

SURVIVAL VERSUS PROFIT MAXIMIZATION IN A DYNAMIC STOCHASTIC EXPERIMENT

BY RYAN OPREA¹

Subjects in a laboratory experiment withdraw earnings from a cash reserve evolving according to an arithmetic Brownian motion in near-continuous time. Aggressive withdrawal policies expose subjects to risk of bankruptcy, but the policy that maximizes expected earnings need not maximize the odds of survival. When profit maximization is consistent with high rates of survival (HS parameters), subjects adjust decisively towards the optimum. When survival and profit maximization are sharply at odds (LS parameters), subjects persistently (and sub-optimally) hoard excess cash in an evident effort to improve survival rates. The design ensures that this hoarding is not due to standard risk aversion. Analysis of period-to-period adjustments in strategies suggests instead that hoarding is due to a widespread bias towards survival in the subject population. Robustness treatments varying feedback, parameters, and framing fail to eliminate the bias.

KEYWORDS: Economic survival, dynamic stochastic decision making, economics experiments.

“Learning is not compulsory. . . neither is survival.”

W. Edwards Deming

1. INTRODUCTION

THE RELATIONSHIP BETWEEN PROFIT MAXIMIZATION AND SURVIVAL has long held a central role in economics. An influential and long standing conjecture (sometimes called the Market Selection Hypothesis²) holds that deviations from profit maximization cannot long survive in the field because profit maximizing firms must also be the most likely to survive. This idea—the classical justification for the profit maximization assumption in economics—is deeply intuitive, but it is often wrong.

In a provocative paper, Dutta and Radner (1999) pointed out that strategies that maximize profits need not maximize the odds of survival. In fact, in many

¹I am grateful to the National Science Foundation for support under Grant SES-0925039, to seminar audiences at UC Berkeley, Florida State University, the University of Michigan, New York University, Purdue University, UC San Diego, UC Santa Barbara, Stanford University, and attendants at the 2011 North American ESA Meetings. I thank Douglas Bernheim, Andrew Caplin, Tim Cason, Gary Charness, David Cooper, Sean Crockett, Guillaume Frechette, Daniel Friedman, Uri Gneezy, Sebastian Goerg, Aspen Gorry, Michael Kuhn, Muriel Niederle, Roy Radner, Jonathan Robinson, Andrew Schotter, Vernon Smith, Charles Sprenger, Richard Thaler, and Nat Wilcox for helpful comments and useful discussion, James Pettit and Jeremy Hewett for developing the software, and Keith Henwood, Chad Kendall, and Jacopo Magnani for outstanding research assistance. All mistakes are my own.

²See, for example, Alchian (1950) and Friedman (1953).

realistic settings, agents can significantly improve their survival odds by deviating systematically from wealth maximizing strategies. For example, a firm's cash serves both as the reservoir out of which it pays profits and as a buffer against bankrupting runs of bad luck. A (discounted, expected) profit maximizing firm must regularly pay out cash as profits, but doing so necessarily erodes its defenses against bankruptcy, exposing it to a potentially severe hazard of going bankrupt.³ On the other hand, a firm biased towards survival can improve its odds simply by hoarding excess cash, losing expected discounted profits in the bargain. Dutta and Radner showed that biases towards survival, if they exist, can short circuit the forces of market selection, allowing serious deviations from profit maximization to survive in the real world.⁴

Do such biases towards survival actually exist? There are reasons to suspect they might. Economists long overestimated the linkage between survival and profit maximization, and it stands to reason that economic agents might be prone to the same mistake. Given the deep-rooted importance attached to survival in most spheres of life, it seems plausible that people may persistently interpret failures to survive as failures to optimize, even when this interpretation is highly misleading. Firms often do hoard cash especially when cash flows are volatile (Han and Qui (2007), Bates, Kahle, and Stulz (2000)), though it is difficult to tell to what degree this hoarding is driven by excessive concern for survival. Key variables necessary for benchmarking optimal behavior are unobservable in the field, and decision problems are far more complex than in our models. In order to shed some empirical light on this question, we conduct a controlled laboratory experiment directly testing Dutta and Radner's model.

In our experiment, subjects choose when to withdraw earnings from a cash reserve that evolves according to an arithmetic Brownian motion in near-continuous time. If the reserve ever falls to zero, the subject goes "bankrupt" and is unable to make further withdrawals. Dutta and Radner (1999) showed that an expected profit maximizing agent will withdraw all cash overflowing a threshold level τ^* (determined by parameters of the model) and withdraw nothing when cash reserves are lower.⁵ In our Core treatments, subjects directly choose such a threshold, τ .

In order to test for bias towards survival, we study two sets of Brownian parameters that vary the degree to which survival and profit maximization conflict. Under High Survival (HS) parameters, profit maximizing thresholds are

³Although we focus on a setting in which firms have access only to retained cash, the logic sketched here and in the model we implement easily extends to firms with access to credit. In this case, bankruptcy occurs not when cash reaches zero, but when it reaches some negative threshold value.

⁴A distinct theoretical literature studies the very different issue of the survival of mistaken beliefs in financial markets (see, e.g., Blume and Easley (1992, 2002, 2006) and Sandroni (2000)). This literature shows (among other things) that the Market Selection Hypothesis can fail in this setting if markets are incomplete. Blume and Easley (2008) provided a useful survey.

⁵Shubik and Thompson (1959) and Radner and Shepp (1996) showed that such overflow policies are also optimal in related environments.

consistent with near-certain survival (96 percent). Under Low Survival (LS) parameters, optimality and expected survival rates are in sharp conflict: optimizing subjects face a survival rate of only 15 percent. In order to improve survival odds in either treatment, a subject must hold a larger than optimal reserve of cash by setting a higher than optimal threshold, $\tau > \tau^*$. However, since subjects nearly always survive at the optimum under HS parameters, even subjects biased towards survival have little reason to deviate from the optimum. Survival biased subjects under LS parameters, by contrast, have strong reasons to deviate. Thus, under the hypothesis that subjects have a bias towards survival, we expect to see substantially higher than optimal thresholds under LS parameters but not under HS parameters.

This is just what we observe in our Core treatments. Thresholds under HS parameters move decisively towards the optimum and settle near the optimal level. Under LS parameters, this learning process fails and subjects persistently set thresholds far higher than optimal ones, hoarding large reserves of excess cash. Evidence from period-to-period adjustments in thresholds suggests that most subjects have some degree of bias towards survival in both HS and LS treatments. This analysis suggests that differences in survival rates (not differences in rates of bias) across the treatments drive our main treatment effect: without low survival rates, we would expect LS subjects to behave much like HS subjects.

Critically, our design ensures that hoarding cannot be an outgrowth of standard risk aversion over earnings. In fact, hoarding actually exposes subjects to *greater* earnings risk each period (see Appendix A of the Supplemental Material (Oprea (2014))). Risk aversion over profits will induce optimizing subjects to withdraw more cash than is optimal, not less, increasing the likelihood of bankruptcy. Indeed, an infinitely risk averse subject in our experiment will put no weight on survival, liquidating her account at the first moment.

We interpret our results instead as an outgrowth of an inferential error: subjects have trouble disentangling the signals provided by survival and earnings when they conflict. A deeply ingrained (and usually reliable) heuristic towards survival leads subjects to associate survival with optimality, leading to large and persistent failures in environments where the two goals are in conflict. This misapplied heuristic is surprisingly resistant to learning, continuing to guide behavior even after dozens of periods of feedback, experience, and even social learning (applied in a robustness treatment).

Our results suggest this error is persistent and broadly distributed, but how well should we expect it to survive corrective forces in the economy? A second component of the Market Selection Hypothesis suggests that investors will be systematically drawn to profit maximizing firms, starving biased firms of resources and driving them from the market. This mechanism, though familiar in economics, hinges on a critical assumption: that investors will not suffer from the very same biases when evaluating investment options. Our final set of treatments (the Investment treatments) evaluates by having subjects

repeatedly choose whether to link their earnings to accounts (“firms”) programmed to hold optimal levels of cash or alternative accounts programmed to persistently hoard a suboptimal, excess cash buffer. Because subjects have no control over the survival prospects of either firm, we argue that survival is considerably less salient here than in the Core treatments, and it is therefore even more natural to focus attention on profits in making decisions. Nonetheless, we find strong evidence that survival bias persistently influences investment choices. Under HS parameters, subjects systematically invest in profit maximizing firms, as we might expect, but under LS parameters (where profits and survival rates conflict), subjects make little distinction between profit maximizers and hoarders: hoarding firms are nearly as likely to receive investment funds as their profit maximizing competitors.

Although we observe evidence of survival bias in several distinct environments, it is important to emphasize that this bias is only capable of actually causing deviations from optimality under a special set of economic circumstances. It is only when cash flows are very low and/or highly volatile (as with our LS parameters) that profit maximization and high survival rates conflict, motivating survival biased agents to abandon profit maximization. Economies in the grips of stagnation, upheaval, or crisis thus seem especially likely to suffer the ill effects of the bias.

Likewise, it is important to emphasize that our experiment only studies a small sampling of the many decision environments in which profit maximization and survival might clash. We chose our experimental environments to facilitate clean comparison with economic theory and to generate crisp and interpretable evidence of the effect of survival rates on optimization. Though these are natural diagnostic settings and generate clean evidence of survival bias, there are many institutional features of rich, real-world markets that may temper (or intensify) the bias. We consider some possibilities in Section 6.

To our knowledge, our work is the first to identify biases towards survival in the laboratory. In a very distantly related paper, Schotter, Weiss, and Zapater (1996) examined whether a frame of economic survival can alter decisions in bargaining games. Oprea (2008) studied principal–agent problems and corporate takeover in dynamic cash management games that resemble ours in some respects. More generally, our paper contributes to a relatively small literature on dynamic-stochastic decision problems that has focused so far primarily on savings problems (e.g., Ballinger, Palumbo, and Wilcox (2003), Brown, Chua, and Camerer (2009), Hey and Dardanoni (1988)) and investment problems (e.g., Oprea, Friedman, and Anderson (2009)).

