and the LEARN treatment allows us to ask whether they extend to settings in which beliefs spring from a different (and less cognitively difficult) source (see Section 3.3 for hypotheses about how the different treatments might change behavior).

¹⁰Specifically, we run a binomial test on the set of state realizations observed by subjects and end the CORE treatment once subjects have seen enough draws to reject the naive heuristic of $\alpha = 0.67$ at the five percent level. This occurs at the end of period 4 (recall we use the same draws across all sessions) which happens to be very long (42 rounds). Prior to this period subjects had seen only 4, 8 and 2 realizations in the preceding 3 periods of play.

3.2 Implementation Details

We ran the experiments at the EBEL laboratory at UC Santa Barbara in November-December 2015. A total of 130 subjects participated (51 in CORE/LEARN and 79 in END).¹¹ Subjects participated in four periods of each treatment, preceded by eight periods of practice.¹² In the practice periods, asset values were determined not by a Monty Hall problem but rather by simple urn draws with known probabilities explicitly communicated to subjects (varied between 0.5 and 0.25 across practice periods). These practice periods allowed subjects to learn about the software and environment without exposing them to feedback about the Monty Hall problem. Subjects were drawn randomly from across majors using the ORSEE (Greiner (2015)) online recruiting tool and consisted of undergraduate students at UC Santa Barbara. Subjects were read instructions aloud, asked to pass a computerized quiz, and then participated in an unpaid practice period before the paid practice began. After the practice periods, we paused to provide new instructions for the Monty Hall task (see Appendix D) and then began the main experiment. Average subject earnings were \$26.31, including a \$5 showup fee.

3.3 Questions and Hypotheses

An earnings maximizing subject will consume $c = 0.1$ of her wealth each round, and allocate $\alpha = 0.4$ of the remainder to asset 0 (0.6 to asset 1) in our main experiment. We say a subject *over-consumes* if she sets $c > 0.1$, is *Bayesian* if she systematically invests in the more likely asset by setting $\alpha < 0.5$, and is *biased* if she chooses $\alpha > 0.5$. Based on the evidence in previous experiments on consumption-savings decisions (e.g. Hey and Dardanoni (1988); Ballinger et al. (2003); Carbone and Hey (2004); Ballinger et al. (2011)), we expect subjects, on average, to over-consume even after a great deal of experience. Based on previous experiments in Monty Hall-type environments (e.g. Friedman (1998) and Kluger and Wyatt (2004)), we expect most subjects to be biased, and the modal subject to invest close to the naive heuristic level of 0.67, particularly early in the experiment.

¹¹We observe most subjects for four full periods in each of these treatments. In several periods software glitches caused subjects' decisions to be improperly submitted, often by causing subjects to accidentally submit $c = 0$. Subjects reported 7 instances of this occurring during the experiment and we identified another 5 cases (in which submissions of $c = 0$ occurred). We dropped these 12 observations (out of 724 total) from the dataset.

¹²The CORE and LEARN sessions were run within-session back-to-back – sessions thus ran for eight periods of actual play. We also extended END sessions for an additional four periods but software errors occurred in these later periods that plausibly interfered with learning in these periods. We therefore report results from these extension periods of the END sessions in Appendix B.

Our main motivating question is whether the joint distribution of c and α is such that strongly biased subjects will tend to run out of wealth while their less biased (or perhaps even unbiased) competitors survive. As shown in Section 2, this question is equivalent to asking whether the survival index, $SI \equiv \log(1 - c) - I(\alpha)$, formed of subjects' c and α choices, tends to be lower for subjects who make biased α choices than for those who make unbiased α choices. As we emphasize in Section 2.3, this question is not trivial: its answer depends on the empirical relationship between α and c in the population. If α and c are strongly negatively related – if subjects setting $\alpha > 0.5$ also tend to hoard wealth by setting a low c – the market selection hypothesis will fail. If, however α and c are positively related, market selection will occur, and do so at a faster rate than if α and c were statistically unrelated (for example, if everyone consumed optimally).

What might we expect based on prior research? The evidence is mixed. For instance, Ballinger et al. (2011) provide evidence that subjects that perform poorly in cognitive measurement batteries are systematically more likely to over-consume. To the degree that the outcomes of these batteries predict biases in belief formation, we might expect α and c to be positively related, supporting the market selection hypothesis. On the other hand, Oprea (2014) provides experimental evidence that subjects in a cash management task tend to be prone to a serious error of hoarding cash due to an over-attachment to survival. If such “survival bias” is related to inferential biases, we might expect the opposite, with higher α being paired with lower c . Finally, some types of overconfidence might lead subjects to bet too much wealth on the state, under-consuming in the process. If overconfidence is related to biased belief formation then, again, α may be inversely related to c , preventing market selection from occurring.

Our treatment design also allows us to examine whether the answer to this question is sensitive to the intensity of competition (reduced in the END treatment) or the difficulty of forming rational beliefs (reduced in the LEARN treatment). We have several reasons to think that these factors might matter. With imperfect competition – the END treatment – an additional reason for setting a high value of c (beyond intertemporal reasoning errors like myopia) arises: doing so is a collusive strategy in the Blume and Easley (1992) setting we implement in the lab.¹³ This additional motive for overconsuming may weaken the relationship between c and α in such a way as to weaken market selection – even highly rational subjects who would otherwise avoid over-consuming might be willing to do so for cooperative purposes.¹⁴ Likewise, forming unbiased beliefs in the LEARN treatment is

¹³Because the aggregate wealth is fixed in each period, the higher the aggregate consumption rate, the higher the aggregate return. Thus, if subjects could collude, they would want to consume all but epsilon of their wealth in each period. Of course, such a strategy is not individually rational.

¹⁴A subject who understands that she affects prices, rationally overconsumes relative to the perfect competition

less cognitively demanding than it is in CORE – in CORE one must engage in complex conditional reasoning while in LEARN one need only count well to eliminate severe biases. We might thus expect the pool of unbiased subjects in the LEARN treatment to be less sophisticated (relative to biased subjects) than the pool in the CORE treatment, making the relationship between α and c less conducive to selection.

4 Results

In Sections 4.1 through 4.3, we report results for our main treatment, CORE. In Section 4.1, we examine the raw distributions of c and α , reporting highly biased, but heterogeneous, outcomes for each of these choice variables. In Section 4.2, we report our main results, showing that c and α are positively related, that survival indices are maximized near the Bayesian benchmark level of 0.4, and that, based on these measurements, biased traders will tend to be driven from markets. In Section 4.3, we compare these results to those from counterfactual transformations of the dataset, showing that (i) the market selection hypothesis could easily have failed under alternative joint relationships between our observed distributions of α and c , and (ii) our estimates imply that market selection will occur at a much faster pace than under the counterfactual in which all agents save optimally. We illustrate these results with simulations using measurements from the data. In Section 4.4, we repeat the same analysis for the robustness treatments END and LEARN and report similar, though noisier, results.

Prior to the beginning of the main experiment, subjects ran through eight practice periods. In these periods, rather than solving Monty Hall problems to infer the probability each asset pays off, subjects were directly told the relevant probabilities (which alternate between values of 0.25 and 0.5 for the probability asset 1 pays off) and asked to make investment and consumption choices. These practice periods serve two functions. First, they allow us to assess whether or not our efforts to induce log utility were successful. If so, we expect subjects to bet their beliefs by allocating wealth across assets according to the true probabilities given to them. This is just what we observe in the experiment: in every period in which the state probabilities are equal, the median subject splits wealth evenly across assets 1 and 2; in every period in which state 1 occurs only with 0.25 probability, the median subject allocates 25% of their wealth to asset 1 and 75% to asset 2. Second, the practice periods allow subjects to learn about their environment, in particular how benchmark. If the subjects that understand this fact are also those that have better beliefs, it would also tend to weaken market selection.

consumption choices affect their payoffs. We find that subjects' median consumption decisions fall from 0.25 points in the first practice period to 0.15 in the first period of the actual experiment, and are relatively steady after. This evidence suggests that subjects' consumption decisions when they enter the real periods represent not random behavior but a learned and considered response to their environment.

