The Disposition Effect on Optimal Stopping Decisions: a Direct Test*

Jacopo Magnani†

December 2, 2014

Abstract

This paper develops a laboratory test of a distinctive prediction of reference-dependent preferences, testable only in dynamic settings: decision makers suboptimally delay realizing disappointing outcomes (procrastination) but suboptimally rush to realize outcomes that are better-than-expected (rushing). In the experiment, subjects invest in a risky asset, whose price evolves in near-continuous time, and they are provided with the option to liquidate it at a fixed salvage value. Optimal behavior is characterized by an upper and a lower stopping thresholds in the asset price space, thus producing a clear rational benchmark and eliminating known confounds. Most subjects indeed tend to delay liquidating losing assets beyond the optimal point and to sell winning assets before reaching the optimal stopping time. Among subjects who show the effect, the median stopping points imply the probability of realizing a winner conditional on stopping is 70% larger than optimal. Such behavior is shown to be consistent with a model of a decision maker who evaluates payoffs relative to an expectation-based reference point, is risk-averse over gains and risk-seeking over losses.

Keywords: Laboratory experiments, disposition effect, expectation-based reference-dependent preferences, prospect theory, optimal stopping decisions.

JEL codes: G02 D81 C91

*I am grateful for helpful comments from Dan Friedman, Ryan Oprea, Meir Statman, Colin Camerer, Lawrence J. Jin, as well as from seminar participants at the Western Economics Association International Graduate Student Workshop in Seattle, Santa Clara University, University of Technology - Sydney, New York University, Einaudi Institute for Economics and Finance, Bay Area Behavioral and Experimental Economics Workshop in Stanford, Yale Whitebox Advisors Graduate Student Conference, Experimental Finance Conference in Zurich, Foundations of Utility and Risk Conference in Rotterdam. I am thankful to Matthew Browne for developing the experiment software. I acknowledge funding from the National Science Foundation Grant SES 1260813.

†Wang Yanan Institute for Studies in Economics, Xiamen University, Xiamen, 361005 Fujian, P.R. China. Email: jacopo.magnani@gmail.com

1 Introduction

The disposition effect - the tendency of individual investors to sell assets whose price has increased and hold on to assets that have dropped in value - is a cornerstone of behavioral finance.[1] More generally, the disposition effect is an important notion for behavioral economics as it exemplifies the idea that decision makers suboptimally delay realizing disappointing outcomes (procrastination) but suboptimally rush to realize outcomes that are better-than-expected (rushing).[2] However, previous evidence has been mostly indirect. For instance, the typical procedure used in studies on the disposition effect in stock markets involves testing whether the frequency of sales is larger for winning stocks than for losing stocks.

I propose and conduct a sharper test of the disposition effect using a laboratory experiment that produces an unambiguous benchmark of rational stopping decisions. I design a security that bundles a risky asset, whose price follows a stochastic process in continuous time, with the option to liquidate the asset at a constant salvage value. I analyze an impatient decision maker who makes a decision about when, if ever, to liquidate the investment. Optimal behavior entails maintaining the current position in the security until either:

1. the asset price reaches an upper threshold B^* above which the expected benefit from waiting is outweighed by the immediate reward from selling the asset and the decision maker reaps this opportunity, or

2. the asset price reaches a lower threshold b^* and the decision maker capitulates, liquidating at an exogenously fixed salvage value and forgoing potential future price increases.

Thus my model provides a clear rational benchmark against which to evaluate the disposition effect. Indeed, the original formulation of the disposition effect by Shefrin and Statman [1985] as the tendency to sell winners *too early* and ride losers *too long* can be formally characterized in terms of the stopping times induced by the optimal thresholds. There is a disposition effect if the actual thresholds (B and b) are both lower than optimal ($b < b^* \& B < B^*$). Results from my experiment strongly support the hypothesis that individual investors have a preference for realizing winners vis-à-vis realizing losers. Both at the aggregate and at the individual level, subjects tend to sell the asset as soon as its price has increased to a point significantly below B^*, while they wait for the price to fall considerably below b^* before

[1]Evidence surveyed in Barber and Odean [2011] suggests the disposition effect in financial markets; see also Genesove and Mayer [2001] for evidence regarding the US real estate market.

[2]Problems where the timing of uncertainty resolution is endogenous are ubiquitous in modern economics. In section 5 I review two potential applications of the disposition effect in the field of information economics, pertaining to experimentation and signaling respectively.

capitulating. This departure from rational behavior is economically significant as the median stopping points imply a probability of realizing a winner conditional on stopping 60% larger than optimal.

Laboratory research provides an important complement to studies on the disposition effect that use financial market data. Obtaining and skillfully analyzing field evidence is a necessary step for understanding the behavior of individual decision makers, however important aspects of the decision making process are unobservable in naturally occurring settings. It is often difficult to identify a normative benchmark against which the disposition effect can be measured. In practice, field data studies rely on the implicit assumption that a rational decision about how long to hold on to an investment should be independent of whether the investment is a winner or a loser. Thus, the typical test of the disposition effect involves checking whether sales of winners are more likely than sales of losers.[3] However, these tests are not grounded in any specific theory and the results may be subject to different interpretations.[4] In my experiment I define an unambiguous benchmark and eliminate confounding factors by design.

While important laboratory work has been conducted on the disposition effect, such as Weber and Camerer [1998], most face potential confounds that make the effect difficult to interpret and difficult to distinguish from rational behavior. At a basic level my design differs from Weber and Camerer [1998] in the mechanism used to induce stopping decisions: while in Weber and Camerer [1998] liquidation decisions arise from portfolio choice motives, here stopping is induced by inter-temporal trade-offs, which more closely matches the environment contemplated in recent theoretical work on the effect. In (lab and field) environments where decision makers are actively engaged in portfolio choice, risk aversion and other diversification motives may lead to differences in the propensity to sell winners and losers.[5] In the multiple heterogenous asset framework of Weber and Camerer [1998] the optimal, expected-value maximizing behavior is to hold on to the single asset that the decision maker identifies as the winner based on her beliefs at a point in time. However, this is also a very risky strategy, because of uncertainty about which asset is the actual winner. Thus holding on to losers

[3] While early evidence on the disposition effect was based on such statistics as the fraction of realized gains relative to realized losses (Odean [1998]), recent studies take a more dynamic view of trading and study how the hazard of a sale is affected by a paper gain or loss (see the survey by Barber and Odean [2011]).

[4] Researchers have interpreted similar findings in opposite ways. A number of recent papers, summarized in Barber and Odean [2011], show that the hazard rate of stock sales as a function of return since purchase is much steeper for gains than for losses and argue that this supports the existence of the disposition effect. Ben-David and Hirshleifer [2012] estimate similar hazard functions, but they argue that there is no clear evidence of preferences for selling a stock by virtue of having a gain versus a loss since there is no upward jump in selling at zero profits.

[5] For example, portfolio rebalancing after price changes is known to be a potential explanation for disposition effect-like behavior in empirical studies (Odean [1998]).

might just be a way to hedge some risk.[6] In my design, standard risk preferences induce behavior that is qualitatively different from the disposition effect. Risk aversion produces a narrowing of the inaction band ($b > b^* \& B < B^*$): since a risk-averse decision maker prefers less payoff variability, she liquidates losers sooner than predicted by risk-neutrality and by the disposition effect.

This paper provides empirical content to theoretical investigations on microfounded models of the disposition effect. Seminal works on the disposition effect, such as Shefrin and Statman [1985], Odean [1998] and Weber and Camerer [1998], tended to ascribe the disposition effect to reference-dependent preferences, featuring risk aversion over gains and love of risk over losses, as hypothesized by Kahneman and Tversky [1979]'s prospect theory[7]. A further contribution of this paper is to formalize this argument for the stylized optimal stopping problem I consider. I present a model of expectation-based reference-dependent preferences and show that it reliably predicts the disposition effect: a decision maker with these preferences procrastinates realizing disappointing outcomes but rushes to realize better-than-expected ones. Moreover, I show that there is a clear mapping between the model's parameter (the elasticity of utility to gains and losses) and the stopping thresholds. The model qualitative and quantitative predictions are borne out in the data for a majority of subjects and the experiment allows me to rule out competing belief-based explanations. This is the first application of expectation-based reference-dependent preferences to stopping decisions I am aware of and the problem formulation and solution methods can be readily extended to other one-time stopping problems.

The paper is organized as follows. In the next section I provide a theoretical description of the decision problem I implement in the lab. I first describe the optimal strategies of a risk-neutral decision maker (section 2.1). I then discuss how to measure the disposition effect in this setting (section 2.2). I study the behavior induced by standard risk-preferences (section 2.3) and I show that reference-dependent preferences are predicted to result in the disposition effect (section 2.4). Section 3 describes the experimental design. Section 4 presents the results and section 5 concludes.

[6] The potential confounds in Weber and Camerer [1998] are not limited to risk aversion. Failures in Bayesian learning of asset qualities is another concern that I eliminate in my design. Furthermore in Weber and Camerer [1998] optimal behavior entails never selling a winner before stock holdings are liquidated in an exogenously fixed final period, so any variability in sales of winners is interpreted as support of the disposition effect.

[7] The relation between the disposition effect in financial markets and prospect theory remains a controversial issue, see for example Barberis and Xiong [2009] and Ingersoll and Jin [2013]. The aim of this paper is not to exactly replicate the environment real investors face. For example, I ignore the possibility of reinvestment. I show that reference-dependent preferences are a powerful explanation of the disposition effect in a stylized optimal stopping problem and therefore can be expected to apply to similar situations.

