

The Disposition Effect and Realization Preferences: a Direct Test*

Jacopo Magnani[†]

March 7, 2014

Abstract

This paper develops a new laboratory test of the hypothesis that individual investors have an irrational preference for selling winning stocks vis-à-vis selling losing stocks. In the experiment, subjects invest in a security that bundles a risky asset, whose price evolves in near-continuous time, with a perpetual put option. Optimal behavior is characterized by an upper and a lower selling thresholds in the asset price space, thus producing a clear rational benchmark and eliminating known confounds. Subjects indeed tend to delay selling losers beyond the optimal point and to sell winners before reaching the optimal liquidation point. The median liquidation points imply the probability of realizing a gain conditional on a sale is 56% larger than optimal. Such behavior is shown to be consistent with a realization utility model and structural estimates reveal that the sensitivity of realization utility to gains and losses decreases faster than what is implied by canonical estimates of prospect-theoretic value functions. A direct estimate of the degree of diminishing sensitivity is an important input for behavioral finance theory as even qualitative results of realization utility models, such as whether investors voluntarily realize losses or not, depend on the value of the sensitivity parameter.

Keywords: Laboratory experiments, disposition effect, realization utility, reference-dependent preferences, option exercise.

JEL codes: G02 D81 C91

*I am grateful to Dan Friedman and Ryan Oprea for guidance and support, to Matthew Browne for developing the experiment software, to the participants of the WEAI graduate student workshop for useful comments. I acknowledge funding from the National Science Foundation Grant SES 1260813.

[†]Economics Department, University of California, Santa Cruz, CA 95064. Email: jacopo.magnani@gmail.com

1 Introduction

The disposition effect - the tendency of individual investors to sell assets whose price has increased and keep assets that have dropped in value - is a cornerstone of behavioral finance. It exemplifies the kind of anomalies that we may expect from the trading behavior of unsophisticated individuals in a large and growing number of asset markets.¹ However, previous evidence has been mostly indirect. For instance, the typical procedure used in field data studies involves testing whether the frequency of sales is larger for winners than for losers and trying to eliminate competing explanations for this fact, such as portfolio rebalancing and speculative motives.

I propose and conduct a sharper test of the disposition effect using a laboratory experiment that produces an unambiguous benchmark of rational selling for both losers and winners. My experiment is informed by a model of liquidation decisions that I build explicitly for studying the disposition effect. I design a security that bundles a risky asset, whose price follows a stochastic process in continuous time, with a perpetual put option. I analyze an impatient investor 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 investor reaps this opportunity, or
2. the asset price reaches a lower threshold b^* and the investor 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 liquidation 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 an irrational preference for selling winning stocks vis-à-vis selling losing stocks. 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 capitulating. This departure from rational behavior is economically significant as the median liquidation points imply a probability of realizing a

¹Evidence 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.

gain conditional on a sale 56% larger than optimal. The relatively simple environment of my experiment also allows me to structurally estimate a microfounded model of the disposition effect at the individual level. Structural estimates confirm that for most subjects the utility weight on gains and losses is large and reveal other important properties of the subjects' preferences for realizing winners.

Laboratory research provides an important complement to field data studies on the disposition effect. Obtaining and skillfully analyzing field evidence is a necessary step for understanding the behavior of individual investors, however important aspects of the decision making process are unobservable in naturally occurring settings. It is often difficult to identify a normative benchmark of trading against which the disposition effect can be measured. In practice, field data studies rely on the implicit assumption that rational behavior implies symmetry between sales of winners and sales of losers. Thus, the typical test of the disposition effect involves checking whether sales of winners are more likely than sales of losers.² However, these tests are not grounded in any specific theory and the results may be subject to different interpretations.³ In my experiment I define an unambiguous benchmark and eliminate confounding factors by design. My experimental design features buying and selling decisions, near-continuous time asset price processes and rules for random termination of play that implement time-discounting.⁴ I use this framework to provide a direct test of individuals' aversion to realize losses and preference to realize gains. I believe that this set of tools will also prove valuable to researchers in studying other aspects of the (partial equilibrium) dynamic behavior of unsophisticated investors in the lab, such as momentum in buying decisions.

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 sales: while in Weber and Camerer (1998) sales arise from portfolio choice motives, here sales are induced

²While 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)).