The remainder of the paper is organized as follows. Section 2 introduces a simple case of Dutta and Radner’s (1999) model, discusses the relationship between survival and profit maximization, and describes the discrete approximation we take to the lab. Section 3 describes our experimental design and implementation. Section 4 presents results from the Core treatments and robustness results from our Social Learning treatment. Section 5 reports results

from our Investment treatments. In Section 6, we discuss the bias, consider its robustness, and offer a plausible psychological interpretation, and in Section 7, we conclude the paper.

2. MODEL, SURVIVAL IMPLICATIONS, AND DISCRETIZATION

Consider an entrepreneur in control of a firm with a cash reserve, $Y(t)$. At any moment, t , the firm's cash is given by

$$(1) \quad Y(t) = Y(0) + X(t) - W(t),$$

where $Y(0)$ is the firm's initial level of cash and $X(t)$ is the firm's net inflow of cash by time t . The cash inflow begins at time $t = 0$ and evolves according to an arithmetic Brownian motion:

$$(2) \quad dX = \mu dt + \sigma dz,$$

where z is a standard Weiner process, $\mu > 0$ is the cash flow's drift, and σ is its volatility. The entrepreneur sets a withdrawal policy $W(t)$, a description of the cumulative amount of cash withdrawn from the reserve by time t . We assume $W(t) \leq Y(0) + X(t)$ (the entrepreneur cannot withdraw more cash than has accumulated in her account).

The entrepreneur is impatient and discounts the future at a rate δ over an infinite horizon.

The entrepreneur faces an important constraint on her withdrawals: if $Y(t)$ ever reaches zero, the firm is bankrupt, meaning $Y(t)$ will never again rise above zero, and the entrepreneur will make no future withdrawals.

The firm's problem is to maximize the discounted expected total withdrawals given by

$$(3) \quad \pi = E \left\{ \int_0^T e^{\delta t} dW(t) \right\},$$

where T is the (stochastic) time at which the firm is finally bankrupt. In order to maximize withdrawals, the entrepreneur must balance the immediate consumption value of cash with its insurance value against bankruptcy.⁶ Withdrawing cash now directly increases profits, but it also increases the probability that a run of bad luck will bankrupt the firm, doing away with future profit opportunities.

Dutta and Radner showed that the value function $V^*(Y(t))$ is strictly concave over $[0, \tau^*]$ for some threshold τ^* , has a slope $V'(\tau^*) = 1$, and linearly increases (with slope less than 1) at $Y(t) > \tau^*$. A withdrawal at time t of

⁶As we emphasize in Appendix A of the Supplemental Material, though cash insures against bankruptcy risk, it does not insure against earnings risk in our experiment.

size w yields a gain of w , a loss of $wV'(Y(t))$, and therefore a net gain of $w(1 - V'(Y(t)))$. The optimal withdrawal is thus 0 if $Y(t) < \tau^*$ (since $V'(Y(t))$ is >1) and $y - \tau^*$ if $y > \tau^*$ (since $V'(Y(t))$ is <1). We will call this optimal withdrawal pattern an *overflow policy*: the firm sets a stationary threshold τ , withdraws nothing when $Y(t) < \tau$, and withdraws $Y(t) - \tau$ whenever $Y(t) > \tau$ (subjects directly choose τ in the experiment). The optimal level of this threshold, τ^* , is determined entirely by the parameters of the cash process and the discount rate.⁷ Defining θ and λ as the absolute values of the positive and negative roots of the expression $\sigma x^2 + 2\mu x - 2\delta = 0$, Dutta and Radner showed that the expected profit maximizing threshold level is

$$(4) \quad \tau^* = \frac{(\lambda/\theta)^2}{\lambda + \theta}.$$

Moreover, given an initial level of cash $Y(0) > \tau^*$ (always true in the parameterizations used in our experiment), the expected discounted returns are given by $\frac{\mu}{\delta} - \tau^*$. We refer the reader to Dutta and Radner (1999) for proofs and additional details.

2.1. Survival and Profit Maximization

The choice of τ affects not only expected profits but also the survival prospects of the firm. Because τ is a reflecting barrier while 0 is an absorbing one, cash will repeatedly fall from τ towards 0 in a series of downward cycles. In our experimental design, we reinterpret δ as the instantaneous probability that the game ends at each moment in time, transforming an infinite horizon problem into a theoretically isomorphic indefinite horizon problem. In this setting, the lower the value of τ , the more likely one of these runs will reach the absorbing barrier of zero and bankrupt the firm before the game stochastically ends. If a decision maker manages to avoid bankruptcy during the game, we will say she has *survived*. The decision maker's *survival rate* as a function of τ can be expressed using the same positive and negative roots used above:

$$(5) \quad s = 1 - \frac{(\lambda + \theta)e^{\tau(\theta - \lambda)}}{\lambda e^{-\tau\lambda} + \theta e^{\tau\theta}}.$$

This expression is simply the probability that the expiration time of the game—an exponentially distributed random variable with parameter δ —occurs before y reaches the absorbing state of 0.

An important feature of the profit maximizing threshold, τ^* , is that it need not maximize this survival rate, s . Indeed, the firm's expected survival rate at

⁷The profit maximizing threshold is both stationary in time and independent of the current level of cash, Y . An optimizing agent will therefore hold a stationary τ throughout the period.

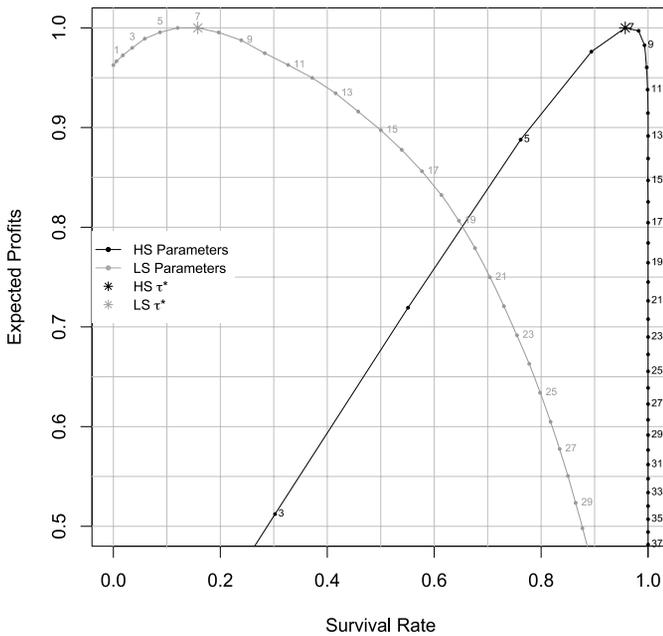


FIGURE 1.—Curves plot the tradeoff between expected profits and survival rates for LS and HS parameters. Numbers along the curve denote thresholds generating each point on the curve.

the profit maximizing optimum can vary dramatically across parameters, a fact we exploit in our design. However, for any set of parameters, the entrepreneur can reduce expected profits and (weakly) increase the likelihood and duration of survival simply by raising her threshold to $\tau^* + \varepsilon$, hoarding excess cash in the firm as insurance against ruin.

Figure 1 illustrates the relationship between survival rates and expected profits by plotting the tradeoff between the two (over threshold choices) for two distinct sets of parameters. Under each set of parameters, profits are maximized at the exact same threshold choice ($\tau^* = 7$), but in each case, very different survival rates result from this profit maximizing choice. Under High Survival (HS) parameters, the firm is virtually guaranteed to survive after choosing τ^* (the survival rate is over 95%), while under Low Survival (LS) parameters, survival will be rare (only 15% of the time) at τ^* .

A profit maximizing agent will make the exact same choice under each set of parameters ($\tau = \tau^* = 7$). However, an agent motivated not only by high profits but also by high survival rates—we call such agents *survival biased*—will make very different choices under each. Under HS parameters, τ^* simultaneously (nearly) maximizes both survival rates and expected profits. An agent with mixed motives over the two types of outcomes will therefore have little reason to deviate much from τ^* : holding a threshold more than a few points

higher than τ^* (rightward along the curve) will cause profits to plummet without yielding any benefits to survival. However, the same agent under LS parameters will have a strong reason to hold a threshold higher than the optimum: doing so (moving rightward along the curve) will reduce profits but will reward this loss with large increases in the survival rate. The higher the weight the agent places on survival relative to expected profits, the larger this upward deviation from τ^* will be.

Our experimental design (described in more detail below) and empirical strategy are rooted in the comparison of behavior under these HS and LS parameters. Under the hypothesis that subjects are survival biased, we expect to see nearly profit maximizing threshold choices under HS parameters but a substantial upward deviation (i.e., substantial hoarding of cash) under LS parameters.

2.2. Discrete Implementation

Continuous time Brownian motion is an idealization unsuitable for the laboratory, so we implement near-continuous time binomial approximations (e.g., Dixit (1993)) that closely mirror the look, feel, and distributional qualities of Brownian motion.⁸ Time is divided into a number of discrete “ticks,” each lasting Δt seconds, and in each tick the cumulative cash flow $X(t)$ increments up or down by a fixed amount $h > 0$. The direction of the cash movements is stochastic: with probability $p > 0.5$ cash increases to $X(t) + h$, and with probability $1 - p$ decreases to $X(t) - h$.

