Econometrica Supplementary Material

SUPPLEMENT TO "SURVIVAL VERSUS PROFIT MAXIMIZATION IN A DYNAMIC STOCHASTIC EXPERIMENT": ONLINE APPENDICES

(*Econometrica*, Vol. 82, No. 6, November 2014, 2225–2255)

BY RYAN OPREA

APPENDIX A: RISK AND SURVIVAL

IN THIS APPENDIX, WE SHOW that our results cannot be driven by standard risk aversion over earnings: risk aversion will not induce rational subjects to hoard (hold too-high thresholds). In fact, subjects in our experiment bear *more* earnings risk by hoarding cash as they do than they would by optimizing. To illustrate, Figure S1 plots the standard deviation of summed period earnings as a function of threshold choice for each treatment.[1] In each case, a star marks the optimum and a dot marks the median threshold chosen (during the final 10% of periods) in the Core treatments.

The plot reveals that earnings risk rises with the threshold between the optimum and the median threshold choice made under each set of parameters. The median subject in each case could have lowered her earnings risk (and simultaneously increased her expected profits) by moving to the optimum. In order to reduce earnings risk, a subject must do the exact opposite of hoarding by setting lower-than-optimal thresholds; an infinitely risk averse subject would liquidate at the first instant (by setting a threshold of 0), going bankrupt for sure but completely insulating herself from risk.

(Of course, at extremely high thresholds, making *any* withdrawals becomes rare, causing earnings risk to fall again but for a highly suboptimal reason. These high thresholds would not be chosen by rational risk averse agents, as the same earnings risk could be achieved with much higher expected earnings by setting a lower than optimal threshold instead. For instance, an HS subject would face similar earnings risk by setting a threshold of 40 or 8, but would, in the latter case, expect to earn twice as much!)

To illustrate further, consider the optimal threshold choices for agents with a standard exponential utility function, $U(x) = 1 - e^{-x/R}$. As R approaches infinity, preferences approach risk neutrality, but at low levels, agents are risk averse. Table SI presents optimal thresholds for various values of R (columns) and for each parameter set (rows). Focusing first on the LS treatment, as R

[1]Results are by simulation using the following procedure for each plot: (1) Draw a period length T from a geometric distribution, governed by parameter q; (2) draw a vector of cash movements, governed by binomial parameter p; (3) for each possible threshold choice, accumulate cash movements, withdrawals, and survival status; (4) repeat steps (1)–(3) 20,000 times and take the standard deviation for each threshold level. Note that this procedure replicates the exact lottery induced by threshold choices made by subjects each period in the actual experiment and so gives us a direct calculation of the earnings risk generated by each period's choice of threshold.

 DOI: 10.3982/ECTA10751

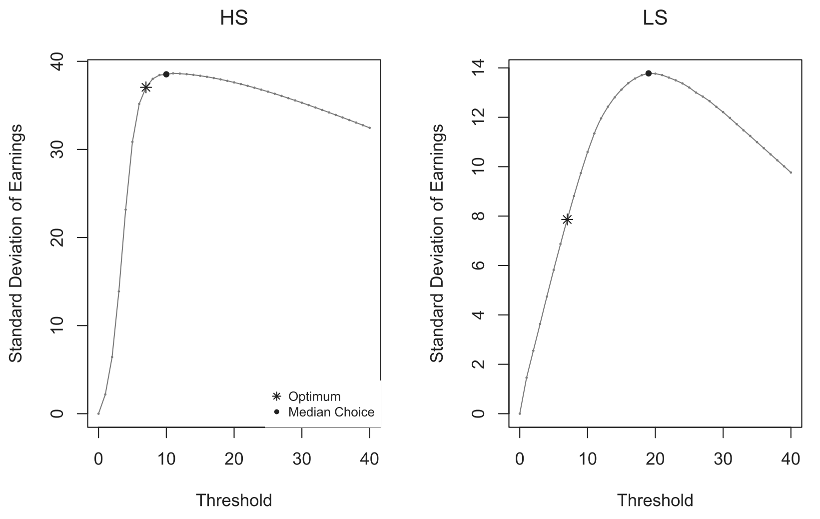

FIGURE S1.—Standard deviation of earnings as a function of the threshold for HS and LS parameters.

falls (as agents grow more risk averse), thresholds fall below the profit maximizing level of 7: risk aversion induces the opposite pattern of behavior to that observed in the experiment. At extremely high levels of risk aversion (e.g., $R = 1$), agents should simply liquidate immediately by setting a threshold of zero. The pattern in the HS treatment is similar though less dramatic: at intermediate levels of risk aversion, subjects should actually choose *lower* thresholds under LS than HS parameters, again the opposite of what we observe in the experiment.[2] That this is exactly the opposite of the comparative static observed in the data further confirms that risk aversion is a poor fit to our data.[3]

TABLE SI

OPTIMAL THRESHOLDS FOR AGENTS WITH UTILITY FUNCTION $U(x) = 1 - e^{-x/R}$ AS A FUNCTION OF R (COLUMNS) AND PARAMETERS (ROWS)

Parameters	1	30	60	90	1000
LS	0	2	4	6	7
HS	0	6	7	7	7

[2]The difference in predictions for HS and LS agents arises from the flatness of payoffs to the left of (below) the optimum for LS subjects. LS subjects lose very little in expected profits by setting too-low thresholds and therefore risk aversion induces relatively large changes in behavior. That we observe virtually no average thresholds below the optimum in the LS treatment is particularly striking given this—subjects lose far more in expected earnings (while simultaneously bearing more risk) by setting too-high thresholds than by setting too-low thresholds!