³Researchers 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.

⁴The laboratory environment that is closest to my experiment is that of Oprea, Friedman and Anderson (2009). While Oprea et al. (2009) study the decision to start an investment project, I design an experiment where an asset can be bought and sold at an exogenously changing price.

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 agents are actively engaged in portfolio choice, risk aversion and other diversification motives may lead to differences in the investors' 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 stock that the agent 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 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 ($b > b^*$ & $B < B^*$) because the investor prefers less payoff variability: a risk-averse investor thus liquidates losers sooner than optimal. In contrast, the disposition effect implies the investor holds on to losers for too long.

This paper provides empirical content to recent theoretical investigations on micro-founded models of the disposition effect. A theoretical explanation for the disposition effect that has recently been put forward is realization utility, i.e. the idea that investors obtain utility from realizing gains and disutility from realizing losses.^{7,8} Important properties of the selling behavior predicted by these models depend on the shape of the realization utility function. Ingersoll and Jin (2013) study the implications of an S-shaped realization utility function that is concave over gains but convex over losses, a property also known as diminishing sensitivity. Ingersoll and Jin (2013) show that investors who experience realization utility voluntarily sell losers only when the sensitivity of realization utility is diminishing and decreases sufficiently fast in the size of realized gains and losses. Since even qualitative results of realization utility models depend on the value of the sensitivity parameter, precise

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

⁶The 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 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.

⁷Seminal 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 risk aversion over gains and love of risk over losses as formalized by prospect theory. However, Barberis and Xiong (2009) show that in a fully dynamic model of trading prospect-theoretic risk-preferences do not necessarily lead to a tendency to realize gains more readily than losses and can even have the opposite effect.

⁸Barberis and Xiong (2012) explore the implications of this assumption in a dynamic model of reinvestment where an investor obtains utility proportional to the realized gains and losses. Frydman, Barberis, Camerer, Bossaerts and Rangel (2012) provide evidence that investors get a burst of utility when they realize gains by using brain imaging (fMRI) during an experiment based on the design of Weber and Camerer (1998).

evidence on this property of utility is necessary. The simple framework of my experiment allows me to produce such evidence.

In order to structurally estimate the sensitivity to gains and losses, I start from a rational benchmark of an investor who cares only about consumption utility from liquidation and then I develop a model of an investor who additionally experiences realization utility.⁹ I show that such a hybrid model predicts the disposition effect relative to the benchmark. Thus the results of the experiment are qualitatively consistent with the realization utility explanation. Then I use the evidence to make inference about the shape of the realization utility function. I show that in my model when sensitivity diminishes faster, it is possible to jointly obtain loss-taking and a stronger disposition effect. The intuition is that when sensitivity diminishes faster the investor cares more about the relative probability with which gains and losses are realized rather than the size of these gains and losses. This is the key identification assumption. I estimate a discretized version of the model by numerical maximum likelihood using tick-level data on inaction and liquidation decisions. I find that the sensitivity to realized gains and losses is diminishing and decreases faster than what is implied by canonical estimates of prospect-theoretic value functions. Given the different role that the sensitivity parameter plays in static lottery choice tasks and in dynamic realization utility models¹⁰, it seems important to calibrate realization utility models from experiments that more closely resemble their theoretical environments, such as the one presented here.

The paper is organized as follows. In the next section I provide a brief theoretical description of the environment I implement in the lab. I first describe the optimal strategies of a risk-neutral investor who is impatient, but otherwise rational. I then discuss how to model the disposition effect in this setting, together with risk- and realization-preferences. Section 3 describes the experimental design. Section 4 presents the results and section 5 concludes.

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 (Ω, \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

⁹In both Ingersoll and Jin (2013) and Barberis and Xiong (2012) the investor obtains only psychological realization utility when he sells an asset, due to continuous reinvestment.

¹⁰As pointed out by Ingersoll and Jin (2013), while static models suggest that an S-shaped utility function is responsible for the disposition effect, in dynamic realization utility models diminishing sensitivity actually reduces the disposition effect, explaining why any voluntary losses are realized rather than none.

Brownian motion:

$$ds_t = \mu s_t dt + \sigma s_t dw_t \quad (1)$$

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:

$$y_t = \delta s_t \quad (2)$$

where δ is an exogenous parameter¹¹.