2 Theoretical Framework

2.1 Optimal Asset Liquidation with a Constant Salvage Value

The theoretical framework considers an asset (hereafter stock) that is a claim to a single random dividend. I work on the standard probability space $(\Omega, \mathcal{F}, P)$, with filtration $\{\mathcal{F}_t, t \geq 0\}$ supporting a Wiener process $w = \{w_t, t \geq 0\}$. The stock price s follows a geometric Brownian motion:

$$ds_t = \mu s_t dt + \sigma s_t dw_t \tag{1}$$

The stock pays out the dividend according to a Poisson process with intensity λ and after the dividend is paid out the stock expires. The dividend (conditional on its realization), y_t, is given by:

$$y_t = \delta s_t \tag{2}$$

where δ is an exogenous parameter[8].

In order to obtain a simple model that is consistent with liquidating the stock at both high and low prices, I consider the behavior of an impatient and risk-neutral decision maker who is endowed with a unit of the stock and the option to liquidate it at a constant salvage value. At each point in time, the decision maker can choose one of the following actions: 1) hold on to the investment and wait, 2) cash the stock and receive its price s_t or 3) exercise the option and obtain the salvage value (or strike price) x. Let τ_o be the time at which the decision maker exercises the option and τ_c the time at which he cashes the stock. Similarly, let τ_λ be the random time of the dividend arrival, at which point the decision maker obtains the dividend according to equation (2). Let

$$\tau \equiv \min\{\tau_o, \tau_c, \tau_\lambda\}$$

Then the problem of the decision maker is described by the following value function:

$$v(s_t) = \max_{\tau_o, \tau_c} E_t \left\{ e^{-\rho(\tau - t)} \left[\mathbb{1}_{\{\tau=\tau_o\}} x + \mathbb{1}_{\{\tau=\tau_c\}} s_\tau + \mathbb{1}_{\{\tau=\tau_\lambda\}} y_\tau \right] \right\} \tag{3}$$

where $\mathbb{1}$ is an indicator function and I assume the decision maker discounts future payoffs at an instantaneous rate of ρ[9].

[8] In a general equilibrium setting δ is the inverse of the pricing kernel, but here I take it as an exogenous parameter since I am studying the disposition effect at the individual level.

[9] As noted above, in a general equilibrium setting δ and ρ should be related by the pricing mechanism. In fact, I will deliberately calibrate the experiment in such a way that the decision maker discounts future at a higher rate than what would be implied by a general equilibrium interpretation of the asset price process. This is necessary to avoid that the decision maker always waits for the dividend to realize. Many behavioral finance models assume that investors are more impatient than what is otherwise standard, see for example Ingersoll and Jin [2013].

The solution to this optimal stopping problem is characterized by an inaction region (b, B). As soon as s reaches B the decision maker sells (cashes the stock): $\tau_c = \inf\{t : s_t \geq B\}$. As soon as s reaches b the decision maker sells by exercising the salvage option: $\tau_o = \inf\{t : s_t \leq b\}$. Then the value function satisfies:

$$v(s) = s, \forall s \geq B \tag{4}$$

$$v(s) = x, \forall s \leq b \tag{5}$$

As shown in appendix A.1, inside the inaction region, the value function is given by:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}s + C_1 s^{R_1} + C_2 s^{R_2}, s \in (b, B) \tag{6}$$

where R_1 and R_2 are algebraic functions of the parameters, while C_1 and C_2 are constants to be determined (see appendix A.1 for details).

I denote the optimal risk-neutral thresholds by b^* and B^*. The associated optimal stopping times that achieve the maximum in equation (3) are denoted by τ_o^* and τ_c^*. For the optimal thresholds, the value function is continuous and differentiable at the boundaries of the inaction region (see Dixit [1993]) and thus the following value matching and smooth pasting conditions must hold:

$$\lim_{s \searrow b^*} v(s) = x \tag{7}$$

$$\lim_{s \searrow b^*} v'(s) = 0 \tag{8}$$

$$\lim_{s \nearrow B^*} v(s) = B^* \tag{9}$$

$$\lim_{s \nearrow B^*} v'(s) = 1 \tag{10}$$

From this system of equations it is possible to determine that the constants C_1 and C_2 are related to the optimal thresholds, b^* and B^*, by the following conditions:

$$C_1 = -\frac{x}{{b^*}^{R_2}{B^*}^{R_1} - {b^*}^{R_1}{B^*}^{R_2}} \frac{{B^*}^{R_2}}{1-R_1} \tag{11}$$

$$C_2 = \frac{x}{{b^*}^{R_2}{B^*}^{R_1} - {b^*}^{R_1}{B^*}^{R_2}} \frac{{B^*}^{R_1}}{1-R_2} \tag{12}$$

The two optimal thresholds can then be found by solving (7),(8),(9),(10) numerically. The value function of the decision maker's problem is illustrated in Figure 1.

The intuition behind these optimal threshold rules is the following. At the upper threshold B^* the expected benefit from waiting for the price to rise further is outweighed by the

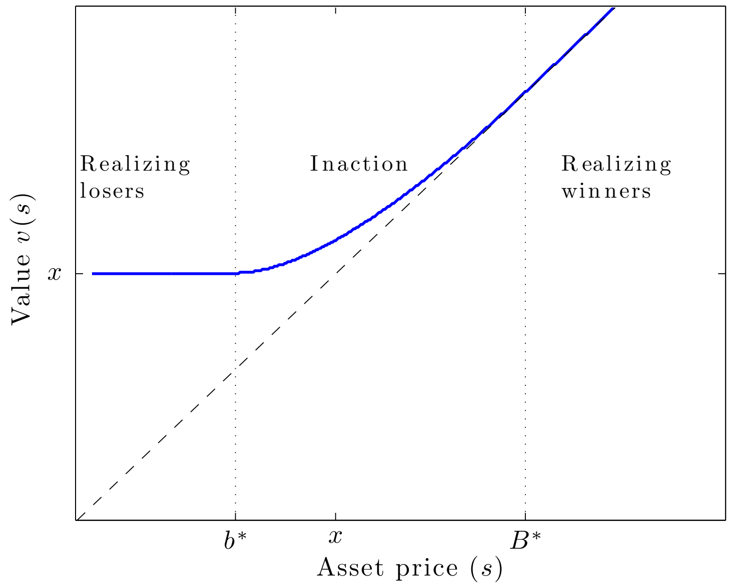

Figure 1: Value function and inaction region of the problem

immediate reward from selling the asset and the decision maker reaps this opportunity. At the lower threshold b^*, the salvage value x exceeds the value of waiting for the price to reach the upper threshold and the decision maker optimally capitulates, forgoing potential future price increases. In a static setting (or when $\sigma \to 0$ or $\rho \to \infty$), the problem reduces to a standard protective put strategy with the following optimal liquidation rule: cash the stock if $s > x$, exercise the option if $s < x$. I summarize the analysis of the risk-neutral case in the following:

Remark 1. *Optimal behavior involves holding on to the investment until the asset price reaches either an upper threshold B or a lower threshold b. The two optimal thresholds for a risk-neutral decision maker, B^* and b^*, solve equations* (7),(8),(9),(10).

2.2 The Disposition Effect: Definition and Measurement

The model presented above provides a clear benchmark for measuring the disposition effect. The classical definition of the disposition effect is the tendency of investors to hold on to losing stocks for too long and to realize winning stocks too early, where winners and losers are defined relatively to the purchase price s_0. In the following discussion I will assume that the purchase prices coincides with the salvage value, as in the lab implementation of the model, i.e. $s_0 = x$. This implies that an agent exercises the safe option if and only if she liquidates a loser.

In the current setting the disposition effect will result in the lowering of both boundaries of the inaction region relative to the optimal benchmark, as illustrated in Figure 2, where I

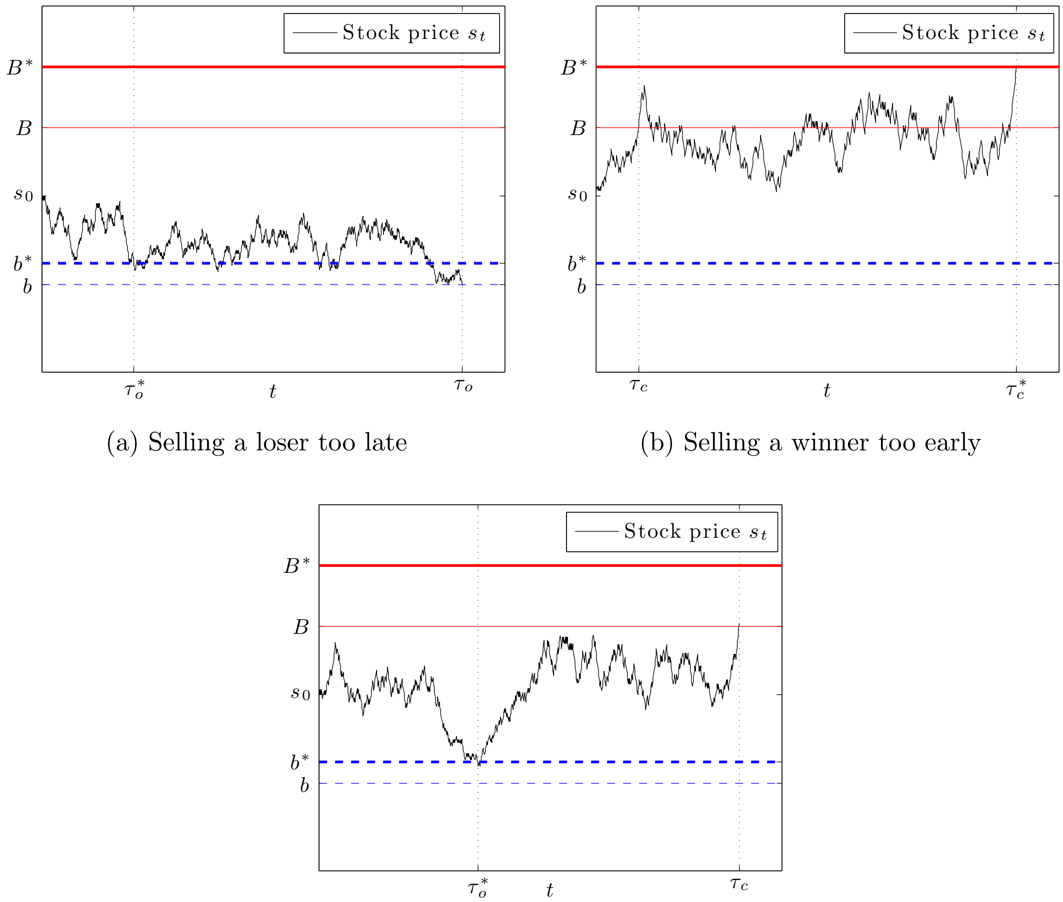

(a) Selling a loser too late

(b) Selling a winner too early

(c) Selling a winner instead of a loser

Figure 2: Disposition Effects

denote by B (b) the actual upper (lower) liquidation threshold used by the decision maker. Panels 2a and 2b illustrates how the classical definition of the disposition effect applies in this setting: in the former the decision maker liquidates a loser too late ($\tau_o > \tau_o^*$), while in the latter he sells a winner too early ($\tau_c < \tau_c^*$). Panel 2c illustrates the notion that the disposition effect leads to realizing winners more frequently and losers more rarely than optimal. The situation illustrated in Figure 2 is summarized in the following:

Definition 1. *Let $\{b^*, B^*\}$ be the optimal thresholds for a risk-neutral decision maker and $\{b, B\}$ be the actual thresholds (or their empirical counterpart). $\{b, B\}$ satisfy the disposition effect condition if:*

$$b < b^* \wedge B < B^* \tag{13}$$

.