[3]Results are very similar using power utility instead of exponential utility.

This pattern seems, at first, quite counterintuitive—cash, after all, has a sort of insurance role against bankruptcy in this setting. However, it is important to make a clear distinction between survival risk and earnings risk. The two are generally inversely related: every unit of cash kept in the firm to insure against bankruptcy is left at risk of being evaporated by downward Brownian runs or random expiration. In order to safeguard cash from risk, a subject must withdraw it from the firm, increasing her risk of bankruptcy in the process.[4] Thus while it is quite possible that subjects systematically *believe* that hoarding is safer from an earnings perspective, it is not possible for this to be a rational response to the environment. Attempts to achieve earnings safety via increased rates of survival require subjects to mistakenly associate earnings and survival—just the type of bias we hypothesize is at work in these data.

Common risk-based alternatives to expected utility theory seem no better suited to these data. Because subjects do not actually lose any retained earnings by going bankrupt, loss aversion seems a poor fit to our findings.[5] Prospect theory is no more useful in explaining the data, as it simply predicts risk aversion over changes in wealth in the positive domain (as all wealth changes in this environment are).

APPENDIX B: INSTRUCTIONS TO SUBJECTS

In this section, we reproduce the instructions handed out and read aloud to subjects at the beginning of each session. Section B.1 reproduces instructions used in the Core and Social Learning treatments. Instructions were identical in these two cases except that: (i) the subsection labeled "Cohorts" was omitted in the Core treatment instructions and (ii) the screenshot used in the Core treatment was a variation containing only the top row of the feedback table below the main plot. Section B.2 reproduces instructions provided in the Investment treatments. These instructions reproduce the Core instructions but add a final section titled "Investment." Both instructions sets included Figure S2 and the Investment treatment instructions included Figure S3.

In order to avoid influencing subjects' decisions, screenshots used in all of these instructions feature both a cash scale and a set of binomial parameters that differ from those used in the experiment. Subjects were allowed to play with the experimental software using these same "test parameters" during the reading of the instructions and were told that these parameters were different from the ones to be used in the actual experiment.

[4]Of course, by withdrawing *nothing* (or nearly nothing) from cash reserves, a firm can simultaneously raise survival rates and reduce earnings risk but only by driving profits to (or very near) zero.

[5]It is possible that subjects are sub-optimally averse to the loss of future earnings possibilities. Under this (very expansive) interpretation, survival bias might be a relative to loss aversion.

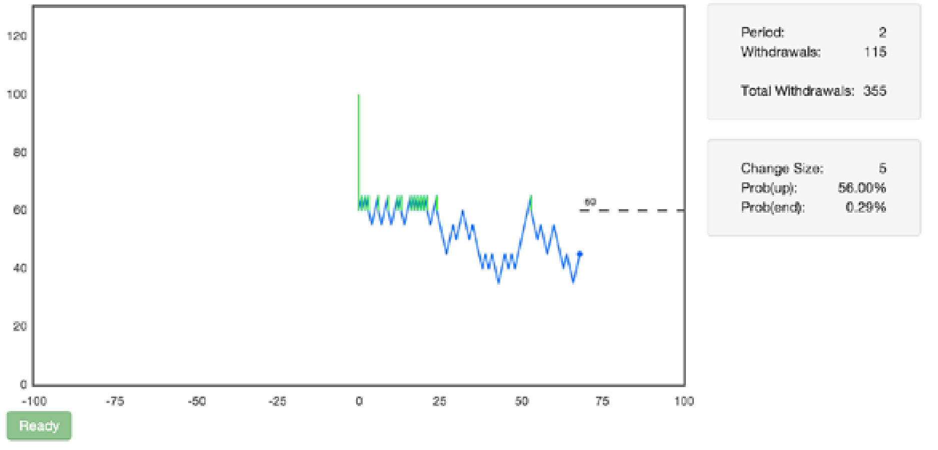

Current period:

Participant	Barrier	Withdrawals	
A (You)	60	115	
B	80	210	
C	40	135	bankrupt

FIGURE S2.—Screenshot from experiment software.

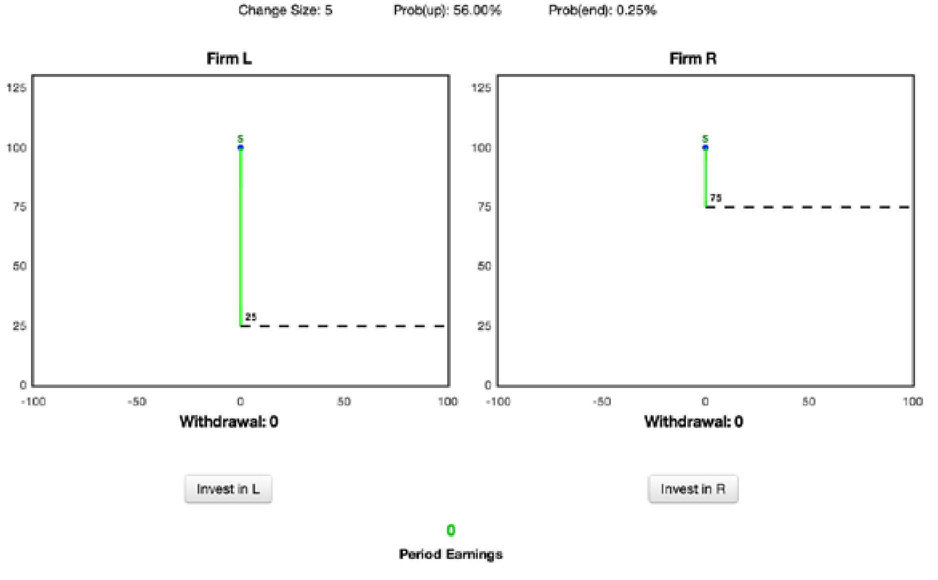

FIGURE S3.—Screenshot from experiment software.