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 investor who is endowed with a unit of the stock and a perpetual put option. At each point in time, the investor 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 a put option which allows him to sell the stock for a strike price x . Let τ_o be the time at which the investor exercises the put 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 investor obtains the dividend according to equation (2). Let

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

Then the problem of the investor is described by the following value function:

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

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

The solution to this optimal stopping problem is characterized by an inaction region (b, B) .¹² As soon as s reaches B the investor sells (cashes the stock): $\tau_c = \inf\{t : s_t \geq B\}$. As soon as s reaches b the investor sells by exercising the put option: $\tau_o = \inf\{t : s_t \leq b\}$. Then the value function satisfies:

$$v(s) = s, \forall s \geq B \quad (4)$$

$$v(s) = x, \forall s \leq b \quad (5)$$

¹¹In 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.

¹²The random arrival of a dividend is not necessary in order for the solution to be described by an inaction region (while time discounting is necessary). However, the random dividend process plays an important role in the experimental design, as discussed below.

Inside the inaction region, the value function is given by:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}s + C_1s^{R_1} + C_2s^{R_2}, s \in (b, B) \quad (6)$$

where R_1 and R_2 are algebraic functions of the parameters (given in Appendix A.1), while C_1 and C_2 are constants to be determined.

I denote the optimal 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:

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

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

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

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

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 relations:

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

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

The two optimal thresholds can then be found by standard numerical methods. The value function of the investor'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 immediate reward from selling the asset and the investor 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 investor optimally capitulates, forgoing potential future price increases. In a static setting (or when $\sigma \rightarrow 0$ or $\rho \rightarrow \infty$), the problem reduces to a standard protective put strategy with the following optimal liquidation rule: cash the stock if $s > x$, exercise the put if $s < x$. For a full description of the comparative statics of the model see Appendix A.2.

Remark 1. *Optimal behavior in the problem of liquidating a risky asset with a constant*