4.1 Marginal Distributions

We begin our analysis of the main data by reporting the raw, marginal distributions of subject-wise median values of α (the rate of investment in asset 0) and c (the rate of consumption), the two variables subjects choose in the experiment.

As expected, most subjects in the CORE treatment show evidence of highly biased beliefs over the state generating process, leading to highly biased investment choices. The left panel of Figure 2 plots a CDF of α (as a solid line), revealing that over 80% of subjects systematically invest in the asset *less likely* to be valuable given the Monty Hall process (by setting $\alpha > 0.5$), and that the median and modal subjects employ the naive heuristic commonly employed in similar inference problems, setting α near 0.67. Crucially, however, there is also heterogeneity: just under 20% of subjects make unbiased investments and a small set of subjects reveal Bayesian beliefs, choosing $\alpha = 0.4$.

The middle panel of Figure 2 plots CDF's of the consumption rate, c , revealing that subjects also tend to make sub-optimal consumption/savings decisions. Over 80% of subjects over-consume and the median subject overconsumes by over 50% (choosing $c = 0.16$) but, again, heterogeneity exists: just over 15 % of subjects consume at the optimal level of 0.1 (or less). These characteristics of the marginal distributions of c and α provide us with a first result:

Result 1. *Most subjects exhibit severely biased beliefs, systematically betting on the less likely state by setting α near the naive heuristic level of 0.67. Subjects also tend to overconsume, with the median subject setting c over 50% higher than the optimal level.*

4.2 Main Results: Joint Distribution

Subjects tend to form highly biased beliefs (set α considerably too high); our main question is whether α is jointly distributed with c in such a way as to allow these suboptimal beliefs to *survive* competitive wealth dynamics. In order for this to happen, subjects with biased beliefs must hoard

more wealth (choose a lower c) than subjects with relatively unbiased beliefs, generating a higher survival index, SI, than their competitors. To find out whether this is the case, we first directly measure the relationship between consumption choices and beliefs by estimating:

$$c_{ij} = \nu + \omega\alpha_{ij} + \varepsilon_{ij} \quad (7)$$

where i indicates subject, j indicates period, and ε_{ij} is a normally distributed error term, clustered at the subject level (recall that subject i makes only one decision per period j).¹⁵ Given that almost no subject sets $\alpha < 0.4$, if more biased subjects hoard wealth, ω will tend to be negative. Estimating (7) with standard errors clustered at the subject level, we find the opposite: ω is instead significantly *greater* than zero (0.15, $p = 0.021$), suggesting that more biased subjects in fact consume systematically *more* than unbiased subjects.

We show this result non-parametrically in the center panel of Figure 2 by plotting as dotted lines the CDF of the subsample of subjects who have near-Bayesian beliefs (set α within a 5 percentage point band of 0.4). This subsample makes nearly optimal consumption choices (0.118 at the median) and the CDF is significantly to the left of the sample as a whole ($p = 0.025$, Mann Whitney test), confirming that consumption among Bayesian subjects is substantially smaller than consumption overall.

The implication of this estimate for the survivability of Bayesian beliefs is illustrated in the right-most panel of Figure 2, which graphs a scatter plot of observed α choices against survival indices (calculated via (6) using subjects' α and c choices). Overlaying a fit to the data using our estimate from (7), we observe a clear relationship between survival indices and beliefs. The curve of the fitted relationship is maximized at 0.36, far below the sample average of 0.61 and close to the Bayesian benchmark $\alpha = 0.4$, with subjects with biased beliefs ($\alpha > 0.5$) tending to have substantially smaller survival indices. Thus, based on the joint distribution of α and c , we expect Bayesian beliefs to tend to out-survive other beliefs.

In the left-hand panel of Figure 2 we show a similar result non-parametrically by plotting a dotted line for the CDF of α for the sub-sample of subjects whose choices place them in the top decile of survival indices. Strikingly, subjects in this sub-sample are uniformly unbiased (choose $\alpha \leq 0.5$) and over 80% are perfectly Bayesian, in strong contrast with the highly biased CDF from the overall sample. This result is highly significant ($p < 0.001$ by a Mann Whitney test) and confirms that the subjects that are most likely to survive are considerably less biased than the

¹⁵The Bayesian Information Criterion prescribes a linear relationship between c and α here.




Figure 2: CORE Treatment Results.

subject population as a whole.

Result 2. *Bayesian subjects have systematically higher survival indices than biased subjects and high survival index subjects tend to be Bayesian, supporting the market selection hypothesis.*

Figure 3 illustrates the implications of these estimates. We populate 200 simulated markets with agents employing (c, α) pairs based on all of our observations of α and their corresponding fitted values of c . We then generate prices and returns over time as in Blume and Easley (1992).¹⁶ We plot, in a solid black line with dots, the mean share of wealth held by weakly biased agents (i.e. agents who do not bet systematically on the better asset, choosing $\alpha < 0.5$) across all of our simulated markets for 150 rounds of play. The results show that the wealth share of biased subjects starts high (most subjects in our sample are highly biased) but converges to zero over time (with unbiased subjects acquiring all of the wealth), illustrating how our estimated results support the market selection hypothesis.

¹⁶Specifically, we collect the 202 α choices and the corresponding fitted values of c from (7) in the CORE treatment, assign each pair to an agent in the market, give each an equal share of initial wealth, and randomly draw realizations of the state for each of 150 rounds. We repeat this entire process 200 times, generating 200 markets of simulated data.




Figure 3: CORE Treatment Simulated Wealth Dynamics.

4.3 Counterfactuals

The data supports the market selection hypothesis – in this subsection we provide evidence that neither the degree nor speed of market selection suggested by the data are trivial outgrowths of the observed belief and consumption distributions, but are instead strongly dependent on the specific joint distribution we observe. To make this case, we compare our data to two *counterfactual* joint distributions that feature the exact same marginal distributions as in our data but that co-vary α differently with c . Doing so allows us to demonstrate that a different relationship between the two would have generated dramatically different results.

First, we show that market selection could have fundamentally failed if the joint distribution between α and c had been different, even if the marginal distributions themselves had been identical. In the “Negative Counterfactual” we order α and c in order to generate the strongest possible negative correlation between the entropy of α and c , generating an environment in which particularly low consumption choices tend to be paired with particularly poor beliefs.¹⁷

¹⁷We order α from lowest to highest entropy value, $I(\alpha)$, and do the same for c . Pairing them with these orderings generates the highest possible positive correlation between entropy and log-savings given the marginal distributions.

To examine how this counterfactual changes results, we re-estimate (7) on the resulting dataset and plot the fitted relationship between α and the survival index under the Negative Counterfactual as a dotted red line on the right-hand panel of Figure 2. The results reveal that under the Negative Counterfactual, survival indices are maximized at a highly biased belief ($\alpha = 0.613$) rather than at near-Bayesian levels, suggesting that if the joint distribution had been different, biased agents could have had systematically higher survival indices and therefore have systematically out-survived unbiased agents. In Figure 3, we plot, in red, the results of simulations using these counterfactuals which show that, in contrast to simulations based on the actual data, biased agents' wealth shares quickly rise to one, as rational agents are systematically driven from the market. Thus, again, if the joint distribution had been different, the market selection hypothesis would have failed in a rather fundamental way.

Result 3. *An alternative relationship between the marginal distributions of α and c in the joint distribution could have generated a failure of market selection in which rational behavior is systematically driven from the market. Our data's support for the market selection hypothesis is thus non-trivially driven by the shape of the joint distribution.*

Second, we show that the market selection hypothesis could have operated significantly more slowly if, instead of exhibiting the heterogeneous consumption rules observed in our sample, all subjects had consumed optimally (chosen identical consumption rules). In the "Optimal Counterfactual" we retain the exact same marginal distribution of α as in the actual data but assign each agent an optimal $c = 0.1$. As Blume and Easley (2006) establish, under optimal consumption, market selection must function: relatively biased agents must be driven from the market eventually. However, as we prove in Section 2, the speed with which market selection works depends upon the magnitude of the difference between the survival indices of relatively biased and unbiased agents. Conditional on the joint distribution being such that market selection functions, greater differences in c 's across biased and unbiased agents (as in our data) tends to cause market selection to go faster than if everyone had consumed optimally.