B.1. *Core and Social Learning 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 insure 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 one of the experimenters will come.

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 you will control an *account of cash* the size of which will constantly change. Sometimes it falls and sometimes it rises (but it rises slightly more often than it falls). At the beginning of each period, you will decide on the maximum amount of cash you will allow to sit in your account—whenever the cash account randomly rises above this level, the excess will be *withdrawn as a payment to yourself*. If the account ever goes below zero, you will be *bankrupt*, and your cash account will stay at zero for the remainder of the period. At some random point the period will end and any cash left in the account will *disappear*. You will, however, keep whatever you have withdrawn so far. We will pay you based on the total amount of withdrawals you have made by the end of the period (NOT based on how much or how little cash you have accumulated in your account).

Screen Information

The *cash in your account, C*, is shown as a blue circle. Once the period starts, the past history of C will trail as a jagged blue line, moving left as on a ticker tape.

You will control a *barrier, B*, that determines the maximum amount of cash you will keep in your account during the period. The current barrier is shown on your screen as a horizontal dotted line. Any time the current cash level, C, rises above the barrier, B, all cash in excess of B will be withdrawn and the cash level will drop to the level of the barrier. The screen will show these withdrawals as vertical green lines (and at the moment you make a withdrawal a dollar sign will flash just above). The taller the green line, the more you have withdrawn, and the sum of all withdrawals is accumulated over the course of the period; the more you withdraw, the more you will earn. (Notice there will also be a green line and an initial withdrawal if you set the barrier below the initial level of cash when you start the period.)

At the beginning of the very first period, you will set the barrier by clicking and dragging on the blank screen. At the beginning of all other periods, you will first see the barrier from the previous period and can click and drag to adjust the barrier either up or down. Once you have set the barrier for the period, you can start the period by simply pressing the button labelled “Ready.”

If your cash account ever drops to zero, you will receive a notification at the top of the screen that you are bankrupt. The cash line will turn red and the cash will stay at zero for the remainder of the period (though any withdrawals you have already made during the period are yours to keep).

You will also be notified when the period *randomly ends*. When the new period starts, we will reset the level of cash and give you another opportunity to play. There will be a number of periods.

On the top right side of the screen, you will be shown the current period, the total level of withdrawals made during the period so far (Withdrawals), and the total level of withdrawals made during the entire experiment (Total Withdrawals, added up across all periods so far). On the bottom is a summary of your barrier, withdrawals, and bankruptcy status.

The higher the withdrawals you make during each period, the more you will earn. You will be paid, for each period, 5 cents for each point of period withdrawals above 30. We will sum up your earnings across periods and pay you that amount in cash at the end of the experiment.

Cohorts

During the experiment, you will be *randomly grouped* in a cohort with two other participants in the room who will be anonymously labeled “B” and “C.” Your decisions will have no impact on the earnings, cash accounts, or bankruptcy of other members of your cohort (and their decisions will have no impact on you). However, during the period you will be able to see the barrier choices and outcomes for these other participants in a table at the bottom of the screen. (Your decision and status will also be shown, highlighted in blue.)

You will stay in the same cohort with the same other participants for the entire experiment. Cohort assignments will be kept completely anonymous—you will never learn other participants’ identities and they will never learn yours. Members of the same cohort must start the period at the same time; after completing a period, you may have to wait for others to submit their barriers before the next period begins.

Details

Here are a few details of how C unfolds.

- The period is a series of many “ticks” (there are 5 ticks per second).

• Each tick, the current size of the cash account C moves randomly up or down by a fixed number of points, e.g., 1.1. The actual value will be provided on the lower right hand side of your screen, labeled as *Change Size*. This number will never change during the experiment.

• Upticks are slightly more likely than downticks; e.g., each tick is up with probability 56% or down with probability 44%. The actual value will be provided on the lower right hand side of your screen as *prob(up)*. This number will never change during the experiment.

• The Change Size and prob(up) levels will be identical for everyone in the experiment, but the computer will independently randomly determine actual changes to cash in each tick *independently* from person-to-person.

• The period ends (and the cash account evaporates) with a small probability each tick, e.g., 0.25% ($\frac{1}{2}$ of 1%). The actual value is provided on the right hand side of your screen as *prob(end)*. This number will never change during the experiment.

• The random process governing period lengths (described above) is mathematically *memoryless*. That means the expected number of ticks left in the period never changes, regardless of how many ticks have already passed. The expected number of ticks left in the period is exactly the same at tick 100 as it is at tick 0.

The prob(end) level will be identical for everyone in the experiment. The computer will randomly determine the same period length for everyone in your cohort.