I introduce a measure of the bias towards realizing winners vs. losers. First, I compute the probability of realizing a winner conditional on stopping. This is equal to the probability that s_t hits B rather than b, conditional on s_t hitting one of the two stopping thresholds. In the case $\mu = 0$ considered in the experiment, this probability is given by: $\Pi \equiv \frac{x-b}{B-b}$. This provides a compact and meaningful way of combining the two threshold values and can be readily compared with the optimal benchmark value of the probability: $\Pi^* \equiv \frac{x-b^*}{B^*-b^*}$, leading to the following:

Definition 2. *Let Π and Π^* be the probability of realizing a winner conditional on stopping induced by the actual and optimal stopping thresholds respectively. The winner bias, Ψ, is defined as the increase in the probability of realizing a winner conditional on stopping relative to the optimal benchmark:*

$$\Psi \equiv \frac{\Pi - \Pi^*}{\Pi^*}$$

While a positive winner bias does not necessarily signal a disposition effect as defined above, Ψ can be used to gauge the economic significance of the effect. More importantly, the empirical analysis will reveal that subjects who show the disposition effect are indeed those with a strongest winner bias.

2.3 Standard Risk Preferences

It is important to note that standard risk aversion (or love of risk) does not generate a disposition effect. The decision problem was carefully designed to ensure this. I illustrate this point in the case of CRRA utility. In general the solution for a decision maker with Bernoulli utility function $f(m)$ is given by:

$$\max_{\tau_o,\tau_c} E_t \left\{ e^{-\rho(\tau-t)} \left[\mathbb{1}_{\{\tau=\tau_o\}} f(x) + \mathbb{1}_{\{\tau=\tau_c\}} f(s_\tau) + \mathbb{1}_{\{\tau=\tau_\lambda\}} f(y_\tau) \right] \right\}$$

I solve the modified problem for a CRRA decision maker: where $f(m) = \frac{m^{1-\gamma}}{1-\gamma}$ (see appendix A.2 for details). In figure 3 I plot the liquidation thresholds for different values of the relative risk aversion coefficient γ. Risk aversion does not lead to a disposition effect. Risk aversion shrinks the inaction region, violating condition (13) as $b > b^*$. Love of risk has the opposite effect, widening the inaction region. Thus standard risk preferences induce behavior that is qualitatively different from the disposition effect. This is a major advantage of my design. I summarize the previous discussion in the following:

Remark 2. *The disposition effect in this stopping problem cannot arise from standard risk preferences.*

2.4 Reference-Dependent Preferences

Standard preferences cannot generate behavior that resembles the disposition effect in this environment, and this is critically important for any clean test of the effect. The disposition

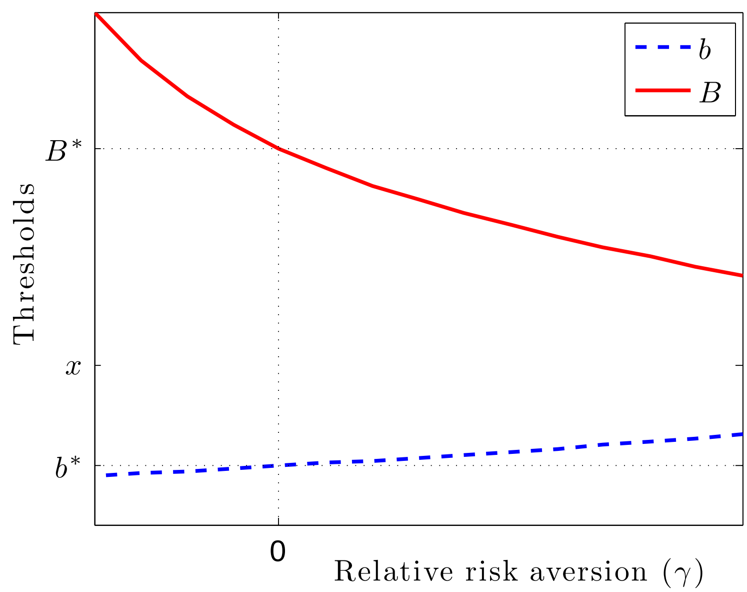

Figure 3: The effect of risk aversion

effect has always been attributed to non-standard, "behavioral" preferences and it is therefore important to look for it in an environment where it cannot be confused with standard economic behavior. In light of the results presented in the previous subsection, I note that a decision maker whose stopping decisions match the disposition effect behaves as if she is risk-seeking when holding a loser but behaves as if she is risk-averse when holding a winner. Indeed, this was the informal explanation of the disposition effect in financial markets offered by the seminal works on this topic (such as Shefrin and Statman [1985], Odean [1998] and Weber and Camerer [1998]).

The notion that a decision maker has asymmetric risk-attitudes is consistent with a long-standing theory of how individuals evaluate risky prospects, namely prospect theory (Kahneman and Tversky [1979] and Tversky and Kahneman [1992], see also the survey of Barberis [2013]). One of the components of prospect theory is an S-shaped utility function, convex over losses and concave over gains, relative to some reference level[10]. This property is also known as diminishing sensitivity because "it implies that, while replacing a $100 gain (or loss) with a $200 gain (or loss) has a significant utility impact, replacing a $1,000 gain (or loss) with a $1,100 gain (or loss) has a smaller impact" (Barberis [2013]). There is much evidence in favour of the diminishing sensitivity hypothesis from static lottery-choice experiments (see for example Tversky and Kahneman [1992], Camerer and Ho [1994] and Wu and Gonzalez [1996]). It is a reasonable assumption for this experiment: ordinary laboratory

[10]Another component of prospect theory is loss aversion. Loss aversion is not sufficient to generate the disposition effect in this task: indeed loss aversion raises the lower threshold because of the presence of expiration risk (as discussed below).

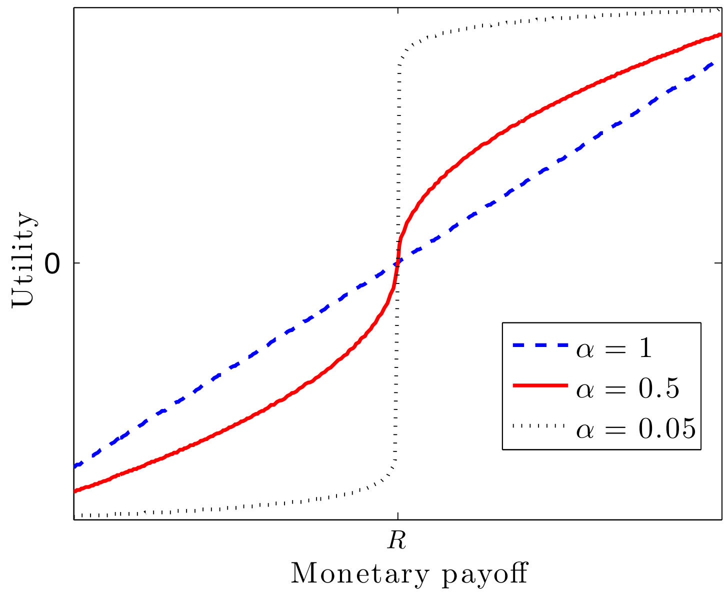

Figure 4: S-shaped utility functions

subjects are likely to experience disappointment from the realization of a loser and elation from the realization of a winner, but disappointment and elation are unlikely to be very sensitive to the size of the loss or gain.

In what follows I set up a model of a decision maker with S-shaped reference-dependent risk preferences and show that it reliably predicts the disposition effect. I assume that the decision maker evaluates a monetary payoff m according to the following S-shaped Bernoulli utility function, standard in the literature on prospect theory:

$$u(m;R) = \begin{cases} (m-R)^{\alpha} & \text{if} \quad m \geq R \\ -(R-m)^{\alpha} & \text{if} \quad m < R \end{cases} \tag{14}$$

Here R is the reference point against which the monetary payoff m is assessed. The parameter $\alpha \in (0,1]$ is the elasticity of utility to gains and losses. When $\alpha < 1$, utility has the property of diminishing sensitivity. Figure 4 illustrates how the parameter α controls the shape of the utility function. A lower α implies stronger risk-aversion over gains and stronger love of risk over losses. The classical estimate of α, obtained by Tversky and Kahneman [1992], is 0.88. Wu and Gonzalez [1996] obtained a lower estimate, around 0.5 (for other estimates see Camerer and Ho [1994], Abdellaoui [2000] and Bruhin et al. [2010]).