To illustrate, we use estimates from (7) to calculate the expected time it would take for unbiased agent i to accumulate m times the wealth of biased agent j , given by the expression,

$$E \left[\tilde{T} \right] = \frac{\log m}{SI^i - SI^j} \quad (8)$$

derived in Proposition 2. Specifically, we consider how long it would take for a perfectly Bayesian agent with $\alpha = 0.4$ to acquire 99% of the market's wealth ($m = 100$) when matched with a

hypothetical subject with beliefs given the sample average (i.e. with $\alpha = 0.61$) . We compare the value implied by our estimates to the value generated if subjects had each consumed optimally, selecting $c = 0.1$. Our estimates imply the market would take 35.80 rounds on average, while the same degree of selection would require 51.36 rounds – over 50% longer – under the Optimal Counterfactual. Using the delta method and expression (8), we can reject the hypothesis that the two time requirements are identical at the 0.001 level via a non-linear Wald test.

We further illustrate the difference in speeds in Figure 3 by running a simulation of wealth dynamics as above except with Optimal Counterfactual agents. The simulation confirms that unbiased agents accumulate wealth (and biased agents lose wealth) much faster under the parameters observed in our data (the solid black line) than under the Optimal Counterfactual agents (broken gray line).

Result 4. *The form of heterogeneity in consumption observed in our data improves the speed at which market selection operates relative to the counterfactual case in which subjects optimally consume.*

Together, these counterfactuals emphasize the influence of the measured joint distribution on our conclusions. If c and the entropy of α had been negative related, market selection would have failed, driving rational agents from the market; if they had been less positively related (for instance if all agents had been near-optimal savers), market selection would have occurred but at a significantly slower pace.

4.4 Robustness: Imperfect Competition and Learning

We report two additional treatments designed to assess the robustness of our main results. In this subsection, we repeat the previous analysis for these treatments to assess how our findings change under less intensively competitive conditions (the END treatment) and in an environment in which forming beliefs is less cognitively demanding (the LEARN treatment).

4.4.1 Robustness: Marginal Distributions

Figures 4 and 5 reconstruct Figure 2 for the END and LEARN treatments. In the CDFs in the left and middle panel of each we also superimpose the CDF of the corresponding behavior from the CORE treatment in gray for reference.

In the END treatment, the belief distribution α is virtually identical to that in the CORE

treatment, but the c distribution is shifted far to the right, revealing that subjects consume considerably more in the imperfectly competitive END treatment than in the perfectly competitive CORE treatment (the difference in distributions is significant by a Mann-Whitney test ($p = 0.014$)). As we describe in section 3.3, overconsumption is in fact a collusive strategy in the END treatment, and our results suggest that subjects take advantage of this fact in the imperfectly competitive environment.

In the LEARN treatment, by contrast, the consumption distribution is similar to that in the CORE treatment (there is a statistically significant but very small 0.01 point reduction at the median), but the belief distribution α is shifted substantially to the left (a median reduction of 10 percentage points in expressed beliefs). While subjects do not significantly learn to consume much more rationally with experience, they do learn via statistical feedback to substantially improve their investment choices ($p < 0.001$). This change in the belief distribution provides evidence that beliefs in the LEARN treatment are not the same sort of beliefs as those in the CORE treatment, just as hypothesized.

We collect these initial observations in our next result:

Result 5. *Under imperfect competition (END) consumption is significantly higher than under perfect competition (CORE). Under learning (LEARN), beliefs are significantly less biased than in the CORE treatment.*

4.4.2 Robustness: Joint Distributions and Selection

We re-estimate model (7) on data from the END and LEARN treatments.¹⁸ As in the CORE treatment, ω is not estimated to be significantly negative in either case, supporting the market selection hypothesis; unlike in the CORE treatment, however, ω is also not significantly positive in either case. Therefore, both relaxing perfect competition and introducing learning weaken the relationship between beliefs and consumption. As above, we show a similar result non-parametrically in the second panels of Figures 4 and 5 by plotting the distribution of consumption choices for the subsample of subjects making near-Bayesian investments using dotted lines. Unlike in the CORE treatment, near-Bayesian subjects' consumption distributions are virtually identical to the distributions for the sample as a whole in both END and LEARN, again showing that rational subjects do not consume much better in these treatments.

¹⁸None of the results reported below change if we cluster standard errors from this model at the market level rather than the subject level in the END treatment, where subjects interact in 8-person markets.




Figure 4: LEARN Treatment Results.




Figure 5: END Treatment Results.

The third panels of Figure 4 and 5 reveal the effect of these estimated relationships on market selection: survival indices tend to reach their maximum near the Bayesian α choice of 0.4 (the estimated maximum is 0.412 in each case), just as in the CORE treatment (though the scatter plots suggest the relationship is somewhat noisier, particularly in the imperfectly competitive END treatment). The left-hand panels of 4 and 5 provide non-parametric evidence by plotting in dotted lines the CDFs of subjects in the top decile of the survival index distribution; the results suggest that high index subjects tend to have better beliefs than the raw sample (this is significant in END and marginally significant in LEARN: $p < 0.001$ and $p = 0.07$ respectively). The red dotted line in the third panel of each of these figures shows estimates from the Negative Counterfactual for each treatment and, as in the CORE treatment, suggests that under different joint distributions of α and c , survival indices would have been maximized at very biased beliefs, generating severe failures of market selection.

Using expression (8) and our estimates, we can also compute the expected time it takes for a Bayesian to drive out the average subject in the LEARN and END treatments (the mean α is 0.54 and 0.64 in LEARN and END, respectively). We estimate these times as 151.38 and 44.6 in the LEARN and END treatment, respectively, and, unlike in the CORE treatment, find that neither is statistically different from the corresponding Optimal counterfactual estimates of 116.9 and 38.86 ($p = 0.789$, $p = 0.548$). This change relative to the CORE treatment is a straightforward consequence of the fact that in both END and LEARN (and in contrast to CORE), unbiased subjects are no better savers than biased subjects, creating smaller differences in the survival indices and eliminating the speed boost observed in CORE.

Result 6. *Market selection continues to operate under both imperfect competition and learning-generated beliefs, though behavior is considerably noisier. However, in both robustness treatments, market selection is no faster than under the optimal consumption benchmark.*

Finally, in the two panels of Figure 6, we again illustrate our findings by running simulations of 200 markets in each treatment using actual subject decisions, fits from estimation, and counterfactual variations (as we do to construct Figure 3 above). As in the CORE treatment, market selection succeeds using actual data but fails in the Negative Counterfactual. However convergence takes considerably longer in both of these treatments than it does in CORE, as we discuss above. Indeed, unlike in the CORE treatment, END and LEARN wealth shares do not converge at a faster rate than under the Optimal counterfactual (and in fact converge a bit slower). The results suggest that market selection continues to operate, eliminating highly biased beliefs even in imperfectly competitive environments and in the presence of significant learning opportunities. In both cases,




Figure 6: Robustness Treatments (LEARN and END) Simulated Wealth Dynamics.

however, a weaker relationship between α and c implies noisier and slower convergence than in the CORE treatment.

5 Discussion

The statistical relationship between beliefs and intertemporal allocation choices we observe in our data suggests that the severe biases that most of our subjects suffer from are unlikely to survive in competitive markets – indeed patterns in our main CORE treatment suggest that market selection is likely to kill off biases at a *faster* rate than it would under natural benchmarks in which agents save optimally. Our design takes some first steps at assessing the robustness of these findings by examining how sensitive they are to (i) the way beliefs are formed and (ii) the intensity of competition in the market: when we make it less cognitively demanding for subjects to form unbiased beliefs (as we do in the LEARN treatment) or when we ease the degree of competition (as we do in the END treatment), market selection continues to operate but is slower and noisier. In this section, we discuss further, natural questions of robustness and offer some interpretations of our results.