The salient feature of Brownian motion is that over a time interval of length g , the change in $X(t)$, $\Delta X(t)$, is a random variable distributed $N(\mu g, \sigma^2 g)$. In our discrete approximation, over any time interval g , the change in $X(t)$ is instead binomially distributed. However, by the Central Limit Theorem, as Δt approaches zero, the distribution of $\Delta X(t)$ approaches the normal distribution and the underlying stochastic process approaches a Brownian motion. Thus, by setting Δt very low (1/5 of a second in our sessions), we can very closely approximate a Brownian motion in the lab.

Limiting Brownian parameters can be recovered from binomial parameters. The drift parameter μ is

$$(6) \quad \mu = \lim_{\Delta t \rightarrow 0} \frac{(2p - 1)h}{\Delta t},$$

while volatility is

$$(7) \quad \sigma^2 = \lim_{\Delta t \rightarrow 0} \frac{4p(1 - p)h^2}{\Delta t}.$$

⁸Indeed, simulations confirm that optimal decisions in discrete time binomial implementations closely mirror predictions generated by their limiting Brownian analogues for all of the parameters reported in this paper.

(Note that h and p are fixed and do not vary with Δt .)

Dutta and Radner’s (1999) model features an infinite horizon that is an impractical fit to the lab for obvious reasons. As discussed above, we interpret the impatience parameter, δ , instead, as an instantaneous hazard of the game ending at each moment, transforming the infinite horizon problem into an equivalent indefinite horizon problem. In our binomial approximation, this is implemented, with a fixed probability q each tick that the tick is the last (e.g., Roth and Murnighan (1978)). As with the other key parameters of the model, we can express the relationship between our discrete discount factor, q , and its continuous time analogue, δ . If expiration has a probability Q per unit of time, then the expected value of 1 unit of payoff to be paid one time unit in the future is $1 - Q$. With $\eta = 1/\Delta t$ time steps per unit of time, and expiration probability q per time step, the discount factor is $e^{-\delta} = 1 - Q = (1 - q)^\eta = (1 - q)^{1/\Delta t}$. Solving for δ , we obtain

$$(8) \quad \delta = \frac{-\ln(1 - q)}{\Delta t}.$$

3. DESIGN AND IMPLEMENTATION

Our experimental design is built around the contrast between the HS (High Survival) and LS (Low Survival) parameters introduced in Section 2.1. Table I lists binomial parameters, the wealth maximizing threshold level τ^* , and the expected survival rate, s^* , at τ^* for each treatment. Central to the design is the fact that the profit maximizing threshold, τ^* , is *identical* under both sets of parameters but the rates of survival at this threshold, s^* , differ significantly. After setting threshold τ^* , subjects will nearly always survive under HS parameters, while subjects setting the same threshold under LS parameters will rarely survive. Thus in HS treatments, profit maximization and survival are mutually consistent, while in LS treatments, they are in conflict. As explained in Section 2.1, under the hypothesis that subjects are biased towards survival, we will observe thresholds near the optimum in HS environments but much higher than the optimum in LS environments.

TABLE I
BINOMIAL PARAMETER VALUES, OPTIMAL THRESHOLDS, AND OPTIMAL SURVIVAL RATES BY PARAMETER SET^a

| Parameterization k | Step Size h | Uptick Prob. p | Expiration Prob. q | Threshold τ^* | Survival Rate s^* |
|-------------------------|------------------|---------------------|-------------------------|-----------------------|------------------------|
| HS | 0.5 | 0.59 | 0.0029 | 7 | 0.96 |
| LS | 1 | 0.5113 | 0.0029 | 7 | 0.15 |

^aIn all sessions, the time step in seconds is $\Delta t = 1/5$.

Our Core treatments—C-LS and C-HS, implemented under LS and HS parameters, respectively—are direct tests of the decision problem discussed in Section 2 and form the centerpiece of the paper. Subjects in these treatments are endowed with an initial amount of cash and directly choose a threshold, τ , that will determine both withdrawals (profits) and survival rates during the game. Subjects are assigned exclusively to either the C-LS or C-HS treatment (the design is entirely between subjects), and by comparing threshold choices across the two treatments, we conduct a sharp test of the existence of survival bias. Here, the HS treatment acts as a direct control for any non-survival causes for deviation from τ^* . A robustness treatment, SL-LS, replicates the C-LS treatment but adds scope for social learning by sharing decisions and outcomes of other subjects with each participant (we discuss this treatment in more detail in Section 4.3). Finally, in our Investment treatments (I-LS and I-HS), subjects act not as entrepreneurs setting their own thresholds, but instead as investors choosing between exogenously programmed accounts. Once again, the contrast between LS and HS parameters (implemented in I-LS and I-HS treatments, respectively) allows us to gauge whether survival bias influences behavior. We provide details and motivation for this treatment in Section 5.

Appendix C of the Supplemental Material describes designs and results for seven additional robustness treatments including five additional sets of parameters and two additional decision environments not studied in the main body of the paper. These results strongly resemble the results of the main design reported in the paper.

3.1. *Laboratory Implementation*

We conducted our main sessions using undergraduate students⁹ at the EBEL laboratory at the University of California, Santa Barbara. No subject was allowed to participate in more than one treatment or experience more than one set of parameters. During each session, subjects sat at visually isolated terminals and made their decisions using a custom piece of software written in Javascript using the Redwood toolkit (Pettit, Hewitt, and Oprea (2014)). In order to ensure independence across decision makers, subjects did not interact with one another during the experiment and were not informed of one another's decisions (except in the Social Learning treatment, described in Section 4.3).

Each session is composed of a series of 40 periods and each period is an independent implementation of the complete dynamic decision problem described in Section 2 (or a variation on it in the case of the Investment treatments). Each period, j , is divided into a series of N_j ticks, each lasting $1/5$ of a second. Forty values of N (the period length) were randomly and independently

⁹Subjects were drawn from an undergraduate subject pool using students from across the curriculum, recruited using ORSEE (Greiner (2004)).

drawn before data collection began (from a geometric distribution governed by parameter $q = 0.0029$) and were assigned to subjects in an independent random order.¹⁰ Thus while subjects experience the same set of period lengths, they experience them in different orders. At the beginning of each period, each subject is endowed with $Y(0) = 40$ units of cash.

In the Core (and Social Learning) treatments, subjects directly choose their threshold, τ , at the beginning of each period and then observe their choice's impact on cash, profits, and survival as the period stochastically unfolds. Constraining subjects to stationary overflow policies (of the sort prescribed by the theory) focuses the task on the size of the buffer established against bankruptcy, while preventing artificial confounding factors like activity bias, boredom avoidance,¹¹ or failure to fully account for the memorylessness of the geometric distribution¹² from influencing outcomes.¹³

Figure 2 shows the computer display subjects used in the Core treatments. The x -axis represents the current time in ticks and the y -axis the level of cash, Y . The current value of $Y(t)$ is represented as a dot, while past values are drawn in blue and move to the left over time as on a ticker tape. On the right side of the screen is a horizontal dotted line—the “decision barrier”—used to choose a threshold for the period. Whenever the current level of cash is above the decision barrier, money is immediately withdrawn to the level of the barrier and the withdrawn amount is represented as a vertical green line. At the beginning of the period, the subject points, clicks, and drags on the screen to adjust the barrier to her preferred level.¹⁴ After setting a threshold, subjects click on a button and the cash level starts its stochastic evolution, driven by the treatment's binomial parameters. If the cash level ever reaches zero, the subject is bankrupt, the line turns red, and the cash stays at zero for the remainder

¹⁰Of course, while subjects are aware of the per-tick period ending hazard, q , they only learn the realized period length once the period ends.

¹¹After going bankrupt, subjects must wait until the period randomly ends before starting the next period. If subjects can adjust their decisions throughout the period prior to bankruptcy, they may avoid going bankrupt simply to maintain the ability to adjust their decisions. This potential confounding motive disappears if subjects can make only one decision at the beginning of the period regardless of their decisions. We thank the editor for pointing this confound out.

¹²Failure to understand the memorylessness of the geometric distribution governing period lengths is the clearest reason a subject might avoid setting a stationary policy as prescribed by the theory. Since this error is unrelated to survival bias (and is artificial to the laboratory implementation), eliminating it by allowing only a single threshold per period reduces a potential source of noise (and perhaps bias) and improves the interpretability of the results.

¹³Appendix C of the Supplemental Material reports results from a series of “Flexible Threshold” robustness treatments that relax the requirement that subjects set stationary overflow policies each period; subjects in these treatments tend to choose relatively stationary overflow-like policies on their own, and the results otherwise broadly resemble those from the Core treatments.

¹⁴In order to avoid influencing decisions, subjects are shown no barrier in the first period and must click in white space to cause the barrier to appear. In subsequent periods, subjects are shown their barrier from the previous period and (quoting the instructions) “can click and drag to adjust the barrier either up or down.”

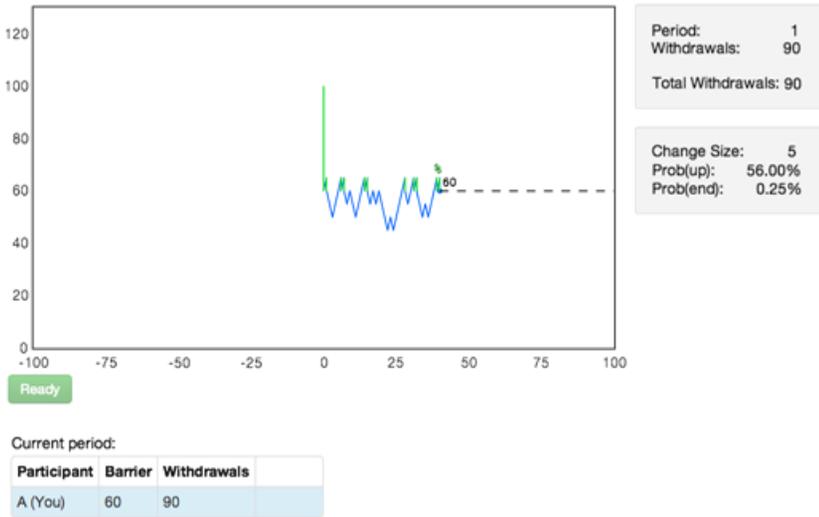


FIGURE 2.—Screenshot from the Core treatments.

of the period (until the random ending time, N_j , is reached).¹⁵ The screen also includes information on withdrawals and a complete list of the binomial parameters for the treatment on the right hand side. The software was designed to provide subjects with complete feedback on the consequences of their decisions during and between periods.¹⁶ (In the other treatments—SL-LS, I-LS, and I-HS—the screen is modified in ways we describe when we introduce the treatments below.)