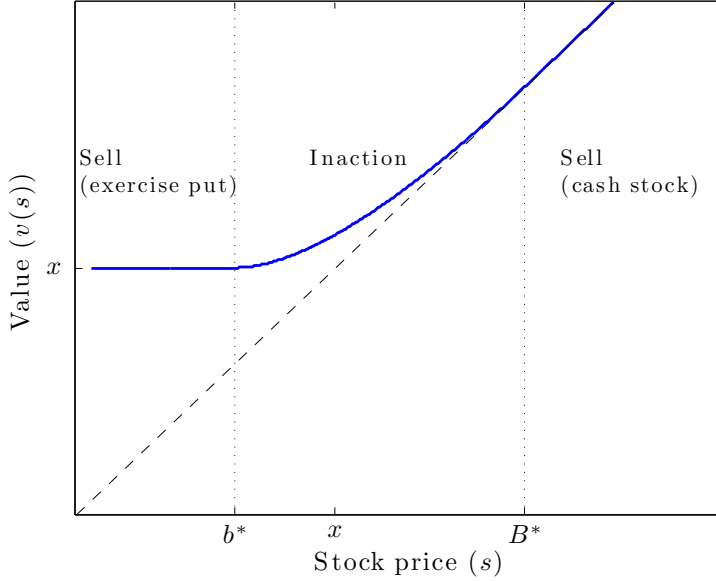


Figure 1: Value function and inaction region of the problem

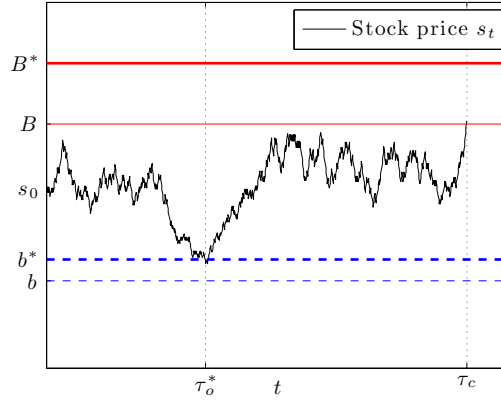
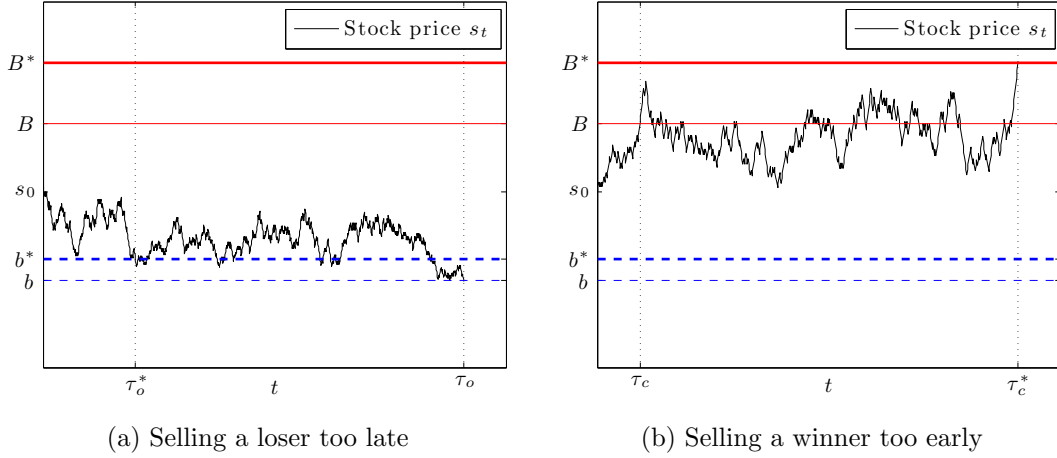
salvage value involves holding on to the investment until the asset price reaches either an upper threshold B^* or a lower threshold b^* . The two thresholds are given by equations (7),(8),(9),(10).

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 gains on winning stocks too early, where gains and losses defined relatively to the purchase price s_0 .¹³

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 denote by B (b) the actual upper (lower) liquidation threshold used by the investor. Panels 2a and 2b illustrates how the classical definition of the disposition effect applies in this setting: in the former the investor liquidates a loser too late ($\tau_o > \tau_o^*$), while in the latter he sells a winner too early ($\tau_c < \tau_c^*$). These two examples also show that the disposition effect leads to “small gains and big losses”, at least relative to the optimal benchmark. Panel 2c illustrates the notion that the disposition effect leads to realizing gains more frequently and losses more rarely than optimal. The situation illustrated in Figure 2 is summarized in the following:

¹³In the following discussion I will assume that the purchase prices coincides with the put strike price, as in the lab implementation of the model. This implies that an agent exercises the put option if and only if she sells a loser.



(c) Selling a winner instead of a loser

Figure 2: Disposition Effects

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

$$b < b^* \wedge B < B^* \quad (11)$$

I introduce a measure of the strength of the disposition effect in this setting. First, I compute the probability of realizing a gain conditional on a sale (or the expected fraction of sales that are sales of a winner). This is equal to the probability that s_t hits B rather than b , conditional on s_t hitting one of the two selling 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 gain conditional on a sale induced by the actual and optimal liquidation thresholds respectively. The intensity of the disposition effect, Ψ , is defined as the increase in the probability of realizing a gain conditional on a sale relative to the optimal benchmark:

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

I remark that Ψ is not used to establish the existence of the disposition effect, but to gauge its economic significance and compare different patterns of behavior within the class of those that satisfy the disposition effect condition (11).

2.3 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. I solve the modified problem for a CRRA investor:

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

where $u(m) = \frac{m^{1-\gamma}}{1-\gamma}$. 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 (11) 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 the stock liquidation problem cannot arise from standard risk preferences.*

2.4 Realization 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 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. Early literature posited that the effect is an outgrowth of asymmetric risk attitudes for gains and losses as formalized by prospect theory. A more recent and robust explanation is that investors may have a direct preference for selling winners as opposed to losers. In the words of Barberis and Xiong (2012), “an investor feels good when he sells a stock at a gain because, by selling, he is creating what he views as a positive investing




Figure 3: The effect of risk aversion

episode. Conversely, he feels bad when he sells a stock at a loss because, by selling, he is creating what he views as a negative investing episode”. Realization preferences of this kind may arise as a simple heuristic that however fails to maximize expected discounted consumption utility, similar to the survival bias studied by Oprea (2012). Unsophisticated agents struggle to associate psychologically unpleasant events (such as failures to survive or loss taking) with optimality and in dynamic stochastic environments they adjust their behavior in order to reduce the frequency with which such events endogenously occur. The realization preference explanation posits a bias that is close to a broad notion of loss-aversion or the sunk-cost fallacy (inasmuch as the purchase price s_0 is treated as a reference point).