• You are paid ONLY based on your withdrawals, NOT based on how large (or small) your cash account becomes (or any other factor). Whatever you withdraw is yours to keep (even if you go bankrupt).

Frequently Asked Questions

Q1: Is this some kind of psychology experiment with an agenda you haven't told us?

Answer: No. It is an economics experiment. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how you earn money in the experiment, and our interest is in seeing how people make economic decisions.

Q2: How long does a period last? Is there a minimum or maximum?

Answer: The length of time is random. For example, suppose the probability is 0.0029 that any tick is the last, and there are 5 ticks per second. In this case, the average length of a period is about 345 ticks or about 69 seconds. Many periods will last less than the average, and a few will last much longer. Periods longer than 7 minutes are so unlikely that you probably will never see one. The minimum length is one tick, but it is unlikely you will ever see a period quite that short!

Q3: How many periods will there be?

Answer: Lots. We aren't supposed to say the exact number, but there will be dozens and dozens of periods.

Q4: Are there patterns in upticks and downticks?

Answer: No. We've tried very hard to make it random. No matter what the recent history of upticks and downticks, the probability that the next tick is up is always the same (as displayed on your screen).

Q5: Do my choices affect the likelihood of an uptick, downtick, or the length of the period?

Answer: No. Your barrier choice obvious affects the amount of cash kept in the account, but the process governing movements of cash and period length are completely independent of your choices.

B.2. *Investment 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 insure 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 one of the experimenters will come.

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, you will control an *account of cash* the size of which will constantly change. Sometimes it falls and sometimes it rises (but it rises slightly more often than it falls). At the beginning of each period, you will decide on the maximum amount of cash you will allow to sit in your account—whenever the cash account randomly rises above this level, the excess will be *withdrawn as a payment to yourself*. If the account ever goes below zero, you will be *bankrupt*, and your cash account will stay at zero for the remainder of the period. At some random point the period will end and any cash left in the account will *disappear*. You will, however, keep whatever you have withdrawn so far. We will pay you based on the total amount of withdrawals you have made by the end of the period (NOT based on how much or how little cash you have accumulated in your account).

Screen Information

The *cash in your account, C*, is shown as a blue circle. Once the period starts, the past history of C will trail as a jagged blue line, moving left as on a ticker tape.

You will control a *barrier, B*, that determines the maximum amount of cash you will keep in your account during the period. The current barrier is shown on your screen as a horizontal dotted line. Any time the current cash level, C, rises above the barrier, B, all cash in excess of B will be withdrawn and the cash level will drop to the level of the barrier. The screen will show these withdrawals as vertical green lines (and at the moment you make a withdrawal a dollar sign will flash just above). The taller the green line, the more you have withdrawn, and the sum of all withdrawals is accumulated over the course of the period; the more you withdraw, the more you will earn. (Notice there will also be a green line and an initial withdrawal if you set the barrier below the initial level of cash when you start the period.)

At the beginning of the very first period, you will set the barrier by clicking and dragging on the blank screen. At the beginning of all other periods, you will first see the barrier from the previous period and can click and drag to adjust the barrier either up or down. Once you have set the barrier for the period, you can start the period by simply pressing the button labelled "Ready."

If your cash account ever drops to zero, you will receive a notification at the top of the screen that you are bankrupt. The cash line will turn red and the cash will stay at zero for the remainder of the period (though any withdrawals you have already made during the period are yours to keep).

You will also be notified when the period *randomly ends*. When the new period starts, we will reset the level of cash and give you another opportunity to play. There will be a number of periods.

On the top right side of the screen, you will be shown the current period, the total level of withdrawals made during the period so far (Withdrawals), and the total level of withdrawals made during the entire experiment (Total Withdrawals, added up across all periods so far). On the bottom is a summary of your barrier, withdrawals, and bankruptcy status.

The higher the withdrawals you make during each period, the more you will earn. You will be paid, for each period, 5 cents for each point of period withdrawals above 30. We will sum up your earnings across periods and pay you that amount in cash at the end of the experiment.

Details

Here are a few details of how C unfolds.

- The period is a series of many "ticks" (there are 5 ticks per second).
- Each tick, the current size of the cash account C moves randomly up or down by a fixed number of points, e.g., 1.1. The actual value will be provided on the lower right hand side of your screen, labeled as *Change Size*. This number will never change during the experiment.

- Upticks are slightly more likely than downticks; e.g., each tick is up with probability 56% or down with probability 44%. The actual value will be provided on the lower right hand side of your screen as *prob(up)*. This number will never change during the experiment.
- The period ends (and the cash account evaporates) with a small probability each tick, e.g., 0.25% ($\frac{1}{2}$ of 1%). The actual value is provided on the right hand side of your screen as *prob(end)*. This number will never change during the experiment.
- The random process governing period lengths (described above) is mathematically *memoryless*. That means the expected number of ticks left in the period never changes, regardless of how many ticks have already passed. The expected number of ticks left in the period is exactly the same at tick 100 as it is at tick 0.
- You are paid ONLY based on your withdrawals, NOT based on how large (or small) your cash account becomes (or any other factor). Whatever you withdraw is yours to keep (even if you go bankrupt),

Investment

In today's experiment, you will be shown, in each period, two separate but similar accounts labeled Firm L and Firm R. You will not set barriers for these accounts, and in fact, will have *no impact* on what happens to either account. Instead, you will decide in each period which of the two firms to *invest in.* If you choose to invest in Firm R, your earnings will be equal to Firm R's withdrawals during that period. If you choose to invest in Firm L, your earnings will be equal to Firm L's withdrawals during that period. (Your precise cash payments will be determined from your earnings each period using the formula described in the earlier instructions.)