We ran 13 sessions, with between 10 and 14 subjects in each.¹⁷ In total, 157 subjects are represented in the data set; 35 in C-HS, 35 in C-LS, 33 in SL-LS, 26 in I-LS, and 28 in I-HS.¹⁸ In order to create strong payoff salience, each pe-

¹⁵It is important to reiterate that subjects who go bankrupt retain any withdrawals made prior to bankruptcy. The only effect of bankruptcy is to prevent subjects from making future withdrawals during the period. We emphasized this point to subjects several times during the instructions.

¹⁶During the period, subjects receive real-time notice of their cash, withdrawals, and survival status as described above. Between periods, subjects are also shown a summary of the earnings and survival consequences of the previous period's threshold choice in a table at the bottom of the screen.

¹⁷In order to sharpen the central test of the paper, C-HS and C-LS treatments were run in joint sessions (though no individual subjects participated in more than one treatment) and each ordering of period lengths was duplicated and assigned to both a C-HS and C-LS subject. These two steps ensure that session effects and differential sequences of period lengths have no influence over our central results.

¹⁸Another 193 subjects are represented in the robustness treatments reported in Appendix C of the Supplemental Material.

riod's withdrawals were transformed¹⁹ into cents using the following formula: $\max(\pi - 30, 0)\psi$, where $\psi = 1.275$ under HS parameters and $\psi = 5$ under LS parameters (this difference in scaling factor roughly equalizes expected earnings at the optimum for each parameter set). The average session lasted roughly 90 minutes and the average subject earned (including a \$5 showup payment) \$22.61.

4. MAIN RESULTS

In Section 4.1, we test our main hypothesis by comparing Core behavior under HS and LS parameters. We show that subjects tend to converge near the optimum under HS parameters but persistently set thresholds nearly three times too large in the LS treatment, indicating a significant bias towards survival. In Section 4.2, we use individual level regressions to measure the bias and we show that it is prevalent in our sample under both HS and LS parameters. A counterfactual exercise using these bias measurements suggests that survival bias is a driver of our treatment level findings. Section 4.3 examines the implications of learning for interpreting our findings and reports the results of the Social Learning treatment, which triples the amount of feedback provided to subjects but does little to alleviate the bias.

4.1. *Hoarding and Survival*

Figure 3 plots time series of median thresholds in the Core treatments, giving us a view of the aggregate dynamics. When profit maximization guarantees near-certain survival (the C-HS treatment), thresholds drop immediately and decisively towards the optimum. By contrast, when profit maximization and survival sharply conflict (the C-LS treatment), thresholds fall slowly and soon stall out at a level nearly three times too high. By the final 10% of the session, the median LS subject sets a threshold nearly twice as large as the median HS subject and hoards four times as much excess cash. A Kolmogorov–Smirnov test on subject-wise median thresholds allows us to decisively reject the hypothesis that subjects set the same thresholds in the two treatments ($p < 0.001$).

RESULT 1: While C-HS subjects learn to set nearly optimal thresholds, C-LS subjects persistently hoard excess cash by setting thresholds far above the optimal level. By the end, this hoarding is four times more severe for the median C-LS subject than for the median C-HS subject.

Aggregates are useful but unavoidably conceal important heterogeneity in individual behavior. To provide a deeper view of the data, we construct a simple

¹⁹The payoff transformation is designed to increase penalties from deviating from the optimum. This makes biased behavior particularly costly and therefore allows us to conduct a more stringent test of survival bias.

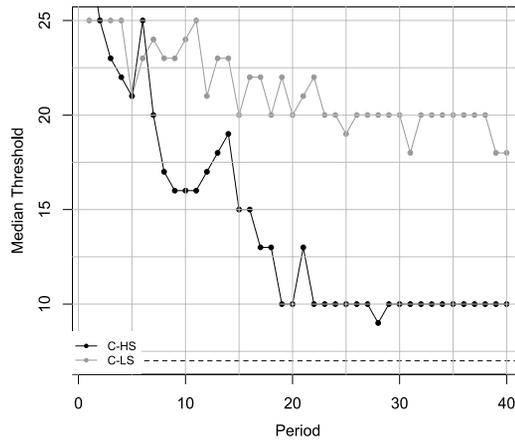


FIGURE 3.—Median thresholds in the Core treatments.

taxonomy of subject types using behavior in the final 10% of the experiment. We classify a subject as “Near-Optimal” if she sets a median threshold within 5 points of the optimum, an “Anti-Hoarder” if she sets a median threshold at least five points too low, and a “Hoarder” if she sets a median threshold at least five points too high. Subjects who set thresholds at least 10 points too high (sacrificing at least 15% of expected profits) are “Extreme Hoarders.”

Rates across subjects for the C-HS and C-LS treatments are shown in Table II. Most C-HS subjects (nearly 75%) converge in the neighborhood of the optimum, but fewer than 20% do so in the C-LS treatment. Conversely, while nearly 75% of C-LS subjects can be classified as Hoarders, only 25% of subjects can be in the C-HS treatment. Perhaps most strikingly, C-LS subjects are more than twice as likely to be Extreme Hoarders as C-HS subjects. All three of these treatment differences in rates are significant by a Fisher Exact test at the 1% level.²⁰

TABLE II
SUBJECT CLASSIFICATIONS (RATES ACROSS SUBJECTS) BASED ON MEDIAN THRESHOLDS IN THE FINAL 10% OF PERIODS

| Treatment | Hoarder | Near-Optimal | Extreme Hoarder | Anti-Hoarder |
|-----------|---------|--------------|-----------------|--------------|
| C-HS | 0.26 | 0.74 | 0.23 | 0.00 |
| C-LS | 0.74 | 0.20 | 0.57 | 0.06 |

²⁰Subjects show little tendency to hold lower-than-optimal thresholds under either parameter set: rates of Anti-Hoarding are nearly zero under both HS and LS parameters.

RESULT 2: By the end of the session, LS subjects are three times as likely as HS subjects to hoard cash and less than a third as likely to converge near the optimum.

What consequences result from this pattern of behavior?²¹ By hoarding cash, the median LS subject nearly quadruples her survival rate, raising it from 15% (the optimal rate) to nearly 60%. This massive improvement in survival comes at a severe cost to profits: the median subject sacrifices over 25% of profits (relative to counterfactual optimal behavior) in the LS treatment. By contrast, HS subjects have virtually no room to improve upon the near-perfect survival rate obtaining at the optimum. With no motive to hoard, the median subject sets a threshold near the optimum, earns 95% of optimal profits, and virtually always survives.

4.2. *Survival Bias*

Why do subjects hoard cash under LS parameters? Standard risk aversion over wealth seems at first an obvious explanation, but as we explain in Appendix A of the Supplemental Material, risk aversion should inspire exactly the opposite behavior in our experiment: in order to shield cash from Brownian runs of bad luck and expiration hazard, a risk averse subject should actually set a lower-than-optimal threshold, decreasing her odds of survival but simultaneously lowering her earnings risk.²² Since almost no subjects make systematic errors in this direction, risk aversion is not a viable explanation for the behavior we observe (nor does it explain the wide differential in rates of hoarding across the two parameter sets). Other preference-based explanations for deviations from wealth maximization like loss aversion and prospect theory seem no better suited to the data.²³

The data instead point to *survival bias*: subjects mistakenly take survival rates into account (above and beyond their impact on profits) when assessing the optimality of their choices. Figure 1, above, illustrates and highlights the in-

²¹The following measurements are obtained using data from the final 10% of the session. We measure optimal profits directly for each period by calculating counterfactual earnings at τ^* using the exact realizations of the binomial process experienced by subjects in the experiment. These calculations are only available for periods in which $\tau \geq \tau^*$ (because bankruptcy censors observation of binomial realizations for $\tau < \tau^*$). Since subjects almost exclusively deviate from the optimum by setting too-high thresholds, 90% of the data meet this criterion, and this is the sample we use to calculate earnings rates.

²²An infinitely risk averse agent will simply liquidate her firm at the first instant by setting a threshold of zero. Under LS parameters, doing this comes with very modest losses to expected earnings; the fact that C-LS subjects rarely pursue this strategy is telling.

²³All changes to earnings are in the positive domain in our experiment—once subjects increase earnings by making withdrawals, these earnings can never be reduced. Standard versions of loss aversion require potential reductions in earnings and therefore cannot apply to our data. Likewise, prospect theory simply predicts risk aversion here because earnings changes are restricted to the positive domain.

ferential logic of our experimental design. In the C-LS treatment, a subject can substantially increase her survival rate (at the expense of expected profits) by increasing her threshold (moving rightward along the curve) near the optimum. The C-HS treatment completely eliminates this motive by generating near-perfect survival rates near the optimum. Survival biased subjects should thus hoard (hold a too-high threshold) under LS parameters and hold a near-optimal threshold under HS parameters. This is just what we see. In the C-LS treatment, the median subject's choice indicates a willingness to sacrifice over 20 percent of expected profits in order to raise his survival rate from 15% to 60%. Together, the treatment level results suggest that subjects' decisions are in fact guided by both survival (subjects deviate strongly from the optimum only when survival conflicts with profits) and profits (subjects could survive even more often than they do in the C-LS treatment by sacrificing even more expected profits).