I therefore introduce realization preferences in the current framework. I assume that when the investor liquidates the asset he obtains a jolt of utility that depends on the price change since purchase ($\frac{s_t - s_0}{s_0}$). As in Ingersoll and Jin (2013), I adopt an S-shaped realization utility function, with diminishing sensitivity:

$$ru(s_t) = \begin{cases} \beta \left(\frac{s_t - s_0}{s_0} \right)^\alpha & \text{if } \frac{s_t - s_0}{s_0} \geq 0 \\ -\beta \left(-\frac{s_t - s_0}{s_0} \right)^\alpha & \text{if } \frac{s_t - s_0}{s_0} < 0 \end{cases} \quad (12)$$

The parameter $\alpha \in (0, 1]$ is the elasticity of realization utility to the realized return. When $\alpha < 1$, realization utility has the property of diminishing sensitivity. This is a plausible assumption as an investor is likely to feel the difference between a small return and a large return more than the difference between two relatively large returns. The linear case, studied by Barberis and Xiong (2012), obtains when $\alpha = 1$. Figure 4 plots the shape of the realization

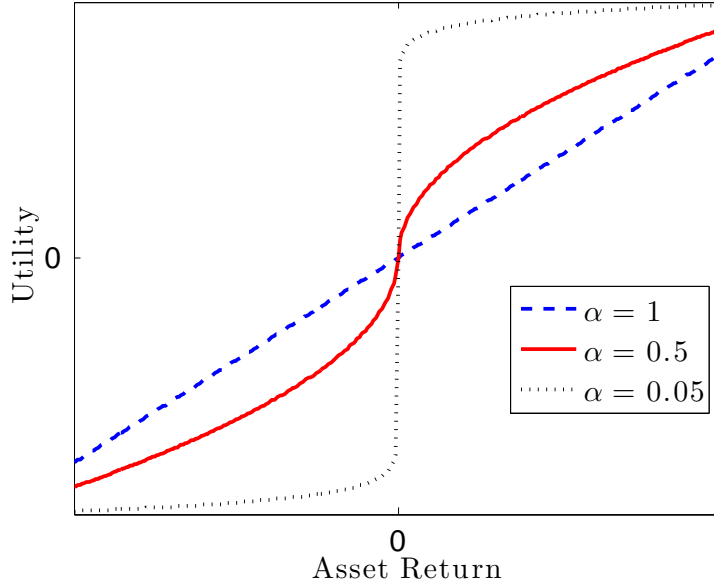


Figure 4: The Shape of the Realization Utility Function

utility function in three cases: the linear case $\alpha = 1$, $\alpha = 0.5$ and $\alpha = 0.05$.

Similarly to Barberis and Xiong (2012), a single parameter, $\beta \geq 0$, determines the importance of realization utility relative to consumption utility. Here β is the marginal rate of substitution between consumption utility and realization utility at a sale. I assume that there is no realization utility when the stock pays out - in other words the psychology of realization utility applies only to voluntary acts¹⁴. This assumption is the same used in Barberis and Xiong (2012), where the asset's dividend yields only consumption utility. The investor's problem then becomes:

$$v(s_t) = \max_{\tau_o, \tau_c} E_t \left\{ e^{-\rho(\tau-t)} \left[\mathbb{1}_{\{\tau=\tau_o\}}(x + ru(s_\tau)) + \mathbb{1}_{\{\tau=\tau_c\}}(s_\tau + ru(s_\tau)) + \mathbb{1}_{\{\tau=\tau_\lambda\}}y_\tau \right] \right\} \quad (13)$$

It is now possible to study the effect of realization utility on the liquidation thresholds. First, for large values of the realization utility weight β the solution to equation (13) does not admit a lower threshold, that is the agent chooses to never realize losses. This is illustrated in Figure 5, where I plot the value of the investment at purchase ($v(s_0)$) as a function of the lower threshold. I consider three cases, $\beta = 5$, $\beta = 10$ and $\beta = 12$ (here I fix the upper threshold and $\alpha = 1$). There is an interior solution for $\beta = 5$ and $\beta = 10$, as value maximization is achieved for $b > 0$. However as β increases the optimal lower threshold discontinuously falls to zero. In the case $\beta = 12$ the maximum is achieved when the lower

¹⁴This is a convenient assumption as it guarantees that the value function in the inaction region has a closed form solution. Indeed the value function is still of the form of equation (6) and the only changes are in the boundary conditions.