The reference point R could be a constant value, for example the reference point could be set equal to the salvage value x. Here I use the notion of an expectation-based reference point, as in the approach originally developed by Bell [1985] and Loomes and Sugden [1986] and recently extended by Kőszegi and Rabin [2007]. One clear advantage of this modeling strategy is that it can be applied to other contexts. In particular, I believe it can provide a microfoundation for the disposition effect in important economic models that involve stopping decisions (as I discuss in Section 5). This notion may also provide a better descriptive model of the actual behavior of subjects in this task. Indeed note that for a decision maker who starts the task at $s_0 = x$ and plans to stop at some future date, the expected monetary payoff is greater than x. For example, this is the case for the optimal risk-neutral policy, which yields $v(x) > x$, as can be seen from Figure 1. It seems plausible that a decision maker with reference-dependent preferences may evaluate the actual realization of the monetary payoff relative to her initial expectation. Thus I assume that, given a pair of stopping times τ_o and τ_c, the reference point is the expected discounted value of the monetary payoffs at the start of the task:

$$R = E\left\{e^{-\rho(\tau-t)}\left[\mathbb{1}_{\{\tau=\tau_o\}}x + \mathbb{1}_{\{\tau=\tau_c\}}s_\tau + \mathbb{1}_{\{\tau=\tau_\lambda\}}y_\tau\right]|s_0\right\} \tag{15}$$

Clearly, this reference point depends on the choice of the two stopping times and in turn on the choice of thresholds. An important question in modeling this type of expectation-based reference-dependent preferences is whether the decision maker is aware of the effect of her choices on her reference point or not. I assume that when choosing the stopping times τ_o and τ_c the decision maker disregards the effect of her choice on the reference-point. In the terminology of Kőszegi and Rabin [2007] the reference point is choice-unacclimating.

Another issue that arises when modeling reference-dependent preferences is how to deal with discounting. The usual interpretation of discounting is that of a time-preference, but in my experiment, as in most experiments on dynamic decisions (e.g., Oprea et al. [2009]), discounting is implemented by expiration risk. It is assumed that at random time, τ_ρ, the investment expires, yielding 0. The expiration time, τ_ρ, is assumed to follow an exponential distribution with parameter equal to the discount factor ρ. Under these conditions, preference-based time-discounting at a rate ρ and expiration risk yield the same utility and behavior in the risk-neutral or risk-averse case. However, under reference-dependent preferences the distinction is important because $u(0; R) \neq 0$: intuitively when the asset expires the decision maker experiences a loss relative to her reference point, but there is no loss when discounting of a future gain is due to time-preferences. In order to provide predictions that apply to behavior in my experiment, I solve the model with expiration risk.

The problem of the decision maker can now be formulated in the following way:

$$(\hat{\tau}_o, \hat{\tau}_c) = \arg\max_{\tau_o,\tau_c} E_t\left[\mathbb{1}_{\{\tau=\tau_o\}}u(x;\hat{R}) + \mathbb{1}_{\{\tau=\tau_c\}}u(s_\tau;\hat{R}) + \mathbb{1}_{\{\tau=\tau_\lambda\}}u(y_\tau;\hat{R}) + \mathbb{1}_{\{\tau=\tau_\rho\}}u(0;\hat{R})\right]$$
$$\hat{R} = E\left\{\left[\mathbb{1}_{\{\tau=\hat{\tau}_o\}}x + \mathbb{1}_{\{\tau=\hat{\tau}_c\}}s_\tau + \mathbb{1}_{\{\tau=\tau_\lambda\}}y_\tau + \mathbb{1}_{\{\tau=\tau_\rho\}}0\right]|s_0\right\}$$

where now $\tau \equiv \min\{\tau_o, \tau_c, \tau_\lambda, \tau_\rho\}$. Details of the numerical solution method are discussed in appendix A.3. Given the optimal stopping times, it is also possible to define the stopping thresholds $\hat{b}$ and $\hat{B}$ by:

$$\hat{\tau}_o = \inf\{t | s_t \leq \hat{b}\}$$
$$\hat{\tau}_c = \inf\{t | s_t \geq \hat{B}\}$$

Figure 5 illustrates how the parameter α affects the solution to the problem. When $\alpha = 1$, the problem reduces to the risk-neutral benchmark. Both the upper and lower liquidation thresholds fall as α decreases. Thus there is a monotonic relation between the elasticity to gains and losses and the disposition effect. The intuition is clearer if the mechanism is analyzed separately for gains and losses. First, a lower elasticity of utility to gains (α) reduces the incentive to wait for a larger gain, thus producing a lower value for the optimal upper threshold B. To understand the effects on the lower threshold, first note that for all values of α the reference point is above x and therefore the salvage value is framed as a loss. So even though liquidating a loser does not entail any nominal loss, it does generate a loss relative to the reference point. The decision maker still chooses to realize losers to avoid incurring larger losses, that may occur, for example, when the asset expires. However, as the elasticity of utility to losses (α) falls, the difference in terms of disutility between realizing a loser and losing the investment decreases, allowing the decision maker to postopone the voluntary realization of losses by choosing a lower b. I summarize the previous discussion in the following:

Remark 3. *S-shaped expectation-based reference-dependent preferences generate the disposition effect.*

3 The Experiment

3.1 Implementing the Model in the Lab

In order to implement the model in the lab, I use a discrete approximation. Each discrete time step or tick has length Δt. I approximate the geometric Brownian motion s_t with the following binomial process:

$$s_{t+\Delta t} = \begin{cases} s_t(1+h) & \text{with probability} \quad p \\ s_t(1-h) & \text{with probability} \quad 1-p \end{cases} \tag{16}$$

As in other dynamic experiments (such as Oprea et al. [2009]), I implement time discounting with random expiration, i.e. termination of the game with no payoff. To approximate the

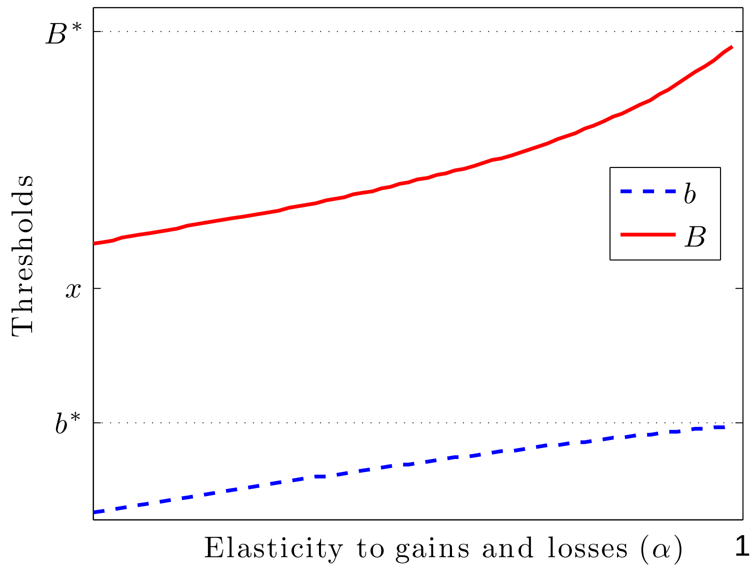

Figure 5: The effect of reference-dependent preferences.

exponential distribution of expiration times with parameter ρ, I use a geometric distribution with parameter r. Similarly the distribution of dividend arrival times, parameterized by λ, is approximated by a geometric distribution with parameter l. The relation between the parameters of the original model and the discrete approximation is discussed in Appendix A.4. Figure 6 summarizes the events that can occur in a tick.

As remarked above, a value of λ different from zero is not necessary to solve the model. However, when $\lambda = 0$, in order to yield a round length of one or two minutes on average, ρ has to be large, and this in turn makes the optimal inaction region narrow. Fixed an average round length, when some of the rounds end in the stock paying out (i.e. $\lambda > 0$) the width of the optimal inaction regions increases. I thus calibrate λ and ρ in order to target a reasonable round duration and a desirable width of the optimal inaction region. Table 1 summarizes the choice of parameter values and the resulting predictions. A tick lasts for 0.2 seconds. The average duration of a round implied by the parameters r and l is around 1.5 minutes. Conditional on termination, the probability that a round expires with no payoff is around 45% and the expected fraction of rounds that end with the arrival of the dividend is 55%. The price process has no drift, so that the uptick probability is $p = 0.5$. The percentage change in the price at each tick is around 5%. The strike price is set at 10, while the optimal inaction region is $(6.2, 17.3)$, and therefore the inaction region is reasonably wide in terms of steps of the price process. Finally the optimal inaction region yields an expected fraction of winners in total sales equal to $\Pi^* = 34\%$.

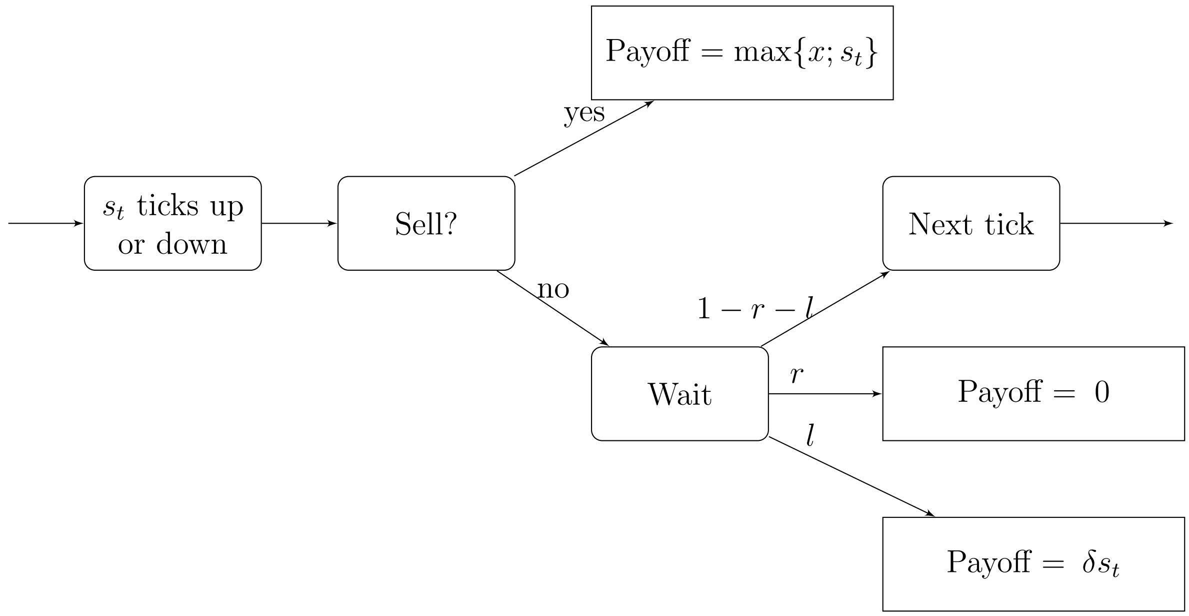

Figure 6: Timeline of a tick.