To better understand this treatment-level finding, we examine how individual subjects' experience of bankruptcy influences their threshold adjustments from period to period. By studying how threshold changes are impacted by bankruptcy events, we access a crucial window into the degree to which survival concerns shape subjects' attempts to optimize their threshold levels over time. Consider the following simple, reduced form specification:

$$(9) \quad \tau_{it} - \tau^* = \alpha + \kappa[\tau_{it-1} - \tau^*] + \beta \text{bank}_{it-1} + \nu_1 t + \varepsilon_{it},$$

where τ_{it} is subject i 's threshold in period t , bank_{t-1} is an indicator variable taking a value of 1 if subject i went bankrupt in period $t-1$, and ε_{it} is a normally distributed disturbance term.

Table III reports the results in column (1). κ , the coefficient on the deviation from the optimum, is estimated significantly below 1 ($p < 0.001$), indicating that subjects tend to adjust in the direction of the optimum over time. However, α is also positive, suggesting that, as in the aggregate data, this adjustment process terminates prior to fully reaching the optimum (though ν , the control for trend, is negative, suggesting that bias in adjustments declines to some degree with experience). Specification (2) adds a control for earnings, centered on subject-wise mean earnings— $\tilde{\pi}$ —and shows that subjects also pursue more aggressive withdrawals (set lower thresholds) after earning lower-than-average earnings.²⁴ Under the hypothesis that subjects are not biased towards survival, we would expect to see no effect of bankruptcy in these sorts of specifications: because the incidence of bankruptcy contains no information about the optimality of behavior in the previous period beyond that summarized by the threshold itself, we should expect it to be estimated at zero. The data, however, allow us to roundly reject this hypothesis: the bankruptcy dummy

²⁴Interacting $\tilde{\pi}$ with a dummy for $\tilde{\pi} > 0$ indicates that $\tilde{\pi}$ influences behavior only when negative. That is, subjects tend to set lower thresholds after underperforming but (sensibly) do not tend to set higher thresholds after earning higher-than-average withdrawals.

TABLE III
 REGRESSIONS EXPLAINING THE DEVIATION FROM
 THE OPTIMUM, $\tau_{it} - t^*$, AS A FUNCTION OF THE
 PREVIOUS PERIOD DEVIATION ($\tau_{it-1} - \tau^*$), A DUMMY
 FOR BANKRUPTCY ($bank_{it-1}$), PERIODS ELAPSED (t),
 AND EARNINGS CENTERED ON SUBJECT-WISE
 AVERAGE EARNINGS ($\tilde{\pi}_{it-1} \equiv \pi_{it-1} - \bar{\pi}_i$)^a

| Variable | (1) | (2) |
|------------------------|---------------------|---------------------|
| (Intercept) | 1.696*** (0.512) | 1.670*** (0.504) |
| $\tau_{it-1} - \tau^*$ | 0.839*** (0.027) | 0.845*** (0.026) |
| $bank_{it-1}$ | 1.848*** (0.515) | 1.784*** (0.511) |
| t | -0.020** (0.010) | -0.022** (0.010) |
| $\tilde{\pi}_{it-1}$ | | 0.007*** (0.002) |

^aStandard errors, clustered at the subject level, are shown in parentheses and one, two and three stars signify significance at the ten, five and one percent levels.

is highly significant and positive in both specifications. Holding both thresholds and earnings constant, the estimates indicate that bias in the adjustment process is over twice as large after bankruptcy than after survival, suggesting that thresholds should settle further away from the optimum in environments of high bankruptcy than in environments of high survival.

Importantly, these estimates do not differ across treatment: interacting an indicator variable for treatment with the dependent variables reveals that none, including the bankruptcy coefficient, β (a useful reduced form estimate of a subject's degree of survival bias), differs in a statistically significant way. Figure 4 plots CDFs of β estimates from *individual level* estimates of (9) for each treatment.²⁵ The results suggest that bias is widespread in both treatments: most subjects are measured with a positive value of β in each case. Moreover, the exercise confirms that the distributions of β parameters are identical across the C-HS and C-LS treatments ($p = 0.247$, Kolmogorov-Smirnov test): subjects appear to be equally biased under both sets of parameters.

²⁵In the HS treatment, a number of subjects actually never experience bankruptcy and they are of necessity dropped from this figure. It is important to emphasize that any selection bias introduced by dropping these subjects should actually work against our results, decreasing the degree of bias measured in HS subjects. This is because severely survival biased subjects are also the most likely to never experience bankruptcy and are thus the ones most likely to be dropped from the sample.

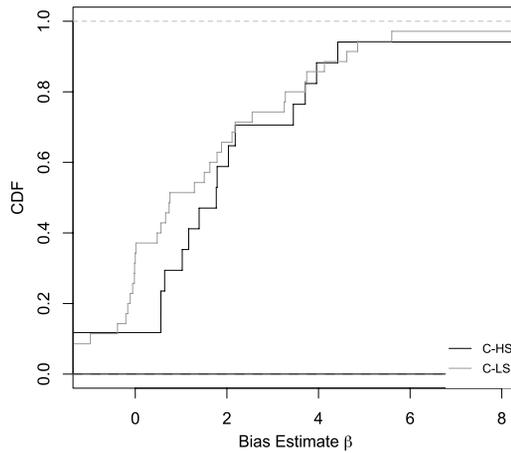


FIGURE 4.—CDFs of survival bias estimates β across subjects in the Core treatments.

RESULT 3: Controlling for previous threshold choices and earnings, subjects set significantly higher thresholds after going bankrupt than after surviving. This bias is measured positive in the majority of subjects and is equally present under HS and LS parameters.

These reduced form estimates suggest survival bias is no more severe in C-LS subjects than in C-HS subjects, yet C-LS subjects set thresholds far higher than the near-optimal ones set by C-HS subjects. To understand why, consider the following simple, back of the envelope counterfactual exercise using a variation of (9) in which all covariates are interacted with a treatment dummy. We fit the model with $\tau_{t-1} = \tau_{HS} \equiv 10$, the threshold level to which the median C-HS subject converges, and ask whether C-LS subjects are predicted to adjust differently in the following period than C-HS subjects. As long as we assume identical bankruptcy/survival outcomes in forming each treatment's fit,²⁶ we cannot reject the hypothesis that C-LS subjects set the same threshold as C-HS subjects in the following period, regardless of whether we impose the assumption that subjects of both type went bankrupt ($p = 0.863$) or survived ($p = 0.229$) in the previous period. Importantly, however, both of these assumptions are clearly counterfactual, as survival rates in fact differ strongly between the two treatments: by design, at a threshold of 10, most C-HS subjects will survive while most C-LS subjects setting the same threshold will go bankrupt. If we adjust the exercise to reflect this feature of the design by comparing a C-HS fit under the assumption of survival (i.e., setting $\beta = 0$) to a C-LS fit under the assumption of bankruptcy (i.e., $\beta = 1$), we predict that C-LS subjects set significantly

²⁶That is, set $\beta = 0$ (indicating survival) for both C-HS and C-LS fits or $\beta = 1$ (indicating bankruptcy) for both C-HS and C-LS fits.

higher thresholds than C-HS subjects ($p = 0.013$), generating a divergence in behavior as in the aggregate data.²⁷ The results therefore suggest that differences in the rates of survival are central to explaining why C-LS thresholds do not set the same thresholds as C-HS subjects. If survival rates were similar across treatments, this exercise would predict similar patterns of adjustment given the similarity of estimates for each treatment. It is because these rates are, instead, dramatically different that statistically similar parameter estimates lead to significantly different behaviors across the two treatments.²⁸

4.3. *Learning and Social Learning*

Our experimental design includes forty periods, each a complete instance of the motivating decision problem. This design choice was intended not to match the learning environment present in naturally occurring markets—few entrepreneurs, after all, have the experience of helming dozens of firms in sequence—but instead to test the robustness of the bias to experience and sophistication. Subjects in unfamiliar environments—especially environments this complex—may make systematic mistakes that reflect not a robust inferential bias but instead a transitory and instinctual initial response to an unfamiliar decision task. After dozens of periods of experience, the risk that systematic errors are due to artificial unfamiliarity with the decision problem is considerably lessened: subjects have a developed sense for the stochastic processes governing outcomes and a firm understanding of the rules of the game (as evidenced by convergence near the optimum in C-HS). Systematic deviations from optimality that survive this level of experience are arguably more telling of the behavior of sophisticated decision makers in the field.

Perhaps the most striking fact about the bias identified in the Core treatments is its robustness to this sort of experience. Even after dozens of repetitions and exhaustive feedback on the consequences of decisions, C-LS subjects

²⁷Similar results obtain across a broad range of previous-period threshold levels. For instance, we get similar results if we fit the models at $\tau_{t-1} = \tau_{LS} \equiv 20$ (the threshold C-LS subjects converge to in the last half of periods) or even at τ^* , the optimum. In all cases, survival outcomes must be different to predict significant differences in subsequent adjustments across treatments.