The two firms will be driven by exactly the same Change Size, prob(up), and prob(end) variables (displayed at the top of the screen). In fact, the only difference between the two will be their barriers. In each period, the computer will select different barriers for each firm, chosen by *actual participants* during previous experiments. These previous participants were shown the exact instructions you just read, played the exact game just described, and were paid exactly as described above.

Figure S2 shows your computer screen with two charts: one for Firm L and another for Firm R. At the beginning of each period, the computer will select a different barrier choice made by a previous participant for each firm. You will see each firm's barrier as a horizontal dashed line prior to making your investment decision. Once the period begins, period withdrawals so far for each firm will be displayed just below each firm's chart in bold.

You will select which firm to invest in by clicking on the button below that firm's chart. When you do, the period will start for both firms simultaneously. The button you pressed will turn green, as will the withdrawals label of the firm you invested in. Your earnings so far for the period will be shown at the bottom of the screen (and will always be equal to the withdrawals of the firm you invested in). If either firm goes bankrupt, you will be notified above the firm's chart (though note: your earnings will be equal to the period withdrawals of the firm you invested in *whether or not* the firm goes bankrupt during the period).

Your investment decisions will determine how much you earn, but will have no effect whatsoever on the barrier, withdrawals, cash movements, or bankruptcy status of either firm either now or in the future. Likewise, your allocation decisions will have no effect on outcomes or payments for any other current or previous participant: they will only affect *your earnings*.

The Change Size, prob(up), and prob(end) variables, shown at the top of the screen, are identical for both firms. The period will end for both firms at the same time. Although the probability of a cash increase (prob(up)) is *identical* for both firms, the computer will randomly choose whether cash goes up or down for each firm *independently* in each tick. The *only difference* between the earnings prospects of the two firms is via their barriers, which are displayed on the screen as you make your investment decision.

Frequently Asked Questions

Q1: Is this some kind of psychology experiment with an agenda you haven't told us?

Answer: No. It is an economics experiment. If we do anything deceptive, or don't pay you cash as described, then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how you earn money in the experiment, and our interest is in seeing how people make economic decisions.

Q2: How long does a period last? Is there a minimum or maximum?

Answer: The length of time is random. For example, suppose the probability is 0.0029 that any tick is the last, and there are 5 ticks per second. In this case, the average length of a period is about 345 ticks or about 69 seconds. Many periods will last less than the average, and a few will last much longer. Periods longer than 7 minutes are so unlikely that you probably will never see one. The minimum length is one tick, but it is unlikely you will ever see a period quite that short!

Q3: How many periods will there be?

Answer: Lots. We aren't supposed to say the exact number, but there will be dozens and dozens of periods.

Q4: Are there patterns in upticks and downticks?

Answer: No. We've tried very hard to make it random. No matter what the recent history of upticks and downticks, the probability that the next tick is up is always the same (as displayed on your screen).

Q5: Do barriers affect the likelihood of an uptick, downtick, or the length of the period?

Answer: No. The barrier chosen obviously affects the amount of cash kept in the account, but the process governing movements of cash and period length are completely independent of the barrier.

Q6: Do my point allocations affect the likelihood of an uptick, downtick or the length of the period? Do they affect which barriers are chosen by the computer?

Answer: No. The process governing movements of cash and period length are completely independent of the point allocations you choose. Likewise, the barriers chosen by the computer from previous participants are not affected by your allocation choices.

APPENDIX C: ADDITIONAL ROBUSTNESS TREATMENTS

In this section, we report results from two additional robustness environments. The Multiple Accounts treatment replicates most of the features of the Core environment (though using different binomial parameters) but applies threshold choices to four separate and independently evolving accounts (instead of only one as in the Core treatments). The resulting fourfold increase in feedback (and increase in the probability that at least one account survives) has no significant effect on survival bias. The Flexible Threshold treatment is similar to the Core treatment but allows subjects to pause the game throughout the period to adjust thresholds. We report results from over 150 subjects in this environment across five parameterizations and find evidence for survival bias.

C.1. *Multiple Accounts Treatment*

In the Multiple Accounts environment, subjects set a single threshold at the beginning of each period just as in the Core treatments from the main body of the paper. However, this threshold is applied to four completely separate and independently evolving accounts each period,[6] each simultaneously displayed in a separate chart on the subject's screen. Subjects therefore observed four independent realizations of outcomes from their decisions each period and earned the sum of withdrawals made by this portfolio of accounts.

We ran 23 subjects through this 4A-LS treatment at the University of California, Santa Cruz and ran an additional 11 subjects through a control, 1A-LS,[7]

[6]Each account is endowed with $Y(0) = 40$ units of cash at the beginning of each period just as in the Core treatments.