Continuous				Discrete					Other		Predictions		
μ	σ	ρ	λ	Δt	p	h	r	l	x	δ	b^*	B^*	Π^*
0.00	0.12	0.005	0.006	0.2	0.50	0.054	0.001	0.0012	10	1	6.2	17.3	34%

Table 1: Parameter values and theoretical predictions.
Note: time is measured in seconds.

3.2 Experiment Details

I implemented the experiment using a custom piece of software programmed in a new Javascript environment called Redwood. Each session is divided into 35 rounds, essentially repetitions of the same task with random ending times as discussed in the previous section. Each round has an initial buying stage. At the beginning of each round each subject is given 100 units of experimental cash. In the buying stage the subject can decide to use her cash to buy shares of a stock at a given price per share. Note that the purchase decision is irrelevant for the rational benchmark. I use this procedure in order to be consistent with the existing literature on the disposition effect and to potentially generate a salient reference point. Also note that a rational decision maker will be always willing to pay s_0 for a security that bundles a put option and a stock priced at s_0 (i.e. $v(s_0) > s_0$).

The purchase price is equal to 10 in every round. The subject can choose to spend up to and no more than the 100 units of cash, but each subject has to buy at least one share. Figure 7 illustrates the screen of the buying stage. After the buying decision is made (or after 20 seconds have passed, with 1 share as default choice), the round moves on to the

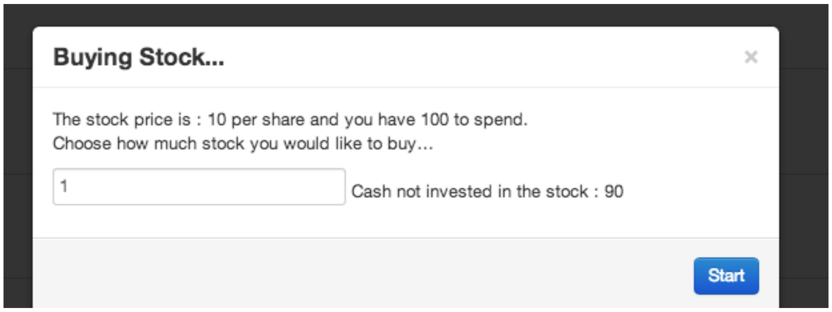

Figure 7: Buying stage.

selling stage.

In the selling stage, the subject display, reproduced in Figure 8 plots the time series of the stock price in real time, with s_0 equal to the purchase price. A subject's only decision is when to press a button labeled "Sell." The choice of whether to cash the stock or exercise the option is automated: whenever the current price is below 10 and the subject decides to sell, she receives the strike price, 10. This feature only eliminates potential noise that is not relevant to the disposition effect. The experiment is run with a semi-strategy method, showing the s_t process up to expiration even after the subject sells. The realization of $\{s_t\}_t$ is determined by the software in real time, so each subject faces different sample paths of the stock price. Similarly, the actual values of the round ending times and of the dividend arrival times are drawn from the same distributions but independently for each subject. After the selling stage is over the computer displays useful summary information about the round, such as the subject's score, her action and whether the round ended with the stock paying out or expiring.

Data was collected in the LEEPS laboratory at the University of California, Santa Cruz between May and October 2013 and in May-June 2014. A total of 108 subjects were drawn from an undergraduate subject pool using students from across the curriculum, recruited using the ORSEE software (Greiner [2004]). Subjects were randomly assigned to visually isolated terminals and interacted with no other subjects during the session. Instructions, reproduced in Online Appendix, were read aloud prior to the beginning of the experiment. Subjects were paid a $5 showup fee and the sum of points earned over all periods converted at an exchange rate of $0.3 for 10 points. Sessions lasted roughly 1 hour and 30 minutes including instructions and subject earnings averaged around $14.

Figure 8: Selling stage.

4 Results

I analyze the experiment results by studying the distribution of stopping points and compare it to the optimal threshold benchmark in order to test for a disposition effect. I leave an analysis of the subjects' purchasing behavior to the Appendix B.1, as it offers no particular insight.

I define a stopping point as a value of the asset price s at which a sale occurs. Figure 9 illustrates the distribution of stopping points pooling all subjects in the study. Letting S represent the random variable that generates stopping points, Panel 11b plots $Prob(S \leq s | S \geq s_0)$, while panel 11a shows $Prob(S \geq s | S \leq s_0)$, as s varies in the relevant range. Confidence bounds at the 99% level are included. The solid vertical lines mark the optimal thresholds, while dashed lines show the actual medians.

The sample medians are $b = 5.14$ and $B = 14.21$. For both thresholds it is possible to reject the hypothesis that they are equal to the optimal level, with a Wilcoxon signed-rank p-value of nearly zero. The aggreagate winner bias is $\Psi = 62\%$. It is possible to conclude the following:

Result 1. *In the aggregate there is evidence of a disposition effect in terms of stopping points: the median lower stopping point is significantly smaller than b^* and the median upper stopping point is significantly smaller than B^*.*

The disposition effect I find in the aggregate data is robust to learning over rounds. To

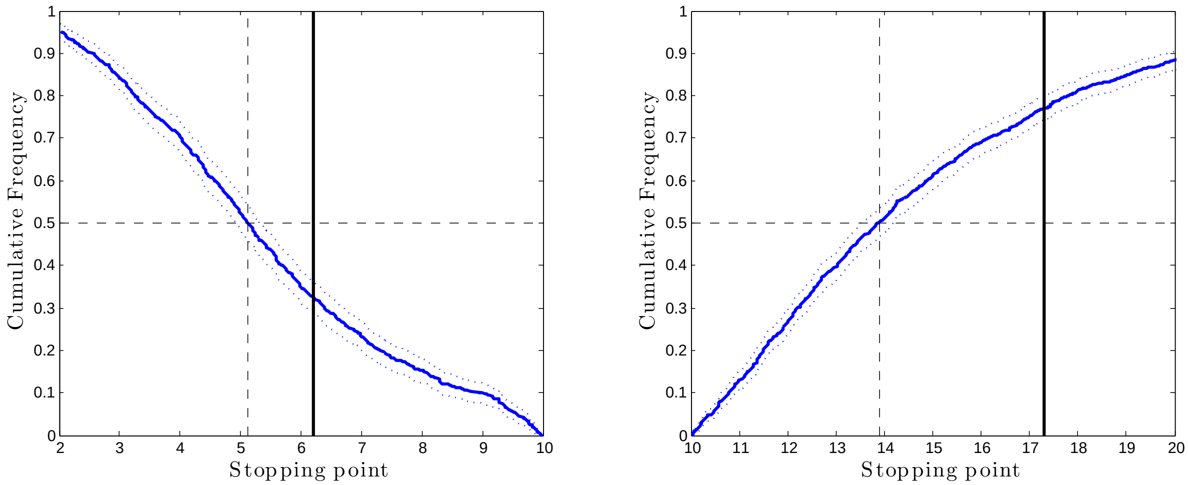

(a) Losers (b) Winners

Figure 9: CDFs of stopping points

show this, I split the whole sample into a sample of early rounds (the first 20) and a subsample of late rounds (the last 15). Figure 10 plots the empirical distributions of stopping points. It is possible to observe that the deviation of the median stopping points (the dashed lines) from the optimal thresholds (the solid lines) does not change much over the experiment (and it is slightly larger in the late rounds sample). Formally, the null-hypothesis that the early and late distributions are equivalent cannot be rejected at standard confidence levels (Kolmogorov-Smirnov p-values: 0.24 and 0.14 for winners and losers respectively).