²⁸Indeed, pushing this exercise a bit further and fitting the model at period 40 (to focus on late behavior), we cannot reject the hypothesis that subjects from either treatment holding a threshold of $\tau_{HS} \equiv 10$ would set τ_{HS} again after surviving (C-HS: $p = 0.167$; C-LS: $p = 0.115$), but we can easily reject the same hypothesis after bankruptcy (C-HS: $p = 0.006$; C-LS: $p = 0.001$) where significant positive fits predict upward adjustment away from τ_{HS} . The results are reversed at the typical low survival threshold $\tau_{LS} \equiv 20$: we cannot reject that subjects from either treatment would again set τ_{LS} after bankruptcy (C-HS: $p = 0.193$; C-LS: $p = 0.624$), but can reject the same hypothesis after survival (C-HS: $p = 0.030$; C-LS: $p = 0.002$) where negative fits predict downward adjustment towards τ_{HS} . The results suggest that high rates of bankruptcy tend to destabilize τ_{HS} and stabilize τ_{LS} , while low rates of bankruptcy do the reverse. That these results hold for both C-HS and C-LS fits lends further support to the idea that C-HS and C-LS behaviors would be similar under similar bankruptcy rates and are different in large part because of differences in the rate of survival.

persist in hoarding a substantial amount of cash in an effort to survive and sacrifice significant earnings by doing so. Data from the C-HS treatment provide a natural control, showing that in the absence of the confounding signal sent by low survival rates, subjects can learn to optimize in an environment this complex with a bit of experience. Together, the treatments suggest that survival bias persists long after subjects have acquired enough experience to learn how to optimize.

In order to push the learning opportunities built into the design even further, we designed a Social Learning treatment—SL-LS—in which subjects are provided feedback on the decisions and outcomes of other participants. At the beginning of each SL-LS session, subjects are randomly assigned to a three player cohort and play a game identical to the Core game save one modification: a table at the bottom of subjects' screens shows, in real time, the thresholds, withdrawals, and bankruptcy statuses of all members of the cohort so far this period.²⁹ Subjects are therefore provided three times the feedback available in the Core treatments. The SL treatment mirrors knowledge transmission that plausibly occurs in some competitive markets and provides a second, distinct environment in which to study the tension between survival and profit maximization.

We ran the Social Learning treatment using the LS parameters and procedures identical to the ones used in the Core treatments on a total of 35 undergraduate students at the EBEL laboratory at UC Santa Barbara. Our primary hypothesis was that the additional feedback provided in this environment and (perhaps) virtuous informal competition over earnings between cohort members would cause subjects to overcome survival bias and converge to thresholds near the optimum.

Surprisingly, we find no effect of social learning on optimization or hoarding rates. In the final 10% of periods, we classify only 18% of subjects as Near-Optimal but 74% of subjects as Hoarders and 48% as Extreme Hoarders using the taxonomy described in Section 4.1. These numbers are virtually identical to those reported for the C-LS treatment and none of these measures can be distinguished from C-LS analogues using Fisher Exact tests. There is a modest reduction in the median SL-LS threshold relative to C-LS (16.5 vs. 19), but the two distributions are statistically indistinguishable ($p = 0.2149$, Mann-Whitney test).

RESULT 4: Social learning is no more effective than individual learning at eliminating survival bias.

Appendix C of the Supplemental Material reports the results of another robustness treatment (the “Multiple Accounts” treatment) that increases the

²⁹Subjects are also shown the decisions and outcomes for all cohort members from the previous period while setting their thresholds at the beginning of each period.

feedback to subjects even more directly. Multiple Accounts subjects make a single threshold choice but this choice is simultaneously applied to four independently evolving accounts each period, each governed by parameters with a low optimal survival rate. This treatment quadruples the level of feedback to subjects and simultaneously reduces the saliency of survival (although survival rates for individual accounts are low, profit maximizing subjects can expect at least one of their accounts to survive 80% of the time). However, this treatment has no significant effect on hoarding, either: subjects persistently set thresholds far above the optimum just as in all of the other low survival treatments we have studied.

5. INVESTMENT AND SURVIVAL BIAS

Even if entrepreneurs suffer from survival bias (as our subjects do), it is far from certain that others would ever invest resources in them. Indeed, a second component of the Market Selection Hypothesis posits that investors will concentrate capital on profit maximizing firms, eliminating biased firms by systematically denying them resources. This claim is intuitive but depends on a critical assumption: that investors are immune to the same biases that lead non-profit maximizing firms astray in the first place. Do investors confuse survival with optimality when choosing between investments, as our Core subjects do when setting their own thresholds?

In our Investment treatments (I-LS and I-HS), subjects are repeatedly given the option of investing in either (i) an account governed by a nearly optimal threshold (the OPTIMAL “firm”) or (ii) an alternative account governed by a much higher than optimal threshold (the HOARDING firm).³⁰ At the beginning of each of 40 periods, subjects are shown two side-by-side accounts, each with a different exogenously assigned threshold, and click a button to invest in one. After making a choice, subjects observe the independent evolution of cash, withdrawals, and survival status for each firm over the course of the period.³¹ When the period ends, cash is reset for each firm (at $Y(0) = 40$), new thresholds are assigned, and the subject makes another investment decision in the following period.

Thresholds assigned to firms are drawn from actual decisions made by C-LS subjects in corresponding periods.³² Subjects thus directly make investment

³⁰Of course, we do not label the firms “OPTIMAL” or “HOARDING” in the actual experiment. Instead, we call them “Firm L” and “Firm R” and randomize which is the OPTIMAL and which is the HOARDING firm from subject to subject.

³¹The plot generated for each account during the period and the timing of the game is exactly as in the Core treatments. The main difference in the display is that subjects observe two independent plots evolving simultaneously during the period instead of one.

³²Specifically, for each period we selected from C-LS choices in the corresponding period (i) the threshold closest to the optimum and (ii) the threshold closest to the sample median, and assigned these to the OPTIMAL and HOARDING firms, respectively. These selection rules

choices over decisions made by previous subjects. In order to approximate the setting of a competitive market, subjects' investment choices have no effect on either of the firms' thresholds or outcomes. Instead, investment choices affect only which firm's withdrawals determine the investor's earnings: Investment subjects earn 1 point for each point of withdrawals made by the target of her investment.³³ We ran 54 subjects through the Investment treatment—26 under LS parameters (the I-LS treatment) and 28 under HS parameters (the I-HS treatment)—at the EBEL laboratory at UC Santa Barbara. Instructions are reproduced in Appendix B of the Supplemental Material.

We designed the Investment environment under the hypothesis that by moving survival status to external parties (i.e., the investor herself never goes bankrupt—only the investment targets go bankrupt) and eliminating subjects' control over whether either firm survives, subjects (even under LS parameters) would cease to be influenced by survival concerns and learn quickly to invest exclusively in the profit maximizing OPTIMAL firm. This outcome—call it the “Selection Hypothesis”—would support the investment component of the Market Selection Hypothesis, as it would suggest that investors are systematically drawn to profit maximizing firms and away from biased ones.

If investors are instead biased towards survival when making investment choices, one of two patterns will emerge under LS parameters. Overwhelmingly strong focus on survival relative to profits will guide subjects to systematically invest in HOARDING firms, generating a perverse inversion of the Market Selection Hypothesis in which profit maximizing OPTIMAL firms are systematically rejected by investors (the “Malselection Hypothesis”). A more likely consequence of the bias (in light of results from Section 4) is that conflicting pursuit of survival and profits will simply lead to confusion over which firm is optimal, preventing investors from systematically investing in *either* type of firm. In this case, we would expect subjects to hold mixed portfolios (constructed over the course of periods) of OPTIMAL and HOARDING firms, funding a significant number of each (the “Non-Selection Hypothesis”). Here, again, HS parameters provide a natural control: we expect even survival biased investors to systematically choose the OPTIMAL firm under HS parameters.

5.1. Results

As expected, subjects in the I-HS treatment reliably invest in the OPTIMAL firm—84% of the time overall and over 90% of the time by the last half of the session. By contrast, subjects in I-LS invest in the OPTIMAL firm only 55% of the time overall and 58% of the time during the last half of the session.

are never revealed to subjects. All subjects in both Investment treatments experience the exact same sequence of OPTIMAL and HOARDING firm thresholds.

³³Points are transformed into dollar payments using exactly the same formulas used in corresponding Core treatments.

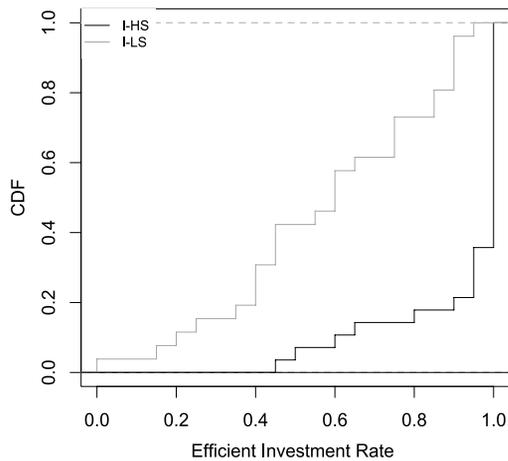


FIGURE 5.—CDFs of subject-level rates of investment in the OPTIMAL firm, by treatment.

Figure 5 plots empirical CDFs of subject-level rates of investment in the OPTIMAL firm during the last half of the session. While over 60% of I-HS subjects invest *exclusively* in the OPTIMAL firm (and over 75% do so at least 95% of the time), not even one I-LS subject concentrates entirely on the OPTIMAL firm. The data thus allow us to decisively reject both the Selection Hypothesis and the Malselection Hypothesis in the I-LS data. When there is a tension between expected profits and survival, subjects invest in the HOARDING firm nearly as frequently as the OPTIMAL firm. Indeed, for I-LS subjects, we cannot reject the hypothesis that subjects are as likely to invest in OPTIMAL firms as HOARDING firms.³⁴ The data thus tend to support the Non-Selection Hypothesis, giving us our next result:

RESULT 5: When (and only when) OPTIMAL firms suffer low rates of survival, subjects invest in suboptimal HOARDING firms nearly as often as they invest in OPTIMAL firms.