First, our experiment was designed to cleanly measure α and c by implementing a very specific decision setting (e.g. induced log rewards, stationary decision rules). As we discuss in 3, this design allows us to directly measure beliefs, avoid severe censoring problems, and directly implement the

competitive premises of the model we implement in the lab – features of the experiment that greatly improve our ability to answer our motivating questions. A natural question is to what degree can we use the joint distribution of α and c measured in the data to extrapolate to richer (e.g. non-stationary) markets with naturally occurring objective functions (e.g. subjects’ own utility functions). Proposition 3 below shows that, under a minimal set of assumptions, we can use behavior measured in our experiment to draw positive conclusions about market selection in these broader settings. Specifically, suppose that the overconsumption mistakes subjects make can be modeled as occurring because subjects put too little weight on the future when optimizing (i.e. they optimize as if future events matter less to earnings than they actually do by optimizing against using some $\delta^* > \delta$).¹⁹ Then we can conclude that our finding that Bayesian beliefs are likely to drive out biased beliefs continues to hold in any market of risk-averse, expected utility maximizing agents employing possibly non-stationary decision rules:

Proposition 3. *Suppose all agents are risk-averse expected utility maximizers and that there is a unique investor i such that for all agents, $j \neq i$, i has the least biased beliefs (lowest relative entropy). Furthermore, suppose that all agents maximize as if they have discount factors weakly less than optimal, and that agent i ’s pseudo-discount factor is closest to the actual discount factor. Then, $\lim_{T \rightarrow \infty} r_T^i \rightarrow 1$ a.s. and $\lim_{T \rightarrow \infty} r_T^j \rightarrow 0$ a.s. for all $j \neq i$.*

Our conclusion in support of market selection follows from the non-negative relationship we observe between consumption and belief biases. Proposition 3 extends Proposition 1 to non-stationary environments, and tells us that we can extend this conclusion, based on our results, to a much more general range of settings than the one studied in our experiment. Nonetheless, sacrificing some of the measurement advantages we achieve with our design in order to directly gather insight into how richer dynamics and possibly non-EU preferences influence market selection in the laboratory (see e.g. Borovička (2015) for a theoretical example) seems a natural avenue for future research.

Second, our experiment focuses on one specific type of biased belief – a conditional reasoning failure that arises in the Monty Hall problem. We chose to study Monty Hall-like problems in part because of their reliable tendency to generate severe and wide-spread biases, without which we would be unable to study the corrective effects of wealth dynamics. However, another important reason we focused on reasoning problems of this type is that they are closely related to a much wider class of reasoning problems in economics. Monty Hall-like failures are closely related to

¹⁹We emphasize here that we have in mind an optimizing mistake rather than variation in subjects’ time preferences. Because subjects do not consume during the experiment, idiosyncratic time preferences are unlikely to explain overconsumption in our data.

the biased reasoning behind the winner’s curse, failures of no-trade theorems, some types of asset market bubbles, and other dynamic phenomena such as the hot hand fallacy. Because of this, our Monty Hall-like task acts as a stand-in for some of the best known and most severe biases studied in behavioral and experimental economics in recent decades.

Nonetheless, it seems probable based on our results that market selection operates with different levels of success over different types of beliefs. Market selection works particularly well in our CORE treatment where subjects must employ sophisticated reasoning to invest in an unbiased fashion, and reducing the difficulty and cognitive load of forming unbiased beliefs seems to influence the success of market selection: when we allow subjects to lean on much simpler learning heuristics to form their beliefs in the LEARN treatment, market selection continues to work but weakens and slows as the positive relationship between beliefs and savings rules weakens. One conclusion we might draw from this result is that biases that arise from cognitively difficult reasoning problems might be particularly good candidates for market selection, perhaps because (as results from Ballinger et al. (2003) suggest), savings errors are directly related to cognitive ability. If so, we might expect mistaken beliefs that are less driven by intelligence (perhaps mistakes related to framing or poorly applied heuristics) to be less correctable by market selection. Likewise, some biases might be linked to savings errors via other channels, creating additional scope for failures of market selection. For instance, one reason a subject might *under-save* is over-confidence: over-confident subjects might bet a greater share of their wealth than is optimal. Belief errors related to ego and self-regard (for instance beliefs concerning one’s own abilities or the precision of one’s interpretation of information) that also have roots in overconfidence may therefore be negatively correlated with consumption rates, creating greater scope for failures of market selection. In principle, the relatively simple individual decision-making task studied in our CORE treatment could be run with many other types of mistakes measured in the laboratory to assess their “survivability” and to evaluate their likelihood of persisting in the face of competitive wealth dynamics.

Third, our results also suggest that characteristics of the market that are unrelated to beliefs may influence the way subjects form savings rules, with important consequences for market selection. In our END treatment, imperfect competition generates a collusive motive to over-save that is not present in the perfectly competitive setting of the CORE treatment, causing unbiased subjects to over-consume as severely as unbiased subjects (a marked change from the CORE treatment). The market institution thus directly impacts the *reason* subjects over-consume, weakening the relationship between c and α , and causing market selection to slow down substantially relative to the CORE treatment. Other features of markets may have similar influence over the way

subjects make consumption errors and therefore the effectiveness of market selection. For instance, our choice of $m_{st} = 2.1$ in the CORE treatment influences whether and at what speed subjects' wealth declines when they invest based on mistakes in beliefs. Oprea (2014) provides evidence from cash management tasks suggesting that subjects are prone to strong and persistent biases towards survival that cause them to under-consume in order to avoid losses of wealth. If survival biases of this sort correlate with errors in belief formation, subjects with biased beliefs might under-save more dramatically in market environments in which wealth declines can be more effectively stemmed by under-saving – a market feature influenced by m_{st} – generating failures of market selection. Studying whether variation in m_{st} influences the effectiveness of market selection via this sort of channel seems a promising avenue for future research. Likewise, incomplete markets (e.g. Blume and Easley (2006) and Beker and Chattopadhyay (2010)), where market selection can fail even if agents uniformly consume optimally, are a natural setting for future study.

Finally, it is useful to consider what our results imply about the persistence and influence of biases in settings external to the lab. Our results suggest that beliefs that survive in highly competitive markets may not look much like beliefs measured in laboratory experiments due to the influence of wealth dynamics. Were we to take the sample average from an experiment on Bayesian reasoning and export it as a parameter into the setting of a competitive market we would producing a misleading portrait of the performance of the market in the long run. In this sense our results suggest caution be taken in projecting results from laboratory settings onto markets that are well suited to market selection (i.e. highly competitive, complete markets) and provides a method for evaluating the survival value of lab-measured biases in future work. However it bears remembering that much of human economic life takes place in settings in which institutional structures shield decision-makers from the disciplining influence of wealth dynamics (e.g. large, hierarchical organizations, political hierarchies, monopolistic industries). In such environments, market selection is likely to fail to operate (or operate very slowly), allowing biased behaviors measured in the lab to survive to influence economic outcomes.

References

- Alchian, A.**, “Uncertainty, evolution, and economic theory,” *Journal of Political Economy*, 1950, 58, 211–221.
- Asparouhova, E., P. Bossaerts, J. Eguia, and W. Zame**, “Asset Pricing and Asymmetric Reasoning,” *Journal of Political Economy*, 2015, 123, 66–122.