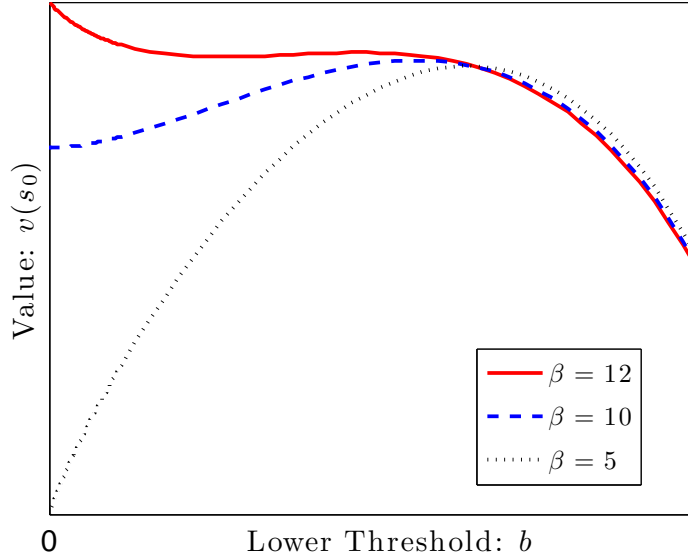


Figure 5: Realization preferences and the lower threshold




Figure 6: Realization utility and the disposition effect

threshold is set to zero. When the lower threshold is zero, no losers are ever realized.

In Figure 6 I plot the optimal thresholds under realization utility as the parameter β varies in the range of values over which the solution admits a lower threshold. I also show how the disposition effect depends on the utility sensitivity parameter α by looking at three cases: the linear case $\alpha = 1$, $\alpha = 0.5$ and $\alpha = 0.05$. The numerical results show that realization preferences lower both the upper thresholds and the lower threshold.

In Figure 7 I illustrate how the weight β affects the intensity of the disposition effect in the realization utility models. I plot the intensity of the disposition effect Ψ as the parameter β varies in the range of values over which the solution admits a lower threshold. The intensity of the disposition effect is monotonically increasing in β .

I summarize the previous discussion in the following:




Figure 7: The intensity of the disposition effect

Remark 3. *Realization preferences produce a disposition effect in the stock liquidation problem. Fixed α , the intensity of the disposition effect is increasing in the realization utility weight (β). Fixed α , for sufficiently large values of the realization utility weight (β) the investor never realizes losers.*

While the realization utility weight β is directly related to whether the investor has an irrational preference for selling winners or not, the sensitivity parameter α also plays a key role in shaping the investor's behavior. In realization utility models there exists an important relation between the sensitivity parameter α and voluntary realization of losses. In the model of Barberis and Xiong (2012) investors reinvest their wealth after each sale, obtain no consumption utility from liquidation and experience linear realization utility. This model generates an extreme form of the disposition effect: the investor sells winners when the gain is sufficiently large but he sells a loser only when he is forced to do so by an exogenous liquidity shock. The model of Ingersoll and Jin (2013) has a similar setup of Barberis and Xiong (2012), but considers an S-shaped realization utility function. Ingersoll and Jin (2013) prove that diminishing sensitivity ($\alpha < 1$) is necessary for an agent with realization utility to voluntarily sell losers in their reinvestment model. The environment described in this paper differs from these two models as liquidation yields consumption utility (a necessary element of an incentivized experiment). Thus the conclusion of Ingersoll and Jin (2013) does not fully hold: Panel 6a shows that the investor chooses to liquidate losers even when the sensitivity of realization utility is constant ($\alpha = 1$) if the weight on realization utility is not too large. However, it is true that with linear realization utility only a weak disposition effect is consistent with voluntary loss-taking. More generally, low values of α are necessary to obtain a strong disposition effect while at the same time having voluntary realization of losses. This can be inferred by inspecting Figures 6 and 7: a model with lower α is able to generate a stronger disposition effect within the region of the parameter space that admits a lower threshold.