Another concern is that the empirical distribution of stopping points may be a biased estimate of the underlying process, as random termination of play implies that stopping points farther from s_0 are less frequently observed. Thus the empirical distribution may be biased towards stopping points closer to the initial value. This has different implications for winners and losers. The result that the average stopping point for losers is below the optimal lower threshold is not affected by this concern, as correcting the bias may only strengthen such finding. However, the bias in favor of smaller stopping points could in priciple be the main driver of the finding that the average stopping point for winners is below the optimal lower threshold. Indeed the behavior of a risk seeking individual (with a wide band) may look like the disposition effect if the censoring bias is large. In order to rule out this possibility I create a sub-sample of stopping points for winners by restricting attention to rounds in which, at some time before expiration, the asset price reaches the optimal upper threshold. Similarly for losers, I generate a sub-sample of stopping points from rounds in which, at some time before expiration, the asset price reaches the optimal lower threshold. In these restriced

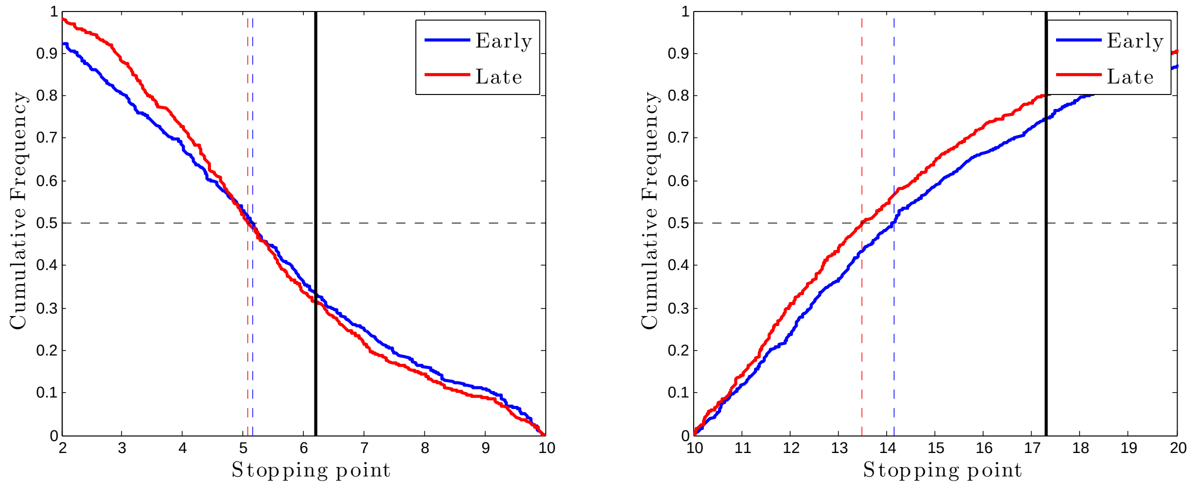

(a) Losers

(b) Winners

Figure 10: CDFs of stopping points in early and late rounds.

samples the censoring bias is eliminated and since the asset price process is exogenous there is no selection bias. This procedure reduces the sample size by 27% for winners and 4% for losers. Figure 11 plots the empirical distributions for the original sample (labeled "All") and the restriced samples. The original and restriced samples of stopping of losers have very similar distributions. There is some evidence of censoring bias in the distribution of stopping decisions for winners, as the median stopping point in the restricted sample is larger than in the original sample. However, even after correcting for censoring the median stopping point in the restricted sample is lower than the optimal threshold. It is still possible to reject the hypothesis that stopping points for winners are clustered around the optimal level, with a Wilcoxon signed-rank p-value of nearly zero.

Result 2. *The finding of an aggregate disposition effect is robust to subject learning and to statistical censoring bias.*

I have shown that the median aggregate stopping points are consistent with the disposition effect condition (13), suggesting that subjects ride losers too long and sell winners too early. However, this is the correct interpretation only insofar as the behavior of subjects is well approximated by threshold rules. In principle a stopping decision at point s may occur much later than the first time the price process has hit s. For example, a subject may observe the price rising from $s = 10$ to $s = 19$, then falling again to $s = 14$ and decide to stop at that point. This is a very different behavior from stopping as soon as the price reaches $s = 14$. Non-threshold behavior will bias tests of the disposition effect based on liquidation points in a precise direction: liquidation points that lie inside the optimal inaction band and

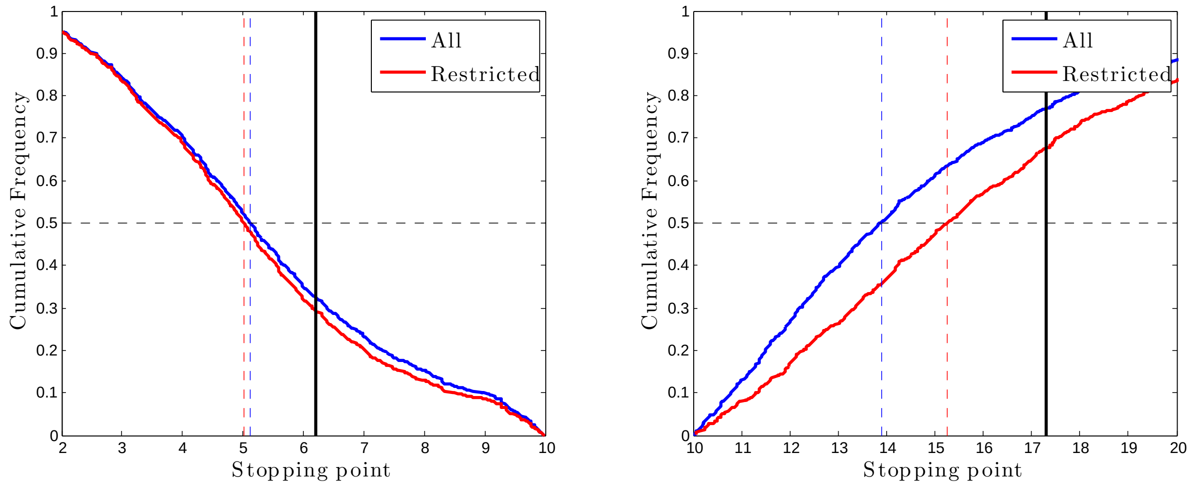

(a) Losers (b) Winners

Figure 11: CDFs of stopping points controlling for censoring bias.

therefore seem to suggest early stopping decisions may in fact represent stopping decisions that happened too late. Thus the result that subjects ride losers too long is robust to non-threshold behavior. On the contrary, the conclusion that subjects tend to sell winners too early may be due to a failure to account for non-threshold behavior. In the current section I address this issue by looking at whether the liquidation decision of a subject occurred before or after the optimal stopping time.[11] For each round in which a stopping decision was made I check whether this decision occurred before or after the optimal stopping time (τ_c^* for winners and τ_o^* for losers). I then report the fraction of stopping decisions that occurred too late, i.e. at $t > \tau_c^*$ for winners and $t > \tau_o^*$ for losers. I compare this to the fraction of stopping decisions that seemed to occur too late based on stopping points, i.e. $s_t > B_c^*$ for winners and $s_t < b^*$ for losers. Figure 12 illustrate the results (for each fraction I also show a two-standard error bar). For both criteria I show the fraction of late stopping decisions in the original sample (labelled "All"), in the sample of late rounds ("Late") and in the sample of late rounds restricted to avoid censoring bias ("Late Restricted"). There is some evidence of non-threshold behavior as using the stopping-time criterion increases the fraction of late decisions with respect to the liquidation-point criterion. However, the disposition effect is robust to this check: the majority of stopping decisions for winners occur too early, while the majority of the stopping decisions for losers occur too late.

Result 3. *Using a stopping-time criterion confirms the finding of a significant disposition effect in the aggregate: on average, subjects tend to sell winners too early and ride losers*

[11]In terms of design, subjects' choices could be restricted to treshold strategies as in Oprea [2012].

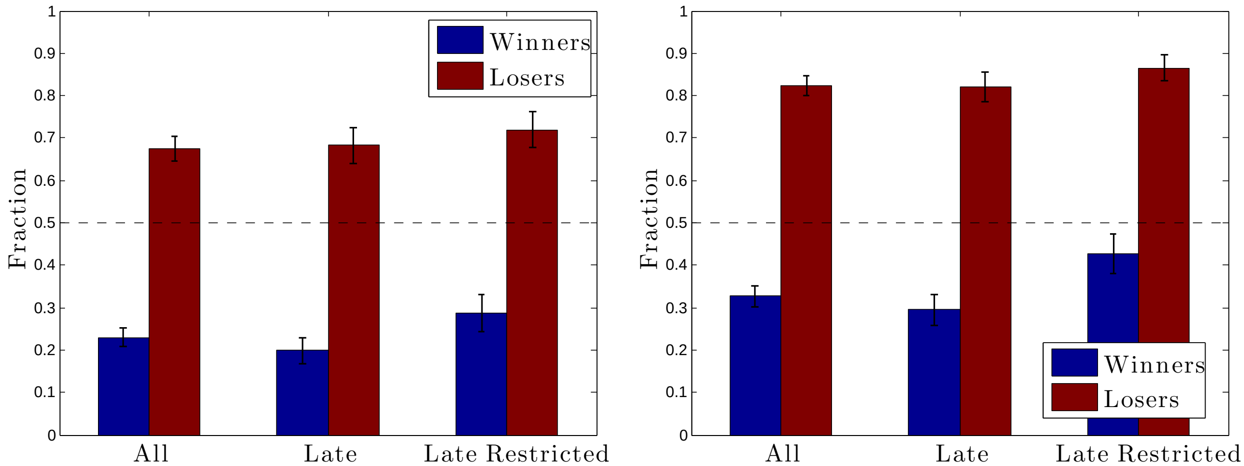

(a) Stopping-point criterion

(b) Stopping-time criterion

Figure 12: Fraction of stopping decisions that occur too late.

too long. The difference in the fraction of stopping decisions that occur later than optimal between winners and losers is highly significant.

In the aggregate 80% of stopping decisions for losers occur later than optimal and similarly the majority of stopping decisions for winners occur earlier than the optimal stopping time. In order to provide further evidence that the aggregate disposition effect reflects a widespread tendency in the subject pool, I look at the individual level. In what follows I use only data from the last 20 rounds for each subject (the results are robust to this choice). I thus define b as the median from the sample of a subject's lower stopping points (and similarly for B). I classify subjects according to the following taxonomy. A subject is classified as risk-neutral if the median lower (upper) stopping points lie within a 10% error band around the optimal lower (upper) threshold (the results are robust to the choice of the error band). A subject's behavior is consistent with the disposition effect ("DE") if the both median stopping points are lower than the respective benchmark (again allowing for an error band). The label "Anti-DE" applies to subjects whose median liquidation points are larger than the respective benchmark. The other two cases are a wider band, labelled "Risk-seeking", and a narrower band relative to the optimal, labelled "Risk-averse". The resulting classification of subjects is summarized in table 2.

For each category, in the first row I show the fraction of subjects. The second and third rows collect the median stopping points for winners and losers respectively. These are medians of the subject-level median points. In the fourth row I show the median of the Ψ measure. Around 62% of the subjects are classified as disposition-effect decision makers.