- Ballinger, T., E. Hudson, L. Karkoviata, and N. Wilcox**, “Saving behavior and cognitive abilities,” *Experimental Economics*, 2011, *14*, 349–374.
- , **M. Palumbo, and N. Wilcox**, “Precautionary savings and social learning across generations: an experiment,” *Economic Journal*, 2003, *113* (490), 920–947.
- Beker, P. and S. Chattopadhyay**, “Consumption dynamics in general equilibrium: A characterisation when markets are incomplete,” *Journal of Economic Theory*, 2010, *145* (6), 2133–2185.
- Blume, L. and D. Easley**, “Evolution and market behavior,” *Journal of Economic Theory*, 1992, *58*, 9–40.
- and —, “If you’re so smart, why aren’t you rich? Belief selection in complete and incomplete markets,” *Econometrica*, 2006, *74* (4), 929–966.
- Borovička, J.**, “Survival and long-run dynamics with heterogeneous beliefs under recursive preferences,” 2015.
- Camerer, C.**, “Do biases in probability judgement matter in markets? Experimental evidence,” *The American Economic Review*, 1987, *77* (5), 981–997.
- Carbone, E. and J. Hey**, “The effect of unemployment on consumption: an experimental analysis,” *Economic Journal*, 2004, *114* (497), 660–683.
- Casari, M., J. Ham, and J. Kagel**, “Selection bias, demographic effects, and ability effects in common value auction experiments,” *The American Economic Review*, 2007, *97* (4), 1278–1304.
- De Long, J., A. Shleifer, L. Summers, and R. Waldmann**, “The survival of noise traders in financial markets,” *The Journal of Business*, 1991, *64* (1), 1–19.
- Friedman, D.**, “Monty Hall’s three doors: construction and deconstruction of a choice anomaly,” *The American Economic Review*, 1998, *88* (4), 933–946.
- Friedman, M.**, “Essays in positive economics,” 1953. University of Chicago Press, Chicago.
- Ganguly, A., J. Kagel, and D. Moser**, “Do asset market prices reflect traders’ judgement biases?,” *Journal of Risk and Uncertainty*, 2000, *20* (3), 219–245.
- Greiner, B.**, “Subject pool recruitment procedures: Organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

- Hey, J. and V. Dardanoni**, “Optimal consumption under uncertainty: an experimental investigation,” *Economic Journal*, 1988, *98* (2), 105–116.
- Kluger, B. and S. Wyatt**, “Are judgement errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem,” *The Journal of Finance*, 2004, *59* (3), 969–997.
- Kogan, L., S. Ross, J. Wang, and M. Westerfield**, “The price impact and survival of irrational traders,” *The Journal of Finance*, 2006, *LXI* (1), 195–229.
- , – , – , and – , “Market selection,” *Journal of Economic Theory*, 2017, *forthcoming*.
- Miller, J. and A. Sanjurjo**, “A bridge from Monty Hall to the (anti-)hot hand: restricted choice, selection bias, and empirical practice,” 2015. IGER Working Paper. Available at SSRN: <http://ssrn.com/abstract=2709837>.
- Nalebuff, B.**, “Puzzles,” *Journal of Economic Perspectives*, 1987, *1* (1), 157–163.
- Oberlechner, T. and C. Osler**, “Survival of overconfidence in currency markets,” *Journal of Financial and Quantitative Analysis*, 2012, *47* (1), 91–113.
- Oprea, R.**, “Survival versus profit maximization in dynamic stochastic games,” *Econometrica*, 2014, *82* (6), 2225–2255.
- Roth, A. and K. Murnighan**, “Equilibrium behavior and repeated play of the prisoner’s dilemma,” *Journal of Mathematical Psychology*, 1978, *17*, 189–198.
- and **M. Malouf**, “Game theoretic models and the role of information in bargaining,” *Psychological Review*, 1979, *86*, 574–594.
- Sandroni, A.**, “Do markets favor agents able to make accurate predictions?,” *Econometrica*, 2000, *68* (6), 1303–1341.
- Slembeck, T. and J. Tyran**, “Do institutions promote rationality? An experimental test of the three-door anomaly,” *Journal of Economic Behavior and Organization*, 2004, *54*, 337–350.

A Appendix A: Omitted Proofs

Proof of Proposition 2

Using (3) and $r_0^i = r_0^j$, the time, $\tilde{T}^{ij}$, at which agent i 's wealth is at least m times the wealth of agent j is given by the smallest time for which

$$\sum_{t=1}^{\tilde{T}^{ij}} \sum_{s=0}^1 1_{st} \left(\log \frac{1-c^i}{1-c^j} + \log \frac{\alpha_s^i}{\alpha_s^j} \right) \geq \log m \quad (9)$$

For convenience, define

$$R_t^{ij} = \sum_{s=0}^1 1_{st} \left(\log \frac{1-c^i}{1-c^j} + \log \frac{\alpha_s^i}{\alpha_s^j} \right)$$

The random variables, R_t^{ij} , are independent and identically distributed, and define a random process, $\mathbf{R}^{\mathbf{j}} = \{\mathbf{R}_t^{\mathbf{j}} : t \geq 0\}$ such that $\tilde{T}^{ij}$ is a stopping time of the process. Using the same algebraic manipulations that lead to Proposition 1 in the main text,

$$E[R_t^{ij}] = SI^i - SI^j$$

which is finite under our assumptions that $c^i, c^j < 1$ and $\alpha_s^i, \alpha_s^j \in (0, 1)$. Assuming $E[\tilde{T}^{ij}] < \infty$ (which we verify below), we can then apply Wald's equation to the expectation of the left-hand side of (9)

$$E \sum_{t=1}^{\tilde{T}^{ij}} R_t^{ij} = E[\tilde{T}^{ij}] E[R_t^{ij}] = E[\tilde{T}^{ij}] (SI^i - SI^j)$$

Therefore, under the assumption that $SI^i > SI^j$, the expected time at which i is expected to have at least m times the wealth of j , is given by

$$E[\tilde{T}^{ij}] = \frac{\log m}{(SI^i - SI^j)} \quad (10)$$

which is in fact finite as assumed. (10) decreases in the difference between SI^i and SI^j so that if i is expected to have m times the wealth of the agent k that has the second highest survival index, then she is also expected to have at least m times the wealth of all other agents. Labeling

SI^k the second largest survival index, the expected time at which i has at least m times the wealth of all other agents is therefore given by

$$E \left[\tilde{T} \right] = \frac{\log m}{(SI^i - SI^k)}$$

□

Proof of Proposition 3

To allow for optimal consumption and investment plans for arbitrary risk-averse expected utility agents, we extend the model to allow consumption and investment decisions to change at every date. The model is otherwise unchanged.

With decisions at every date, the model corresponds to the case of independent and identically distributed states of the more general exchange economy studied in Blume and Easley (2006). As they show, market selection can be studied using the conditions for Pareto optimality, and the results then immediately apply for competitive, complete markets by the First Theorem of Welfare Economics.

To apply the results of Blume and Easley (2006), we assume agents are expected utility maximizers with strictly concave and monotonic utility functions satisfying the Inada condition at 0 (their Axiom 1). In Axiom 2, Blume and Easley (2006) assume that the aggregate endowment at each date is bounded away from 0: this assumption is satisfied in our model provided at least one agent places positive investment on the winning asset, which is the case in our empirical data. Finally, in Axiom 3, Blume and Easley (2006) assume that each agent puts positive probability on any realized path of states that is possible. We impose this assumption noting that, in our empirical data, agents almost always invest positive amounts in both assets in line with believing all states are possible.²⁰

Given Axioms 1-3, Proposition 3 then follows from the analysis in Section 3.1 of Blume and Easley (2006). They show that the agent i with the maximum survival index given by $\kappa_i = \log(\delta^*) - I_q(q^i)$ survives almost surely, where q^i is the agent's belief and δ^* is the agent's pseudo-discount factor. If there is a unique agent i with the least biased beliefs and the closest to rational discount factor, then their survival index is maximal. □

²⁰We observe only 3 cases in which a subject invests all of their wealth in only one asset, and believe these are simply mistakes as they did not repeat this behavior.




Figure 7: Results from END Treatment With Learning.

B Appendix B: Data from Latter Periods in END Treatment

As in the CORE sessions, we allow END sessions to continue even after subjects had received sufficient feedback to learn away the bias, ending the formal END treatment. However, software glitches late in these sessions led a number of subjects' decisions to be mis-submitted (the glitch caused subjects to submit extreme values of c - 0 or 1 - and α), generating price effects (observed by other subjects) that likely provided misleading feedback on others' beliefs in several cases. These errors also caused sometimes severe software lags and a great deal of subject frustration and boredom, with later parts of sessions sometimes lasting considerably longer than intended. For these reasons, we dropped this component of the experiment from the main paper.

Here we report the results and show that in fact the outcomes look much like those from the LEARN and END treatment. Figure 7 shows the results in the same format we use for the other treatments. As in the other treatments, survival indices tend to be highest close to Bayesian beliefs of 0.4, though results are noisier than in the CORE treatment. The Negative Counterfactual continues to imply a systematic failure of market selection (with unbiased agents being driven out of the market), showing, again, how strongly results are driven by the joint distributions of α and c . As in END and LEARN, little evidence exists that near-Bayesian subjects consume much less

than the treatment as a whole. However, the beliefs of high survival index subjects tend to be substantially closer to Bayesian than the treatment overall.