[7]We ran an additional 11 subjects through a second session of this control treatment, but the session suffered a serious failure of experimental control: one subject announced a strategy and interpretation of the game to the entire session of participants. We therefore eliminate this session from the data set. However, reintroducing the session (expanding the control 1A-LS treatment to 22 subjects) has no effect on the statistical conclusions reported below (and in fact strengthens the statistical evidence for our conclusions).

that features identical parameters but only applies the threshold to one account (exactly as in our Core treatments). Both 4A-LS and 1A-LS use binomial parameters $(h, p, q) = (1.5, 0.52, 0.005)$.[8] These parameters generate an optimal threshold $\tau^* = 10.84$ and a low survival rate at this optimum of $s^* = 0.33$. The data set includes 50 periods of the game for each subject.[9]

We designed the Multiple Account treatment 4A-LS to reduce or eliminate the effect of survival bias. The protocol has two features that we hypothesized would lead to this result. First, the treatment quadruples the amount of feedback subjects receive in each period and from each threshold choice. If survival bias can be learned away with feedback, fifty complete repetitions with four independent realizations in each repetition should reveal this. Second, although the survival rate is low at the optimum as in other LS treatments (33% under these parameters), subjects hold a portfolio of accounts and can therefore expect at least one of the accounts to survive about 80% of the time. As in the Investment treatments from the main body of the paper, we expected this to "de-personalize" the bankruptcy event and make it less salient: it is not the subject who fails to survive in this treatment, but some subset of her investments.

The data do not support our motivating hypothesis: survival bias continues to be strongly evident in the 4A-LS treatment, is not significantly reduced relative to the control, and does not show signs of disappearing with experience. Figure S4 summarizes the results by plotting a time series of median threshold choices. The figure reveals a pattern in both the 4A-LS and 1A-LS treatments that *exactly* matches the one reported for the C-LS treatment in the main text. Subjects initially set extremely high thresholds, reduce these thresholds somewhat (and slowly) in early periods, and soon stabilize at a level far above the optimum. By the final 10 percent of periods, the median subject's median threshold is twice as high as the optimum in both the 4A-LS and 1A-LS treatments (23 and 24.5, respectively). We cannot reject the hypothesis that thresholds are identical in these two treatments ($p = 0.222$, Kolmogorov–Smirnov test). Thus although the Multiple Account protocol substantially increases the amount of feedback subjects receive and greatly increases the chances that at least one of the subject's accounts survives, (i) subjects persistently hoard a large amount of cash and (ii) behave no differently than subjects in the standard protocol. The treatment therefore provides further evidence of the robustness of survival bias to feedback and framing.

[8]All four accounts in 4A-LS are driven by the same binomial parameters, but realizations of cash shocks each tick (and therefore withdrawals and the timing and incidence of bankruptcy) are independent from account to account.

[9]We intended to run 60 periods, but due to software errors, subjects were dropped from the experiment after period 50, forcing us to restrict the data set and analysis to the first 50 periods.

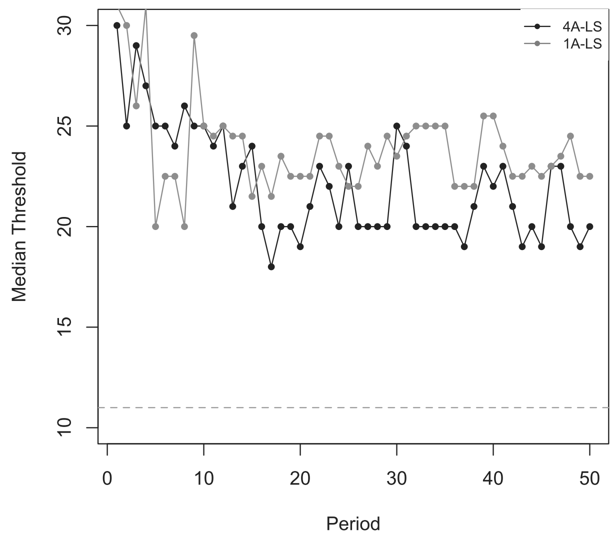

FIGURE S4.—Median thresholds in the 4-Account protocol (4A-LS) and the control treatment 1A-LS.

C.2. *Flexible Threshold Treatments*

We ran 130 subjects through four main additional treatments using a completely different set of parameters[10] from the ones used in the main body of the text and a somewhat different decision protocol (called "Flexible Threshold"). These serve as a useful robustness check for our results. The 2×2 design used in these sessions (summarized in Table SII), varies the optimal threshold level, τ^*, between roughly 10 and 22 and the optimal survival rate, s^*, between 33%

TABLE SII

BINOMIAL PARAMETER VALUES, OPTIMAL THRESHOLDS, AND OPTIMAL SURVIVAL RATES BY TREATMENT[a]

Treatment k	Step Size h	Uptick Prob. p	Expiration Prob. q	Threshold τ^*	Survival Rate s^*
LS-10	1.5	0.52	0.005	10.84	0.33
LS-22	2.5	0.52	0.004	22.10	0.34
HS-10	0.74	0.57	0.004	10.84	0.88
HS-22	1.5	0.565	0.004	22.01	0.88

[a]In all treatments, the time step in seconds is $\Delta t = 1/5$.

[10]Though the initial endowment of cash at the beginning of each period, $Y(0)$, is set to 40 in each case, just as in the paper's main treatments.

(LS treatments) and 88% (HS treatments). The inferential strategy for which this design was constructed is identical to that of the main data set reported in the paper: holding τ^* constant, we compare average HS and LS threshold choices (i.e., we compare LS-10 to HS-10 and LS-22 to HS-22). Under the hypothesis that survival bias exists, threshold choices should tend towards optimal levels under HS parameters but remain far above the optimum under LS parameters.