Does survival bias lie behind this failure of market selection? Once again, observing behavior under HS parameters (where survival and expected profits are mutually consistent) gives us an answer. Subjects are over 50% more likely to invest in the OPTIMAL firm in I-HS than in I-LS, and this holds true both overall and in the final half of the session. Using a Mann–Whitney test

³⁴We test this with a Wilcoxon test conducted on subject-wise rates with a null hypothesis that the rate of investment in the OPTIMAL firm is 50% ($p = 0.136$). The same test using I-HS subjects allows us to roundly reject the same null ($p < 0.001$).

applied to subject-wise rates of efficient investment, we can reject the hypothesis that subjects are equally likely to efficiently invest in the two treatments ($p < 0.001$).³⁵ Overall, subjects in the I-LS treatment sacrifice five times more money (relative to consistent efficient investment) than I-HS subjects do. This gives us a further result:

RESULT 6: I-HS subjects are significantly more likely to invest in the OPTIMAL firm than I-LS subjects, suggesting that survival bias influences investment choices.

The effect of survival bias on investors is to reduce (and nearly eliminate) subjects' ability to make distinctions between optimal and suboptimal investment options. As a result, Market Selection fails to operate: biased firms are nearly as likely to be chosen to receive "capital" by investors as profit maximizing alternatives.

6. DISCUSSION

6.1. *Interpretation*

We have operationalized and measured survival bias, but how should we interpret it, psychologically? The data, of course, do not answer this question directly, though in this section we discuss our favorite interpretation.

We view survival bias as an outgrowth of a deeply rooted heuristic rule to "avoid low survival strategies"—an important (perhaps even hard-wired) rule for any biological organism. Like all heuristics, this "Survival Seeking Heuristic" (SSH) is a reliable guide to optimal action in many (even most) spheres of life. And like all heuristics, it can lead to biased decision making when ported to the wrong domains. Although a heuristic towards survival is of primary importance for a biological entity, it is only secondarily important to an economic entity. This biologically important heuristic can therefore lead to serious biases when ported to economic decision making.

In a thematically related recent example from economics, [Charness and Levin \(2005\)](#) experimentally studied an environment in which reinforcement learning—another useful heuristic in many real-world settings—conflicts with optimal Bayesian decision making. They found that subjects tend to lean heavily on the reinforcement heuristic even though it is a poor fit to the setting, leading to persistent errors. We view survival bias as an instance of this sort

³⁵This is true despite the fact that the earnings differential between OPTIMAL and HOARDING firms is generally not greater in the I-HS than in the I-LS treatment. Indeed, over most of the distribution of period returns, the reverse is true: the earnings difference between OPTIMAL and HOARDING firms is actually greater in most I-LS periods than in most I-HS periods.

of misapplied heuristic. Because the SSH is usually a reliable guide to optimal decisions, subjects have difficulty associating low survival odds with optimality. When subjects fail to survive, they sense that they have made a mistake and adjust their cash holdings to improve survival odds in the future. As in Charness and Levin, optimal decisions simply “feel wrong” to subjects in these cases.

To expand on this notion, it is useful to think of agents’ behavior being shaped to varying degrees by both the SSH and a competing Wealth Seeking Heuristic (WSH). The Market Selection Hypothesis is built on an assumption that these two heuristics will tend to generate similar behavior: the impulse to survive will generally lead agents to seek out profit maximizing strategies and vice versa. Dutta and Radner’s contribution is to show that the two heuristics can, in fact, often generate systematically different behaviors. Our experiment effectively exogenously varies the degree to which the heuristics provide coinciding advice; under HS parameters they generally agree, while under LS parameters they provide systematically different advice. Our results suggest that subjects follow both heuristics to some degree (and we suspect that subjects systematically conflate the two goals).³⁶ The experiment shows, however, that SSH is strong enough relative to the WSH to lead to large and systematic errors when the two sharply conflict (as they only do under LS parameters).

Results from the Investment treatments support this “competing heuristics” interpretation in a particularly sharp way. I-LS subjects are drawn to both low survival/high profit investments and high survival/low profit investments, and to a nearly equal degree. Instead of choosing one type of investment over the other, virtually all subjects elect to simply hold a mixed portfolio (over the course of periods), and the average subject invests nearly as often in the option suggested by the SSH as the WSH. When we remove the influence of the SSH by applying HS parameters, this mixing stops and investment flows systematically to the profit maximizing firm.

6.2. *Robustness*

Our experiment was designed to generate a clean test for survival bias and to assess its robustness on several initial dimensions, but it leaves open a number of questions about the robustness of the bias to institutional forces and selective pressures at work in naturally occurring environments. In this subsection, we review what our results tell us about robustness, speculate on features

³⁶We hypothesize that subjects (like some early proponents of the Market Selection Hypothesis) overestimate the linkage between survival and earnings, attempting to increase expected earnings by increasing survival rates. It is equally possible that subjects mistakenly see a direct connection between survival risk and earnings risk, viewing high survival strategies as somehow “safer” from an earnings perspective. Indeed, most consumers of the research reported in this paper (the author included) have initially drawn the same mistaken conclusion!

of real-world markets that may mitigate (or exacerbate) the pattern of behavior documented here, and suggest some important directions future research might follow to evaluate robustness more fully.

Perhaps the most important direct piece of evidence on the robustness of survival bias provided by our design is its striking resistance to learning and feedback. In our Core Low Survival treatment, subjects severely hoard even after dozens of periods of repetition of the task and exhaustive feedback concerning the consequences of these decisions. Follow-up treatments explicitly designed to stress-test this finding triple (the Social Learning treatment) and quadruple (in Supplemental Material Appendix C's Multiple Accounts treatment) the amount of feedback given to subjects but fail to eliminate the severity of the bias. As we argue in Section 4.3, robustness to heavy feedback and experience suggests that the bias we have measured is not simply an artificial initial reaction by subjects to an unfamiliar environment, but is instead something more durable, perhaps with potential to survive outside of the lab. Future research should explore issues of learning further, and there are many promising directions that might be pursued. For example, future designs might increase the number of participants in social learning cohorts or dramatically increase the number of periods of repetition. Another future design strategy might be to systematically vary parameters in low survival settings in such a way as to vary the degree of noise in feedback concerning profits in order to assess how not only the volume but also the quality of feedback might influence the intensity and longevity of the bias.

Conversely, perhaps the largest questions left open by our data are to what extent the salience of survival differs between the lab and the field and to what degree variation in salience influences survival-seeking behavior. At one extreme, the salience of survival might be so intense in the field that laboratory subjects avoid bankruptcy in part because of the bundle of strong negative associations decision makers attach to bankruptcy in the external world. If our data are in part driven by this sort of analogic reasoning, we would expect the sort of aversion to bankruptcy observed in our experiment to be present also (and perhaps even be intensified) in many naturally occurring settings. One way to explore this in future research would be to modify the framing of the experiment in such a way as to weaken (or strengthen) associations with real-world firms and bankruptcy events, perhaps under the guidance of recent theory on analogy-based expectations and valuation (e.g., Jehiel (2005)).³⁷

At the opposite extreme, in some centrally important parts of the economy, survival may be hardly salient at all. For example, one consequence of decentralized ownership of firms in many developed economies is that no individual firm looms particularly large in a typical investor's portfolio. As the number of firms represented in an investor's portfolio grows, individual firms are reduced

³⁷See also Huck, Jehiel, and Rutter (2011) for supporting evidence for analogy-based expectations.

to a collection of revenue streams and this may reduce the salience of survival events for investors. Moreover, investors are often at an arm's length from (and have little control over) the life and death of individual firms, perhaps reducing the salience of any one firm's survival relative to its function as a profit source. For these reasons, it is difficult to imagine that the holder of an index fund consisting of thousands of firms is much concerned about the life and death of individual firms in its portfolio; it may even be that immunity from biases of this sort is an important (and unheralded) efficiency produced by decentralized capital markets. Our experiment takes some first steps at examining this possibility: our Investment treatment and the Supplemental Material's Multiple Accounts treatment both attempt to reduce the saliency of bankruptcy in ways inspired by mechanisms plausibly at work in financial markets. In both cases we fail to eliminate survival bias, but it seems clear that some of the ideas in these treatments can be pushed much further.

One path future researchers might take is to study environments in which subjects choose between or control considerably more "firms" than in our treatments. It seems plausible that the salience of survival only really diminishes once an appropriate scale of diffusion is reached (i.e., the antidote might be in the dose). It is also possible that the ability to focus on profitability and ignore survival is a relatively specialized ability held only by savvy investors and managers but not by the general population. If so, conducting experiments like ours on non-student populations (perhaps financial professionals) might yield different results. Of course, even if mechanisms in capital markets do exist to counteract survival bias, many firms in developed economies—and most firms in many developing economies—are privately held and have relatively little direct interaction with financial markets. In these regions of the economy we might expect survival bias to be a particularly relevant concern.

Many specific aspects of our protocol were implemented purely in order to achieve clean measurement and interpretable results. Some of the more artificial aspects of our design might be relaxed in order to improve realism in future research. For instance, in order to maximize comparability with benchmark predictions of the theory, we constrain subjects to "overflow policies" and have them directly set stationary thresholds in each repetition of the game. In our Flexible Thresholds treatment, reported in the Supplemental Material, we partially relax this feature of the protocol by allowing subjects to hold non-stationary thresholds, but many other changes to the protocol might be worth studying. For instance, variations in which subjects set constant rates of withdrawal (rather than overflow thresholds) or even discrete time variations in which subjects directly choose a withdrawal amount each period may lead to a more nuanced understanding of the relationship between profits and survival. The task we implement in our Core treatments is also a single-player decision problem, while most firms in the field are embedded in interactive competitive environments. Our Social Learning treatment relaxes the isolation of the Core treatments somewhat by informing subjects of others' decisions and outcomes.