C Appendix C: Motivating the CORE Design

The CORE treatment – our main treatment – is designed to satisfy the following five goals. In this appendix we discuss these goals and how they map into the design we pursued.

Our first goal is to implement an environment in which most subjects form biased beliefs concerning states of the world. We chose a variation of the Monty Hall problem, an inference problem that has inspired a particularly large bias in previous experiments (see, for example, Nalebuff (1987); Friedman (1998); Slembeck and Tyran (2004); Kluger and Wyatt (2004)). Subjects are shown five gray cups and told that there is a coin under one of them (each cup chosen with equal probability). In the first step of the procedure, the computer uniformly randomly selects two cups to be the “green cups.” In the second step, the computer selects two of the three remaining gray cups that do not hide a coin and lifts them up. In the final step, the remaining cup (the one that has not been lifted and was not selected to be green) is the “blue” cup. Subjects must evaluate which color cup the coin is under in order to infer the state (state 0 occurs if the cup is green, and state 1 if it is blue). A Bayesian agent in this task forms a belief that state 0 occurs with probability 0.4, but the vast majority of subjects in previous work use a heuristic – we call it the “naive heuristic” – that causes them to attach probability 0.67 to state 0 occurring.²¹ The bias that this task generates arises due to a failure of conditional reasoning that is similar to the reasoning errors behind the winner’s curse, the hot hand fallacy, and failures of no trade theorems. By studying this task we are therefore studying the survival value of a type of bias with a particularly wide range of economic applications.²²

²¹The *naive heuristic* arises because subjects almost universally neglect the information contained in the computer’s decision rule to deliberately choose to lift cups that do not hide the coin, causing them to improperly update their prior and leading them to evaluate the green cup as being more likely than the blue cup (as there are twice as many green cups as blue).

²²We used a five door variation on the Monty Hall task in order to make it easy to clearly distinguish biased beliefs from the focal action of simply attaching equal weight to each state (a natural strategy for a subject who is fundamentally confused or simply not paying attention to the experiment). Failure to employ conditional probabilistic reasoning in the three door version of the problem (i.e. the naive heuristic described in footnote 21) leads an agent to assign a 0.5 probability to each state occurring. The same reasoning failure in the five door version leads an agent to believe the less-likely-to-occur state occurs with a 0.67 probability – a distinctive pattern that is unlikely to be focal for any other reason.

Second, we want to cleanly measure subjects’ beliefs in order to identify biases. We achieve this by (i) paying subjects according to a log transformation of withdrawals each round and (ii) paying in binary lottery tickets (to be realized and transformed into currency at the end of the experiment) rather than currency directly – a payment protocol we call LOG-LP. Component (ii) theoretically (at least under the premises of EU, see Roth and Malouf (1979)) incentivizes even risk-averse subjects to maximize expected earnings. Component (i) causes an expected earnings maximizing subject to “bet her beliefs”, independently of her consumption decision, a consequence of the fact that with a log utility function, agents maximize the expected growth of their wealth (see Blume and Easley (1992), Theorem 5.1). Thus, we can interpret allocation decisions, α , as directly revealing a subject’s beliefs.²³ Adopting the LOG-LP reward procedure also facilitates two additional, crucial goals of the design described below.

Third, we want to avoid aggregation and censoring problems that naturally arise in dealing with subjects with endogenous survival horizons. Specifically, subjects that adopt low-survival strategies tend to run out of wealth and are effectively censored in later rounds of the period making it difficult to compare the per-round consumption and investment decisions of low and high survival subjects in any principled way. Studying environments in which all subjects make only one decision per period, applied across all rounds of the game, solves this problem and allows for a transparent and unbiased way of comparing subjects’ choices.²⁴ An extremely useful additional feature of the LOG-LP payoff scheme described above is that it makes it optimal for subjects to set a stationary consumption rate, c , and stationary investment rule, α , across rounds. Because of this property, we can have subjects make only one choice at the beginning of each period, making aggregation and measurement straightforward and transparent.²⁵

Fourth, it is important for the experimental environment to satisfy the price-taking assumption at the heart of Blume and Easley’s (1992) general equilibrium model, without which clean interpretation of subjects’ decisions is difficult. One way to do this might be to run extremely large markets, though it is difficult to ensure *ex ante* that any finitely-sized market is in fact large enough to satisfy this central assumption of the model we are studying for our environment. An appealing

²³There is strong evidence in our data that this procedure “works” and that subjects indeed bet their beliefs. In paid practice periods, we directly tell subjects the probability that each asset “wins” and find that most subjects do allocate wealth across securities according to these probabilities (see Section 4)

²⁴In fact, we collected initial sessions in a setting in which subjects were allowed to set non-stationary choices. We redesigned the experiment precisely because we found the resulting aggregation and censoring problems to be intractable.

²⁵In addition to facilitating measurement, this procedure simplifies and speeds up the experiment, reducing subject confusion and allowing us to collect more data.

alternative is to adapt the model to an *individual decision* task with exogenous prices provided by the experimenter, making subjects unambiguous price-takers. Importantly, under the LOG-LP payoff scheme, the optimal decision rule is invariant to prices, allowing us to set stationary prices *exogenously* and induce exactly the same optimal behavior in individual subjects as in a perfectly competitive market. Setting m_{st} to a constant, exogenous level (2.1 in our implementation), we generate a price-taking environment in our CORE treatment. (In a robustness treatment, END, described below, we relax this and allow prices to emerge endogenously).²⁶

Finally, fifth, we want subjects to be quite familiar with the mechanism and decision environment when they enter the dataset so that our measures of behavior early in the experiment are not simply picking up subject confusion and inexperience. However, it is equally important that subjects do not learn away biases because (at least in the CORE treatment) we want to provide as strong of a bias as possible for the market to correct through sheer wealth dynamics. To ensure that subjects enter our dataset comfortable with the basic interface and decision environment, without becoming unbiased via learning, subjects participate in eight paid practice periods prior to the beginning of the experiment (that is prior to period 1). In these practice periods, subjects experience an environment identical to the one in the actual experiment except that states of the world are determined by a simple coin flip in place of the Monty Hall process, allowing subjects to focus their attention on understanding how savings translate into wealth accumulation and earnings.

D Appendix D: Instructions

We provide the instructions for the CORE treatment, followed by those for the END treatment, and finally the instructions for the Monty Hall problem.

²⁶ $m_{st} = 2.1$ was calibrated, along with other parameters, to maximize payoff salience in the experiment.

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

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

In the experiment you will make decisions over several periods. At the end of the last period, you will be paid the sum of your earnings over all periods, in cash.

The Basic Idea.

Each period will be divided into some randomly determined number of rounds. At the beginning of each period (in the first round) you will be given 40 **Francs** and will decide what percentage to:

- **Withdraw** each round
- **Invest** in the blue and green stocks each round

Your **earnings** are determined entirely by the **Francs you withdraw** in each round, but you can only make withdrawals in future rounds by **investing some of your Francs** now. If you choose to withdraw all of your Francs, you would only have Francs to withdraw from in the first round. However if you chose to withdraw none of your Francs, you would earn no money for the period. Regardless, when the period ends, any Francs remaining in your account will simply disappear.

Every round one of the two stocks (**blue** or **green**) will be valuable (we will call it the “**winning stock**”) and the other will be worthless. Any Francs invested in the winning stock will be increased (multiplied by a **multiplier** greater than 1) to determine the number of Francs you start the **next round** with. Any money invested in the losing stock will simply disappear. Thus the more Francs you invest in the winning stock this round, the more Francs you have to withdraw from and invest with in the next round and so on.

Decision Bar.

Figure 1 shows your computer display. In the middle is the **decision bar** which shows the percentage of your beginning Francs you choose to withdraw (shown in red), invest in green (shown in green) or invest in blue (shown in blue) each round. The

- greater the fraction of the decision bar you fill with **red**, the greater the fraction of your Francs you will withdraw each round of the period,
- the greater the fraction filled with **blue** the greater the fraction you will invest in the blue stock each round, and
- the greater the fraction filled with **green** the greater the fraction you will invest in the green stock each round.