	Subject type				
	Risk-neutral	DE	Anti-DE	Risk-seeking	Risk-averse
Fraction	3%	62%	0	13%	22%
Median B	16.0	13.6	-	19.6	13.1
Median b	5.8	4.9	-	4.0	6.9
Median Ψ	20%	72%	-	12%	46%

Table 2: Subject classification.

Notably, no subject behaves in the opposite way of the disposition effect. Note that for many subjects Ψ is positive, meaning that the probability of realizing a winner conditional on stopping is larger than optimal even for subjects who are close to the optimal or who adopt a risk-seeking or a risk-averse strategy. However, this statistics is significantly larger for the disposition effect group, with a mean around 72%. The subject-level analysis summarized in table 2 leads to the following:

Result 4. *The behavior of more than half of the subjects is consistent with the disposition effect. Moreover, for this group of subjects the deviation from optimal behavior is particularly significant, as it implies a probability of realizing a winner conditional on stopping 72% larger than optimal.*

5 Discussion

This paper develops a laboratory test of a distinctive prediction of reference-dependent preferences, testable only in dynamic settings: decision makers suboptimally delay realizing disappointing outcomes (procrastination) but suboptimally rush to realize outcomes that are better-than-expected (rushing). Many empirical studies have found evidence suggesting the disposition effect in financial markets - the tendency of individual investors to sell winners too soon and ride losers too long. However, most of the evidence is indirect and using field data it is difficult to establish whether the effect actually reflects an underlying preference.

To address these issues, I develop and conduct a laboratory test based on a model of optimal stopping decisions. The task involves deciding when, if ever, to liquidate a risky asset, whose price follows a stochastic process in continuous time. The payoff structure is designed in such a way that optimal behavior entails stopping as soon as the asset price reaches either an upper threshold or a lower threshold, realizing a winner or a loser respectively.

Because the ranking of stopping payoffs is highly salient, ordinary laboratory subjects are likely to experience disappointment from the realization of a loser and elation from the realization of a winner. In this situation, the principle of diminishing sensitivity (formalized by an S-shaped utility function) predicts subjects will show the disposition effect. My results indicate that indeed most subjects behave as predicted by a model of S-shaped expectation-based reference-dependent preferences. The disposition effect I observe in the lab is economically significant (implying a probability of realizing a winner conditional on stopping 70% larger than optimal), robust to learning and cannot be confounded by standard preferences, such as risk aversion.

This paper shows that S-shaped reference-dependent preferences are a very powerful explanation of behavior in a stylized optimal stopping problem. The disposition effect that I document in my experiment has obvious implications for financial markets but its potential scope is considerably broader. Any microeconomic setting in which optimal behavior takes the form of a two-sided stopping rule with a salient ranking of stopping payoffs is potentially vulnerable to this sort of effect. Here we consider two applications that are quite distinct from financial markets, pertaining to signaling and experimentation respectively.

1. When stochastic information about the value of a privately-informed seller's asset is gradually revealed to a market of buyers, equilibrium involves a no-trade region in the market posterior belief space. As shown by Daley and Green [2012], as soon as the belief reaches the upper threshold, both the low and high types sell at a high price. At the lower threshold, the low type mixes between accepting a low price and waiting. The results of my experiment suggest that most low types would be willing to wait longer before capitulating and try to realize the higher offer sooner than the theory by Daley and Green [2012] predicts. Clearly the effects of reference-dependent preferences depend on equilibrium considerations in this setting, but I conjecture that S-shaped preferences may create an important amplification mechanism of the inefficiency introduced by asymmetric information.

2. In dynamic R&D models, the decision maker collects costly information about the value of a project in continuous time and stops the experimentation process either by abandoning the research or by building the prototype. As shown by Moscarini and Smith [2001] the optimal solution involves an inaction region in the posterior belief space: when beliefs are sufficiently pessimistic the research is abandoned, yielding a null payoff, and when beliefs are sufficiently optimistic the project is built . However, abandoning the project is likely to be viewed as a loss by most individual decision makers, and this can trigger risk-seeking behavior over a range of pessimistic beliefs, thus delaying the abandonment decision. The implications of reference-dependent preferences for the timing of the building decision are less obvious: on one hand diminishing sensitivity induces a lower building threshold (a slope effect), on the other risk-aversion

implies a lower expected utility from builiding as the prototype value is uncertain (a level effect).

Finally, the results of this paper are likely to apply to individual investors in real-world stock markets, especially those who make small investments (in which case diminishing sensitivity to losses is a reasonable assumption). When investors have multiple opportunities to reinvest, the implications of diminishing sensitivity for behavior are less straightforward than in one-time stopping problems. Nonetheless theoretical work on this topic has already shown the importance of S-shaped reference-dependent preferences (Ingersoll and Jin [2013]). Mine is the first paper to produce evidence about this property of utility in a dynamic stochastic experiment[12].

To conclude, the evidence I provide in this paper suggests that S-shaped reference-dependent preferences are important in shaping behavior when the timing of uncertainty resolution is endogenous. Studying how this kind of preferences affect stopping decisions in specific contexts, especially dynamic learning problems, seems a promising avenue for future research.

[12] Frydman et al. [2012] use brain imaging during a laboratory experiment with multiple trading and reinvestment opportunities. They show that individuals experience bursts of utility from sales, even before final monetary payoffs are assigned. However, Frydman et al. [2012] do not address the issue of whether realization utility is S-shaped.

References

Mohammed Abdellaoui. Parameter-free elicitation of utility and probability weighting functions. *Management Science*, 46(11):1497–1512, 2000.

Brad M. Barber and Terrance Odean. The behavior of individual investors. In George M Constantinides, Milton Harris, and René M Stulz, editors, *Handbook of the economics of finance: Financial Market and Asset Pricing*, volume 2B. Elsevier, Amsterdam, 2011.

Nicholas Barberis and Wei Xiong. What drives the disposition effect? an analysis of a long-standing preference-based explanation. *the Journal of Finance*, 64(2):751–784, 2009.

Nicholas Barberis and Wei Xiong. Realization utility. *Journal of Financial Economics*, 104 (2):251–271, 2012.

Nicholas C Barberis. Thirty years of prospect theory in economics: A review and assessment. *The Journal of Economic Perspectives*, 27(1):173–195, 2013.

David E Bell. Disappointment in decision making under uncertainty. *Operations research*, 33(1):1–27, 1985.

Itzhak Ben-David and David Hirshleifer. Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect. *Review of Financial Studies*, 25(8):2485–2532, 2012.

Adrian Bruhin, Helga Fehr-Duda, and Thomas Epper. Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412, 2010.

Colin F Camerer and Teck-Hua Ho. Violations of the betweenness axiom and nonlinearity in probability. *Journal of risk and uncertainty*, 8(2):167–196, 1994.

Brendan Daley and Brett Green. Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504, 2012.

Avinash Dixit. *The Art of Smooth Pasting*. Harwood Academic Publishers, 1993.

Cary Frydman, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel. Using neural data to test a theory of investor behavior: An application to realization utility. Working paper, National Bureau of Economic Research, 2012.

David Genesove and Christopher Mayer. Loss aversion and seller behavior: Evidence from the housing market. *The Quarterly Journal of Economics*, 116(4):1233–1260, 2001.

B. Greiner. The online recruitment system orsee 2.0 - a guide for the organization of experiments in economics. Working Paper Series in Economics 10, University of Cologne, Department of Economics, 2004.

Jonathan E Ingersoll and Lawrence J Jin. Realization utility with reference-dependent preferences. *Review of Financial Studies*, 26(3):723–767, 2013.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291, 1979.

Botond Kőszegi and Matthew Rabin. Reference-dependent risk attitudes. *The American Economic Review*, pages 1047–1073, 2007.

Graham Loomes and Robert Sugden. Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2):271–282, 1986.

Giuseppe Moscarini and Lones Smith. The optimal level of experimentation. *Econometrica*, 69(6):1629–1644, 2001.

Kumar Muthuraman and Sunil Kumar. Solving free-boundary problems with applications in finance. *Foundations and Trends® in Stochastic Systems*, 1(4):259–341, 2008.

Terrance Odean. Are investors reluctant to realize their losses? *The Journal of finance*, 53 (5):1775–1798, 1998.

Ryan Oprea. Survival versus profit maximization in a dynamic stochastic experiment. Technical report, Working Paper, University of British Columbia 2012., 2012.

Ryan Oprea, Daniel Friedman, and Steven T Anderson. Learning to wait: A laboratory investigation. *The Review of Economic Studies*, 76(3):1103–1124, 2009.

Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of finance*, 40(3):777–790, 1985.

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.

Martin Weber and Colin F Camerer. The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization*, 33(2):167–184, 1998.

George Wu and Richard Gonzalez. Curvature of the probability weighting function. *Management science*, 42(12):1676–1690, 1996.