To this main 2×2 robustness design we added a fifth robustness treatment—HS-32—that features a threshold of 32.58 and survival rate of 88%, generated by parameter vector $(h, p, q) = (2.5, 0.565, 0.005)$. We ran 29 subjects through this treatment for a total of 159 subjects in the Flexible Threshold data set. No subject in the data set experienced more than one set of parameters (this was an entirely between-subjects design).

Several details of the design (beyond the parameter configurations and resulting τ^* and s^*) of these robustness sessions differed from the main experiment in the paper:

- Subjects were allowed to pause the experiment and adjust their barrier at any point (and as often as they wanted) during the period and therefore had an additional way to stray from optimal behavior: by failing to set an overflow policy of the optimal sort.
- Sessions lasted for 60 rather than 40 periods (each, again, an entire run of Dutta and Radner's model).[11]
- While a threshold payoff transformation is used in these sessions, it is applied to total withdrawals accumulated over periods rather than to each period individually.
- Sessions were conducted at the University of California, Santa Cruz and used a piece of custom software (programmed in JAVA) that differed in minor, cosmetic ways from the software used in the main experiments.

These treatments give us an additional related (but distinct) environment in which to test for survival bias. The results of these treatments yield the following main results:

- Most subjects use strategies that resemble overflow policies of the optimal class. In an overflow policy, an agent makes (at most) one initial withdrawal larger than the binomial parameter h during the period, and after, makes only withdrawals of size h, whenever cash overflows the threshold.[12] Overall, 84%

[11]Due to software errors, some additional subjects had to be dropped from the data set. Some of the subjects remaining in the data set are missing data from the final period (period 60), again due to software errors.

[12]When a subject makes multiple withdrawals larger than h, it suggests she is allowing cash to accumulate in her account with the intention of withdrawing it as a lump sum. This exposes cash needlessly to random expiration risk and is therefore suboptimal. We call this suboptimal alternative to an overflow policy an "accumulation policy."

of subjects make only one withdrawal greater than h in their median period[13] and in each period the median withdrawal is precisely of size h.

- In addition to setting overflow-like policies, subjects tend to use relatively stationary thresholds to make withdrawals during the period.[14] In the modal period, subjects withdraw at only a single threshold[15] and the median difference between the highest and lowest threshold within-period is a mere 1.27 points. When subjects do change their thresholds during the period, they are considerably more likely to reduce them than increase them (70% of threshold changes within-period are decreases). This pattern of behavior is consistent with subjects "chasing the surplus"—moving the threshold towards current cash levels during runs of bad luck—, a possible sign that some subjects fail to grasp the "memoryless" property of the geometric distribution governing period lengths. It is also consistent with subjects simply adjusting due to learning within the period.
- Figure S5 plots HS (in black) and LS (in gray) median threshold decisions over the course of 20 three-period bins. The left hand panel plots treatments HS-10 and LS-10 (T-10) treatments, while the right hand panel plots treatments HS-22 and LS-22 (T-22) treatments.[16] The results show that median HS thresholds tend to move decisively towards optimal threshold levels over time and settle near their respective optima in both T-10 and T-22 settings. By contrast, LS thresholds tend to remain well above optimal levels, and after the first few periods are always higher than HS thresholds.
- Figure S6 shows that thresholds are very sensitive to the economic forces of the model and have surprisingly strong convergence properties (given the complexity of the model) when survival concerns are mostly absent (HS treatments). HS-10 and HS-22 time series separate cleanly to two distinct optima. By contrast, optimal predictions seem to have no power to organize LS data. Thresholds in LS-10 and LS-22 are very similar and both are far above their respective optima.
- Regressions using period-to-period threshold adjustments suggest that bankruptcy has a significant impact on period-to-period adjustments even after controlling for withdrawals. Moreover, individual-level regressions suggest that the distribution of HS and LS bias measurements (β) are positive for most subjects and are similarly distributed across the two treatment types.

[13]That is, for each subject we calculate the number of withdrawals greater than h for each period. The median number of withdrawals over periods is exactly 1 for 84% of subjects.

[14]We use only thresholds at which actual withdrawals were made in our data set (in order to filter out threshold changes made for non payoff reasons, such as boredom alleviation).

[15]Interestingly, the degree of stationarity is of roughly the same size early and late in the session; the typical subject does not appear to learn to be more (or less) stable with experience.

[16]To measure a subject's threshold during a period, we look at the sequence of withdrawals made by the subject and take the median threshold choice determining these withdrawals.