You can control the composition of the decision bar using the input boxes below it (not shown in the figure). The left hand box controls your withdrawal amount. The right hand box controls how you invest the remaining (un-withdrawn) Francs across the two stocks. Next to the decision bar are exact percentages showing the total fraction you will withdraw each round and the fraction of the remaining Francs (e.g. those not withdrawn) you will invest in green vs. blue. After making your decision, you can finalize it by clicking the **submit button** (also not shown).




Figure 1: Computer display.

Given the way the experiment is set up, you can expect to **maximize your earnings** by distributing the Francs you choose not to withdraw, according to the true probability that each stock is the winner. So, for example, if the blue stock is the winner 50% of the time, you can expect (on average) to maximize your earnings by investing 50% of the Francs you don't withdraw in the blue stock and 50% in the green stock. If the blue stock is instead the winner 75% of the time, you can expect (on average) to maximize earnings by investing 75% in blue and 25% in green, etc.

Buying Tickets

When you make a withdrawal in a round, it will be used automatically by the computer to buy **lottery tickets**.

The Francs you withdraw will automatically purchase tickets according to the **curved function** shown in Figure 2. A graph to the left of the decision bar on your screen also shows you, as you adjust the decision bar, how many tickets various levels of withdrawals will buy you given your starting Francs. This figure will rescale every round given the number of Francs you started the round with.

Because of the curved function, doubling the amount you withdraw will not double the number of tickets you earn. Instead, the number of tickets you earn increases by a fixed amount each time the amount you withdraw doubles. Specifically, doubling your withdrawal will always increase your tickets by 0.7. For example, consider the points highlighted in Figure 2. If you withdraw 1 franc, you earn 0 tickets, if you withdraw 2 Francs, you earn about 0.7 tickets, if you withdraw 4 Francs, you earn about 1.4 tickets, and so on.

Francs Graph

After everyone has made their withdrawal and investment decisions and clicked the **submit button**, the period will begin and the rounds will advance automatically. The **Francs Graph** (on the right side of the screen) shows what happens in each round. It plots (from left to right) a horizontal black line for each round of the Period so far: the higher the line is, the more Francs you started the round with. Below each line is




Figure 2: Ticket Conversion Chart.

a miniature version of the decision bar showing you how these Francs were allocated, given the choices you made at the beginning of the period. Specifically this bar will show you (1) the amount you withdrew (just under the horizontal line, in red) and (2) the amount you invested in the **winning stock** (at the bottom in whatever color stock won) each round. Between is a blank space representing the amount of money you invested in the losing stock, to remind you that any Francs invested in the losing stock disappear!

The Francs Graph will also visualize for you how your winning investment translates to next round's Francs. At the end of each round you will see an arrow pointing from the amount you invested in the winning stock pointing to the round's beginning Francs and a caption telling you the **multiplier** (i.e. if the multiplier is 2.1, it will say "2.1x") This is to remind you that your beginning Francs **next round** are simply equal to the multiplier times the Francs you invested in the winning stock **this round**

Finally, the Francs graph will show, in red, the number of tickets your withdrawals bought you each round as a number floating above the round's initial wealth level.

Winners and Losers

Every round the computer will determine which stock – the green or the blue – is the winner for the round. In the first part of the experiment, the computer will determine this randomly in the following way.

The lower right hand side of the screen will show you four **overturned cups** each round. Some of the cups will be green and some will be blue. Each round, the computer will randomly determine one of the cups to put a **coin under** (each cup has an equal likelihood of hiding the coin each round). At the end of the round the computer will lift the cups. Whichever color cup the coin is under will be the **color of the winning stock** – if the coin is under a green cup, the green stock will be the winner while if it is under a blue cup the blue stock will be the winner. Thus, if half of the cups are blue and half green, there is a 50% chance each of the blue and green stock being the winner. If three of the cups are green there is a 75% chance of the green stock being the winner and a 25% of the blue stock of the winner. If three of the cups are instead blue there is a 75% chance of blue winning and 25% of green, and so on.

Later on we will pause the experiment, give you new instructions and change the way the winner is determined.

The Multiplier

The multiplier will always be equal to 2.1. Thus, any francs you invest in the winning stock **this round** will always be multiplied by 2.1 to determine the number of francs you start the **next round** with.

Periods and Rounds

The experiment will be divided up into a number of periods. Each period will be divided up into a number of rounds. At the beginning of each period (in round 1) you will be given 40 Francs to divide among withdrawals and investment in stocks. You will then play a number of rounds that will be **randomly** determined by the computer. Specifically each round there is a **10% chance** that the computer will determine that it is the **last round** of the period. Because each round is the last round with probability 10%, periods will last an average of 10 rounds. But, periods can be much shorter or much longer.

When a period ends, any Francs you have accumulated, but have **not withdrawn** to buy tickets, disappear and **become worthless**. Your earnings for the period will be determined by the total number of tickets you have purchased in the period. You will then start a new period with 40 Francs, and it will last a new randomly determined number of rounds.

Your Earnings

Your earnings each period will be determined by the lottery tickets you purchase with your withdrawals. Each lottery ticket gives you a 1 in 35 chance of winning **\$4.00** for the period. So, if you earn 10 lottery tickets in a period, you can expect to make $4 \times \frac{10}{35} = \1.14 for that period, on average. If you earn 25 tickets instead, you expect to earn $4 \times \frac{25}{35} = \2.86 on average. Note that partial lottery tickets count and that if the total number of tickets you earn is negative, you will earn zero.

Lotteries for each period will be conducted at the end of the experiment by the computer, once all periods are complete. In addition to your lottery earnings, you will be paid the \$5.00 show-up fee.

Reducing Randomness in Tickets

In the actual experiment, we will make a slight change to the payments from what is described above. Every round, the computer will randomly determine the period length and the winning stocks and show it to you – and how it impacts your earnings – on your screen to give you a sense of how the game works and how your decisions translated to earnings. However, in the background, the computer will randomly determine the length of the period thousands of more times for each period you play, calculating your tickets based on your withdrawals and decisions for that period each time. Your lottery tickets will be based on the **average lottery tickets** purchased over these thousands of period length determinations. We do this simply to tie your earnings more closely to your decisions (and less to random chance).

Thank you for your participation!

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

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

In the experiment you will make decisions over several periods. At the end of the last period, you will be paid the sum of your earnings over all periods, in cash.

The Basic Idea.

Each period will be divided into some randomly determined number of rounds. At the beginning of each period (in the first round) you will be given 40 **Francs** and will decide what percentage to:

- **Withdraw** each round
- **Invest** in the blue and green stocks each round

Your **earnings** are determined entirely by the **Francs you withdraw** in each round, but you can only make withdrawals in future rounds by **investing some of your Francs** now. If you choose to withdraw all of your Francs, you would only have Francs to withdraw from in the first round. However if you chose to withdraw none of your Francs, you would earn no money for the period. Regardless, when the period ends, any Francs remaining in your account will simply disappear.

Every round one of the two stocks (**blue** or **green**) will be valuable (we will call it the “**winning stock**”) and the other will be worthless. Any Francs invested in the winning stock will be increased (multiplied by a **multiplier** greater than 1) to determine the number of Francs you start the **next round** with. Any money invested in the losing stock will simply disappear. Thus the more Francs you invest in the winning stock this round, the more Francs you have to withdraw from and invest with in the next round and so on.

Decision Bar.

Figure 1 shows your computer display. In the middle is the **decision bar** which shows the percentage of your beginning Francs you choose to withdraw (shown in red), invest in green (shown in green) or invest in blue (shown in blue) each round. The

- greater the fraction of the decision bar you fill with **red**, the greater the fraction of your Francs you will withdraw each round of the period,
- the greater the fraction filled with **blue** the greater the fraction you will invest in the blue stock each round, and
- the greater the fraction filled with **green** the greater the fraction you will invest in the green stock each round.