Online Appendices

A Model Details

A.1 Solving the Risk-Neutral Model

Inside the inaction region (b, B) the value function of the risk-neutral decision-maker defined in (3) can be written in a recursive way. Consider a small time interval Δt, then standard arguments give:

$$\begin{aligned}
v(s(t)) &\simeq \frac{\lambda\Delta t}{1+\rho\Delta t}y(t) + \frac{1-\lambda\Delta t}{1+\rho\Delta t}E_t v(s(t+\Delta t)) \\
(1+\rho\Delta t)v(s(t)) &\simeq \lambda\Delta t y(t) + (1-\lambda\Delta t)E_t v(s(t+\Delta t)) \\
\rho\Delta t v(s(t)) &\simeq \lambda\Delta t[y(t) - v(s(t))] + (1-\lambda\Delta t)E_t[v(s(t+\Delta t)) - v(s(t))] \\
\rho v(s(t)) &\simeq \lambda[y(t) - v(s(t))] + (1-\lambda\Delta t)\frac{1}{\Delta t}E_t[v(s(t+\Delta t)) - v(s(t))]
\end{aligned}$$

Taking $\Delta t \to 0$:

$$\rho v(s) = \lambda[y - v(s)] + \frac{1}{dt}E_t dv(z(t)), s \in (b, B) \tag{17}$$

The left hand side can be interpreted as the instantaneous return to holding on to the stock. The first term on the right hand side is the gain or loss realized when a dividend arrives: the decision maker has to give up $v(s)$ and obtains a dividend y. The second term on the right is the expected change in the value of holding on to the stock. Using Ito's lemma leads to the Hamilton-Jacobi-Bellman equation of the problem:

$$(\rho+\lambda)v(s) = \lambda y + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (b, B) \tag{18}$$

I use (2) to express the dividend in terms of the current price and obtain:

$$(\rho+\lambda)v(s) = \lambda\delta s + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (b, B) \tag{19}$$

The general solution of this second order differential equation is:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}s + C_1 s^{R_1} + C_2 s^{R_2} \tag{20}$$

where:

$$\begin{aligned}
R_1 &= \frac{\frac{1}{2}\sigma^2 - \mu + \sqrt{\mu^2 + \frac{1}{4}\sigma^4 - \mu\sigma^2 + 2(\rho+\lambda)\sigma^2}}{\sigma^2} \\
R_2 &= \frac{\frac{1}{2}\sigma^2 - \mu - \sqrt{\mu^2 + \frac{1}{4}\sigma^4 - \mu\sigma^2 + 2(\rho+\lambda)\sigma^2}}{\sigma^2}
\end{aligned}$$

The boundary conditions at the optimal thresholds b^*, B^* are:

$$\lim_{s \searrow b^*} v(s) = x$$
$$\lim_{s \searrow b^*} v'(s) = 0$$
$$\lim_{s \nearrow B^*} v(s) = B^*$$
$$\lim_{s \nearrow B^*} v'(s) = 1$$

Substituting the expression for the value function gives:

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} b^* + C_1 b^{*R_1} + C_2 b^{*R_2} - x = 0 \tag{21}$$
$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1 R_1 b^{*R_1-1} + C_2 R_2 b^{*R_2-1} = 0 \tag{22}$$
$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} B^* + C_1 B^{*R_1} + C_2 B^{*R_2} - B^* = 0 \tag{23}$$
$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1 R_1 B^{*R_1-1} + C_2 R_2 B^{*R_2-1} - 1 = 0 \tag{24}$$

From these four equations it possible to obtain expressions for the two constants C_1, C_2 in terms of B^* and b^*, leading to equations (11) and (12). In order to solve for the two optimal thresholds it is necessary to solve this system of non-linear equations numerically.

A.2 Solution with CRRA Utility

When payoffs are evaluated by means of a CRRA utility function $f(\cdot)$, the model can be solved in a similar way to the risk-neutral case. Equation (18) is replaced by the following Hamilton-Jacobi-Bellman equation:

$$(\rho + \lambda)v(s) = \lambda f(y) + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (b, B) \tag{25}$$

Since the CRRA utility function $f(\cdot)$ is a power function, equation (25) can be solved analytically. The two thresholds can then be found by solving (numerically) the system of equations given by the four boundary conditions:

$$\lim_{s \searrow b^*} v(s) = f(x)$$
$$\lim_{s \searrow b^*} v'(s) = 0$$
$$\lim_{s \nearrow B^*} v(s) = f(B^*)$$
$$\lim_{s \nearrow B^*} v'(s) = f'(B^*)$$

A.3 Solution with Reference-Dependent Preferences

Let $\tau \equiv \min\{\tau_o, \tau_c, \tau_\lambda, \tau_\rho\}$. The problem under reference-dependent preferences is:

$$(\hat{\tau}_o, \hat{\tau}_c) = \arg\max_{\tau_o, \tau_c} E_t \left[\mathbb{1}_{\{\tau=\tau_o\}} u(x; \hat{R}) + \mathbb{1}_{\{\tau=\tau_c\}} u(s_\tau; \hat{R}) + \mathbb{1}_{\{\tau=\tau_\lambda\}} u(y_\tau; \hat{R}) + \mathbb{1}_{\{\tau=\tau_\rho\}} u(0; \hat{R}) \right]$$

$$\hat{R} = E \left\{ \left[\mathbb{1}_{\{\tau=\hat{\tau}_o\}} x + \mathbb{1}_{\{\tau=\hat{\tau}_c\}} s_\tau + \mathbb{1}_{\{\tau=\tau_\lambda\}} y_\tau + \mathbb{1}_{\{\tau=\tau_\rho\}} 0 \right] | s_0 \right\}$$

Solving this problem involves dealing with a number of issues related to the S-shape of the function $u(\cdot; R)$, defined in equation (14), and to the fact that the reference point is endogenous. Since the function $u(\cdot; R)$ is not differentiable at R, the problem needs to be solved separately for gains and for losses. Let $v_L(\cdot)$ be the value function in the loss domain and $v_G(\cdot)$ be the value function in the gain domain. At the two optimal thresholds $\hat{b}, \hat{B}$ value matching and smooth pasting apply. Additionally, the two value functions must satisfy value matching and smooth pasting at the reference point, R. Proceeding as before, the value functions of the problem can be written in terms of a Hamilton-Jacobi-Bellman equation:

$$(\rho + \lambda) v_L(s) = \rho u(0; \hat{R}) + \lambda u(y; \hat{R}) + \mu s v_L'(s) + \frac{1}{2} \sigma^2 s^2 v_L''(s), s \in (b, \hat{R}) \tag{26}$$

$$(\rho + \lambda) v_G(s) = \rho u(0; \hat{R}) + \lambda u(y; \hat{R}) + \mu s v_G'(s) + \frac{1}{2} \sigma^2 s^2 v_G''(s), s \in (\hat{R}, B) \tag{27}$$

The boundary conditions at the optimal thresholds $\hat{b}, \hat{B}$ and at the reference point are:

$$\lim_{s \searrow \hat{b}} v_L(s) = u(x; \hat{R}) \tag{28}$$

$$\lim_{s \searrow \hat{b}} v_L'(s) = 0 \tag{29}$$

$$\lim_{s \nearrow \hat{B}} v(s)_G = u(\hat{B}; \hat{R}) \tag{30}$$

$$\lim_{s \nearrow \hat{B}} v'(s)_G = u_1(\hat{B}; \hat{R}) \tag{31}$$

$$\lim_{s \nearrow \hat{R}} v_L(s) = \lim_{s \searrow \hat{R}} v_G(s) \tag{32}$$

$$\lim_{s \nearrow \hat{R}} v_L'(s) = \lim_{s \searrow \hat{R}} v_G'(s) \tag{33}$$

Since $u(\cdot; R)$ is not simply a power function, these differential equations cannot be solved analytically. Additionally, they cannot be solved with standard numerical methods, because this is a free-boundary problem. In order to solve them I use the method described in Muthuraman and Kumar [2008]. This method involves starting with arbitrary boundaries, solving the differential equation numerically using the value matching conditions (equations (28), (30) and (32)) and then adjusting the boundaries to ensure that also smooth pasting is satisfied (equations (29), (31) and (33)). This loop is repeated until the boundaries and the value function converge. To this algorithm I add the computation of the endogenous reference point $\hat{R}$. First note that $\hat{R}$ is the value function of a risk-neutral decision maker at s_0. To solve for this value function I use the two value-matching conditions at the given thresholds $\hat{b}, \hat{B}$. Using the method described in A.1, we have a system of three equations in three unknowns $\hat{R}, K_1, K_2$, where K_1 and K_2 are constants that can be obtained by the two value-matching conditions. The three equations are:

$$\hat{R} = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} s_0 + K_1 s_0^{R_1} + K_2 s_0^{R_2} \tag{34}$$

$$x = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} \hat{b} + K_1 \hat{b}^{R_1} + K_2 \hat{b}^{R_2} \tag{35}$$

$$\hat{B} = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} \hat{B} + K_1 \hat{B}^{R_1} + K_2 \hat{B}^{R_2} \tag{36}$$

This computation is nested in each loop of the algorithm.

A.4 Discrete Approximation

The continuous time and discrete time parameters are related by the following standard conditions (see Dixit [1993]):

$$\mu = \frac{(2p-1)h}{\Delta t}$$
$$\sigma^2 = \frac{4p(1-p)h^2}{\Delta t}$$
$$\rho = \frac{-\ln(1-r)}{\Delta t}$$
$$\lambda = \frac{-\ln(1-l)}{\Delta t}$$

B Details of the Empirical Analysis

B.1 Purchasing Decisions

Figure 13 shows the histogram of the number of shares bought in a round by a subject. Only a small fraction of buying decisions involve the one share minimum requirement. The number of shares bought did not seem to be related to behavior in any significant way.

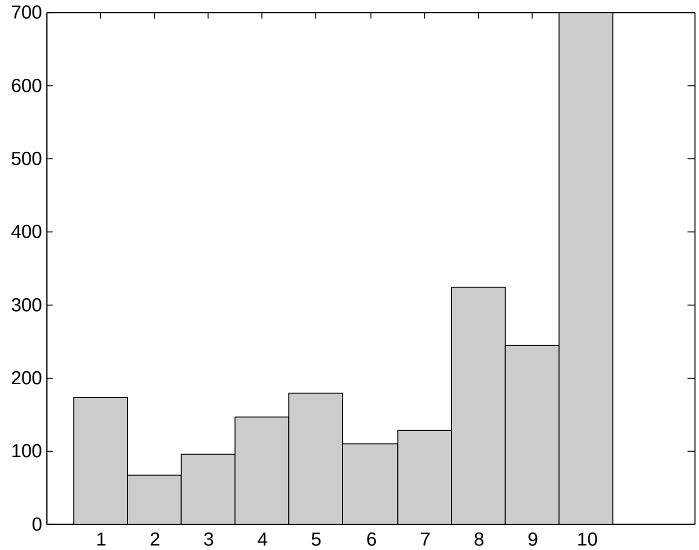

Figure 13: The distribution of the number of shares bought