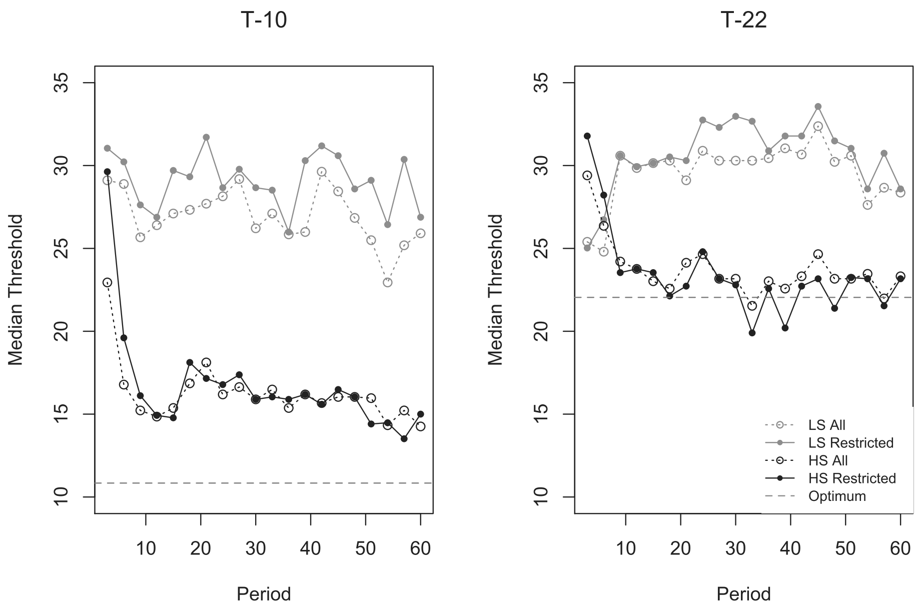

FIGURE S5.—Median thresholds plotted (in 20 bins) over periods, organized by optimal threshold predictions (T-10 and T-22). Solid lines show only the cases (labeled Restricted) that are consistent with overflow policies, while dotted lines show all data.

These robustness treatments support the conclusions reported in the paper using the main design, and suggest these conclusions are robust to changes in the action space, parameterization, threshold predictions, survival rates, and time horizons. Several of these changes are worth highlighting. First, hoarding is unaffected by a change to a Flexible Thresholds protocol, probably because subjects in these treatments make, on average, little use of such flexibility and tend to hold relatively static overflow thresholds similar to the ones implemented by design in the main paper's setup. Second, hoarding in LS treatments persists beyond the 40 periods studied in the main body of the text without any sign of reduction, further supporting the paper's conclusion that survival bias is quite robust to learning. Third, we observe significant hoarding in LS treatments even though it features a somewhat weaker conflict between profit maximization and survival than in the main paper's design (subjects survive only 15% of the time at the optimum under the main paper's LS parameters, but 33% of the time at the optimum under these robustness LS parameters).

Although the results from these treatments tend to support the results from the main body of the paper and provide some additional perspective on the results' durability to changes in environment, they have some inferential dis-

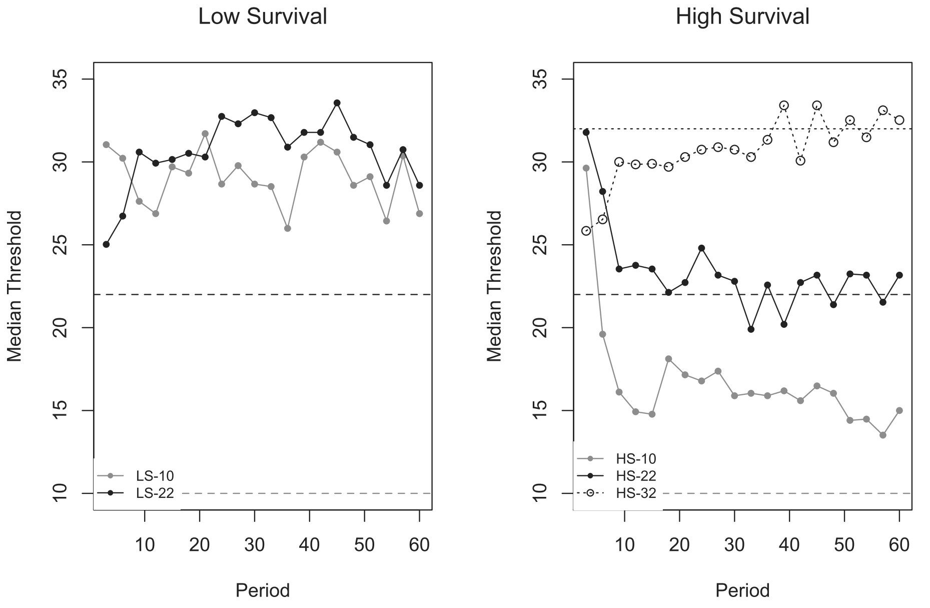

FIGURE S6.—Median thresholds plotted (in 20 bins) over periods, organized by survival rates (Low Survival and High Survival).

advantages relative to our main design.[17] Most importantly, allowing for Flexible Thresholds has the potential to induce an aversion to bankruptcy not due to a heuristic error but simply due to a desire to avoid boredom. By avoiding bankruptcy, subjects can continue making economically relevant threshold choices; after going bankrupt, subjects no longer have anything to do in the experiment until the next period begins. In the Fixed Threshold protocol, featured in the paper, we restrict subjects to choosing a single threshold at the beginning of the period (as predicted in the model and as the typical subject does in these robustness treatments), eliminating boredom-avoidance motives: subjects are unable to make further decisions after the period begins *whether or not* they survive. That we get very similar results in both settings suggests that such boredom motives were not the main driving force of the Flexible Threshold results. Nonetheless, we feature the more directly interpretable (and therefore more reliable) Fixed Threshold protocol in the main body of the paper.

[17]Indeed, these sessions were run prior to the main ones reported in the paper. The Fixed Threshold protocol used in the paper's experiments was designed and implemented largely to eliminate these potential inferential problems. We thank the editor for pointing the most severe of these out.

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.