You can control the composition of the decision bar using the input boxes below it (not shown in the figure). The left hand box controls your withdrawal amount. The right hand box controls how you invest the remaining (un-withdrawn) Francs across the two stocks. Next to the decision bar are exact percentages showing the total fraction you will withdraw each round and the fraction of the remaining Francs (e.g. those not withdrawn) you will invest in green vs. blue. After making your decision, you can finalize it by clicking the **submit button** (also not shown).




Figure 1: Computer display.

Given the way the experiment is set up, you can expect to **maximize your earnings** by distributing the Francs you choose not to withdraw, according to the true probability that each stock is the winner. So, for example, if the blue stock is the winner 50% of the time, you can expect (on average) to maximize your earnings by investing 50% of the Francs you don't withdraw in the blue stock and 50% in the green stock. If the blue stock is instead the winner 75% of the time, you can expect (on average) to maximize earnings by investing 75% in blue and 25% in green, etc.

Buying Tickets

When you make a withdrawal in a round, it will be used automatically by the computer to buy **lottery tickets**.

The Francs you withdraw will automatically purchase tickets according to the **curved function** shown in Figure 2. A graph to the left of the decision bar on your screen also shows you, as you adjust the decision bar, how many tickets various levels of withdrawals will buy you given your starting Francs. This figure will rescale every round given the number of Francs you started the round with.

Because of the curved function, doubling the amount you withdraw will not double the number of tickets you earn. Instead, the number of tickets you earn increases by a fixed amount each time the amount you withdraw doubles. Specifically, doubling your withdrawal will always increase your tickets by 0.7. For example, consider the points highlighted in Figure 2. If you withdraw 1 franc, you earn 0 tickets, if you withdraw 2 Francs, you earn about 0.7 tickets, if you withdraw 4 Francs, you earn about 1.4 tickets, and so on.

Francs Graph

After everyone has made their withdrawal and investment decisions and clicked the **submit button**, the period will begin and the rounds will advance automatically. The **Francs Graph** (on the right side of the screen) shows what happens in each round. It plots (from left to right) a horizontal black line for each round of the Period so far: the higher the line is, the more Francs you started the round with. Below each line is




Figure 2: Ticket Conversion Chart.

a miniature version of the decision bar showing you how these Francs were allocated, given the choices you made at the beginning of the period. Specifically this bar will show you (1) the amount you withdrew (just under the horizontal line, in red) and (2) the amount you invested in the **winning stock** (at the bottom in whatever color stock won) each round. Between is a blank space representing the amount of money you invested in the losing stock, to remind you that any Francs invested in the losing stock disappear!

The Francs Graph will also visualize for you how your winning investment translates to next round's Francs. At the end of each round you will see an arrow pointing from the amount you invested in the winning stock pointing to the round's beginning Francs and a caption telling you the **multiplier** (i.e. if the multiplier is 2.1, it will say "2.1x") This is to remind you that your beginning Francs **next round** are simply equal to the multiplier times the Francs you invested in the winning stock **this round**

Finally, the Francs graph will show, in red, the number of tickets your withdrawals bought you each round as a number floating above the round's initial wealth level.

Winners and Losers

Every round the computer will determine which stock – the green or the blue – is the winner for the round. In the first part of the experiment, the computer will determine this randomly in the following way.

The lower right hand side of the screen will show you four **overturned cups** each round. Some of the cups will be green and some will be blue. Each round, the computer will randomly determine one of the cups to put a **coin under** (each cup has an equal likelihood of hiding the coin each round). At the end of the round the computer will lift the cups. Whichever color cup the coin is under will be the **color of the winning stock** – if the coin is under a green cup, the green stock will be the winner while if it is under a blue cup the blue stock will be the winner. Thus, if half of the cups are blue and half green, there is a 50% chance each of the blue and green stock being the winner. If three of the cups are green there is a 75% chance of the green stock being the winner and a 25% of the blue stock of the winner. If three of the cups are instead blue there is a 75% chance of blue winning and 25% of green, and so on.

Later on we will pause the experiment, give you new instructions and change the way the winner is determined.

The Market and the Multiplier

You will be randomly matched up with 7 other participants in this room at the beginning of the experiment to form a **market**. You will remain matched with this group for the entire experiment. Other members of your market will only influence the **multiplier** – the amount your investment in the winning stock will be multiplied by to determine the number of francs you start **next round** with.

Every round we will give you feedback on the overall market's decisions by calculating the **fraction** of the total francs the market invested in the green and blue stocks. We will show you these fractions (as percentages) in the Total Market Summary box in the middle of the screen. These fractions tell you something about the overall market's beliefs about the relative likelihood of the green and blue stocks being the winner.

These fractions will determine the **multiplier** each round. The multiplier is simply 1 divided by the fraction of francs invested in the winning stock that round. For example, if 50% of francs overall are invested in the winning stock this round, the multiplier would be $1/0.5=2$. If 75% were invested in the winning stock, the multiplier would instead be $1/0.75=1.33$ and if 25% were invested in the winning stock, the multiplier would be $1/0.25=4$. Thus **the less the market invests** in the winning stock overall, **the more your investment** in the winning stock is multiplied by in determining your next round's starting francs.

Periods and Rounds

The experiment will be divided up into a number of periods. Each period will be divided up into a number of rounds. At the beginning of each period (in round 1) you will be given 40 Francs to divide among withdrawals and investment in stocks. You will then play a number of rounds that will be **randomly** determined by the computer. Specifically each round there is a **10% chance** that the computer will determine that it is the **last round** of the period. Because each round is the last round with probability 10%, periods will last an average of 10 rounds. But, periods can be much shorter or much longer.

When a period ends, any Francs you have accumulated, but have **not withdrawn** to buy tickets, disappear and **become worthless**. Your earnings for the period will be determined by the total number of tickets you have purchased in the period. You will then start a new period with 40 Francs, and it will last a new randomly determined number of rounds.

Your Earnings

Your earnings each period will be determined by the lottery tickets you purchase with your withdrawals. Each lottery ticket gives you a 1 in 35 chance of winning **\$4.00** for the period. So, if you earn 10 lottery tickets in a period, you can expect to make $4 \times \frac{10}{35} = \1.14 for that period, on average. If you earn 25 tickets instead, you expect to earn $4 \times \frac{25}{35} = \2.86 on average. Note that partial lottery tickets count and that if the total number of tickets you earn is negative, you will earn zero.

Lotteries for each period will be conducted at the end of the experiment by the computer, once all periods are complete. In addition to your lottery earnings, you will be paid the \$5.00 show-up fee.

Reducing Randomness in Tickets

In the actual experiment, we will make a slight change to the payments from what is described above. Every round, the computer will randomly determine the period length and the winning stocks and show it to you – and how it impacts your earnings – on your screen to give you a sense of how the game works and how your decisions translated to earnings. However, in the background, the computer will randomly determine the length of the period thousands of more times for each period you play, calculating your tickets based on your withdrawals and decisions for that period each time. Your lottery tickets will be based on the **average lottery tickets** purchased over these thousands of period length determinations. We do this simply to tie your earnings more closely to your decisions (and less to random chance).

Thank you for your participation!

Instructions

In the remaining periods of the experiment, everything will be **exactly the same except** we will change how the **winning stock** is determined.

Each round, the screen will show you **five overturned cups**. As in the first part of the experiment, the computer will randomly select one of these cups to hide a coin under. Each cup is equally likely to be selected to hide the coin.


Each round, the computer will first randomly select two of the cups to be the **green cups**. Each cup is equally likely to be selected as a green cup in any given round.

Two cups randomly colored GREEN.




Next, the computer will lift two remaining (non-green) cups that do not have a coin under them.

Show two of remaining cups WITHOUT coin.




Whichever cup is not green and has not been lifted will be the **blue cup**.

Last remaining cup is colored BLUE.



Whichever color cup the coin is under will be the **color of the winning stock** – if the coin is under a green cup, the green stock will be the winner while if it is under the blue cup the blue stock will be the winner. In the example below, the coin was under a green cup and so the green stock is the winner this round.



We will use this procedure for the remaining periods in the experiment. Please think carefully about your decisions as your earnings will depend a lot on the investment decisions you make.