This increase in feedback has surprisingly little effect on survival bias, but it is quite possible that more direct forms of interaction—for example, direct competition between subjects over revenues—might succeed in altering the relative salience of profits and survival and therefore the intensity of survival bias.³⁸

Finally, it is important to emphasize that even if survival seeking behavior survives in the field, we should expect it to lead to biased behavior only under a special set of economic circumstances. In economic environments with relatively high returns and low volatility, profit maximization also generates relatively high rates of survival and the heuristic we have identified therefore need not lead to serious deviations from optimality. It is only when returns are low and/or highly volatile that survival and profit maximization are thrown into conflict and survival seeking behavior leads to failures to optimize. Thus, survival bias may be of greatest concern in times of economic stagnation, financial panic, or political crisis, where poor and uncertain returns prevent profit maximizers from realizing high rates of survival.

7. CONCLUSION

Survival is an important motive in a number of settings, but is of secondary importance to a profit maximizing firm. In fact, in many environments, high rates of survival can only be achieved by deviating seriously from optimality. We report the results of a laboratory experiment studying how subjects manage dynamic cash flows in environments with bankruptcy risk. Our central treatment variable varies whether optimality is consistent with high survival rates. When wealth maximization and high rates of survival are mutually consistent (our HS parameters), subjects show strong and immediate tendencies towards optimal behavior. However, when wealth maximization and high rates of survival are incompatible (our LS parameters), subjects persistently hoard excessive cash in an evident attempt to improve their prospects for survival.

Standard risk aversion over wealth may seem an appealing explanation for hoarding, but it cannot generate the types of deviations from optimality we observe in the data: risk aversion over wealth should induce subjects to withdraw more cash than is optimal, not less. Moreover, subjects do not lose accumulated earnings when they go bankrupt, suggesting that alternatives like loss aversion and prospect theory are poor explanations for our results. Individual level analysis suggests instead that our results are generated by a bias towards survival. Controlling for prior behavior, subjects set higher thresholds after experiencing bankruptcy than after surviving, indicating an aversion to bankruptcy events. This aversion is of roughly equal size in both LS and HS environments and seems to be the main driver of treatment differences.

³⁸Competing for revenue might cause decision makers to focus more heavily on profits relative to survival, tempering survival bias. Of course, an alternative possibility is that subjects in this setting will come to compete over survival, worsening the bias instead!

The body of the paper and the Supplemental Material report a total of ten robustness treatments covering five additional parameter sets and four additional environments. In all of these we observe the same basic pattern: subjects tend towards optimality when survival and profit maximization are in harmony and tend to hoard excess cash when they conflict. Our extensions suggest that the bias is robust to variation in the parameters generating the survival rate, the form and intensity of feedback, and the framing of the decision problem. Nonetheless, as we point out in the discussion, it is quite possible that framings or feedback protocols not explored here exist that are capable of mitigating or eliminating the bias's effects. The search for such protocols, perhaps guided by framing and feedback environments provided in rich, real-world settings, is a critical next step for this research program.

Much other work remains to be done on this topic. We studied two particularly diagnostically valuable parameter combinations in the main body of the paper (and five more reported in Appendix C of the Supplemental Material), but examination of others may provide a sharper sense of the circumstances under which survival bias leads to deviations from optimality. It seems clear from our data that subjects are willing to make decisions that generate *some* risk of bankruptcy, but avoid decisions that generate a great deal of bankruptcy risk. Further exploration of the parameter space may provide sharper evidence on how subjects trade off earnings and survival.

Our experiment abstracts away from a number of important elements in the firm's cash-use decision problem in order to focus crisply on survival. Future research may lift some of these abstractions and give us a deeper picture of how individuals balance competing uses for cash. One important abstraction employed here is that firms are at scale with no investment opportunities or growth prospects. Dutta and Radner's (1999) model extends to firms with growth opportunities and we suspect the model can be usefully adapted into more elaborate experiments. Likewise, our Investment design abstracts away from interactive features of markets for credit that may impact the use of cash within the firm. Examining survival decisions in the setting of more realistic competitive credit markets may be an interesting direction for future research.

Finally, our results may suggest a useful methodological point. The Market Selection Hypothesis—the classical justification for the profit maximization assumption in economic theory—holds that competitive markets are subject to natural-selection-like forces that kill off entities that fail to maximize profits (Friedman (1953)). The hypothesis is sometimes invoked to argue that departures from optimality sometimes observed in economics experiments are unlikely to survive in the field. Dutta and Radner (1999) pointed out that the hypothesis relies on an intimate link between profit maximization and survival and showed that if some agents have a bias towards survival, the hypothesis may fail and departures from optimality may survive. Our experiment shows that subjects can in fact exhibit a bias towards survival and therefore suggests

that the Market Selection Hypothesis may fail in certain settings. If biases towards survival are correlated with other biases, our results may suggest a mechanism by which other deviations from wealth maximization identified in laboratory experiments may survive in the field. This seems a promising avenue for investigation in future work.

REFERENCES

- ALCHIAN, A. (1950): "Uncertainty, Evolution and Economic Theory," *Journal of Political Economy*, 58, 211–221. [2225]
- BALLINGER, T., M. PALUMBO, AND N. WILCOX (2003): "Precautionary Saving and Social Learning Across Generations: An Experiment," *The Economic Journal*, 113, 920–947. [2228]
- BATES, T., M. KAHLE, AND R. STULZ (2000): "Why Do U.S. Firms Hold so Much More Cash Than They Used to?" *Journal of Finance*, 64, 1985–2022. [2226]
- BLUME, L., AND D. EASLEY (1992): "Evolution and Market Behavior," *Journal of Economic Theory*, 58 (1), 9–40. [2226]
- (2002): "Optimality and Natural Selection in Markets," *Journal of Economic Theory*, 107 (1), 95–130. [2226]
- (2006): "If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica*, 74 (4), 929–966. [2226]
- (2008): "Market Competition and Selection," in *The New Palgrave Dictionary of Economics* (Second Ed.), ed. by S. Durlauf and L. Blume. Basingstoke, U.K.: Palgrave Macmillan. [2226]
- BROWN, A. L., Z. E. CHUA, AND C. F. CAMERER (2009): "Learning and Visceral Temptation in Dynamic Savings Experiments," *Quarterly Journal of Economics*, 124 (1), 197–231. [2228]
- CHARNESS, G., AND D. LEVIN (2005): "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity and Affect," *American Economic Review*, 95 (4), 1300–1309. [2248]
- DIXIT, A. K. (1993): *The Art of Smooth Pasting*. Reading, U.K.: Harwood Academic. [2232]
- DUTTA, P., AND R. RADNER (1999): "Profit Maximization and the Market Selection Hypothesis," *Review of Economic Studies*, 66 (4), 769–798. [2225,2226,2228,2230,2233,2253]
- FRIEDMAN, M. (1953): *Essays in Positive Economics*. Chicago: University of Chicago Press. [2225, 2253]
- GREINER, B. (2004): "The Online Recruitment System ORSEE 2.0—A Guide for the Organization of Experiments in Economics," Working Paper Series in Economics 10, Department of Economics, University of Cologne. [2234]
- HAN, S., AND J. QUI (2007): "Corporate Precautionary Cash Holdings," *Journal of Corporate Finance*, 13, 43–57. [2226]
- HEY, J., AND V. DARDANONI (1988): "Optimal Consumption Under Uncertainty: An Experimental Investigation," *The Economic Journal*, 98, 105–116. [2228]
- HUCK, S., P. JEHIEL, AND T. RUTTER (2011): "Feedback Spillover and Analogy-Based Expectations: A Multi-Game Experiment," *Games and Economic Behavior*, 71 (2), 351–365. [2250]
- JEHIEL, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104. [2250]
- OPREA, R. (2008): "Free Cash Flow and Takeover Threats: An Experimental Study," *Southern Economic Journal*, 75 (2), 351–366. [2228]
- (2014): "Supplement to 'Survival versus Profit Maximization in a Dynamic Stochastic Experiment'," *Econometrica Supplemental Material*, 82, <http://dx.doi.org/10.3982/ECTA10751>. [2227]
- OPREA, R., D. FRIEDMAN, AND S. ANDERSON (2009): "Learning to Wait: A Laboratory Investigation," *Review of Economic Studies*, 76 (3), 1103–1124. [2228]

- PETTIT, J., J. HEWITT, AND R. OPREA (2014): “Redwood: Software for Graphical, Browser-Based Experiments in Discrete and Continuous Time,” *Experimental Economics* (forthcoming). [2234]
- RADNER, R., AND L. SHEPP (1996): “Risk vs. Profit Potential: A Model for Corporate Strategy,” *Journal of Economic Dynamics and Control*, 20, 1373–1393. [2226]
- ROTH, A., AND J. K. MURNIGHAN (1978): “Equilibrium Behavior and Repeated Play of the Prisoner’s Dilemma,” *Journal of Mathematical Psychology*, 17 (2), 189–197. [2233]
- SANDRONI, A. (2000): “Do Markets Favor Agents Able to Make Accurate Predictions,” *Econometrica*, 68, 1303–1341. [2226]
- SCHOTTER, A., A. WEISS, AND I. ZAPATER (1996): “Fairness and Survival in Ultimatum and Dictatorship Games,” *Journal of Economic Behavior and Organization*, 31, 37–56. [2228]
- SHUBIK, M., AND G. L. THOMPSON (1959): “Games of Economic Survival,” *Naval Research Logistics Quarterly*, 6, 111–123. [2226]

*Dept. of Economics, University of California, Santa Barbara, 3014 North Hall,
Santa Barbara, CA 93106-9210, U.S.A.; roprea@gmail.com.*

Manuscript received April, 2012; final revision received June, 2014.