Remark 4. *The maximum value of the intensity of the disposition effect consistent with voluntary loss-taking is decreasing in the elasticity of realization utility.*

This property is an extension of the argument made by Ingersoll and Jin (2013) to a model that combines consumption and realization utility. To provide an intuitive explanation of this result, first note that the choice of threshold values has two effects: 1) it affects the probability of realizing a gain vs. a loss, 2) it determines the final payoff when one of the two threshold is reached. For an investor who experiences realization utility, lowering b decreases the probability of having to realize a loss but at the same time it increases the psychological pain of liquidating a loser when the price hits b (since a lower b implies a larger loss). When α is small, however, the negative realization utility from liquidating a loser does not increase much as b is pushed below b^* , thus allowing for a larger deviation from the rational benchmark. At the same time, the positive realization utility from selling a winner does not fall much as B is pushed below B^* , while this allows to enjoy the pleasure of realizing a gain with higher probability. When sensitivity diminishes faster the investor cares more about the relative probability with which gains and losses are realized rather than the size of these gains and losses.

The ability of the realization utility model to generate a disposition effect consistent with voluntary loss-taking is important. Field data shows that individual investors do indeed sell losers, even though more rarely than winners. Moreover, never realizing losers is particularly costly when liquidation yields consumption in addition to psychological utility. Therefore the sensitivity parameter α is a key element of the realization utility model.

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:

$$s_{t+\Delta t} = \begin{cases} s_t(1+h) & \text{with probability } p \\ s_t(1-h) & \text{with probability } 1-p \end{cases} \quad (14)$$

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 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




Figure 8: Timeline of a tick.

parameters of the original model and the discrete approximation is discussed in Appendix A.3. Figure 8 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 $\Pi^* = 34\%$.

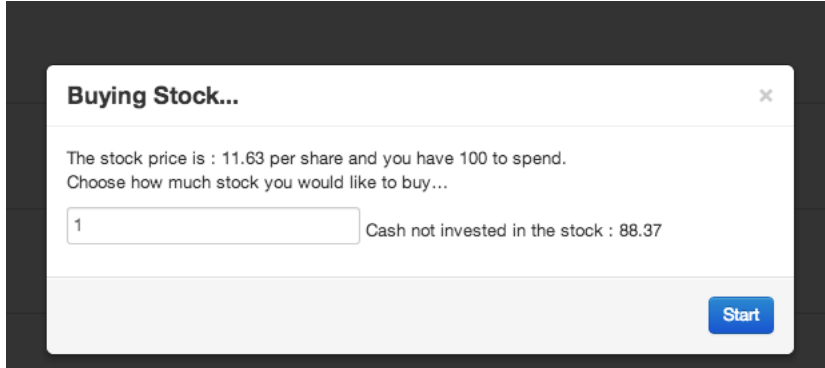


Figure 9: Buying stage.

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 investor 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 9 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 selling stage.

In the selling stage, the subject display, reproduced in Figure 10 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



Figure 10: Selling stage.

the put 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 from May to October 2013. A total of 62 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 \$0.003 for each point earned over all periods. Sessions lasted roughly 1 hour and 30 minutes including instructions and subject earnings averaged around \$14.




Figure 11: CDFs of liquidation points

4 Results

I analyze the experiment results in two subsections. First, I study the distribution of liquidation points and compare it to the optimal threshold benchmark in order to test for a disposition effect. Then I estimate a model of realization utility and discuss the structural results. Other aspects of the data can be found in the Appendix. I leave an analysis of the subjects' purchasing behavior to the Appendix B.1, as it offers no particular insight (and most of the subjects in most of the rounds invested in more than one share, the required minimum, see Figure 17).

4.1 Tests of the Disposition Effect using Liquidation Data

I define a liquidation point as a value of the asset price s at which a sale occurs. Figure 11 illustrates the distribution of liquidation points pooling all subjects in the study. Letting S represent the random variable that generates liquidation points, Panel 13b plots $Prob(S \leq s | S \geq s_0)$, while panel 13a 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 associated measure of the intensity of the disposition effect is $\Psi = 56\%$. It is possible to conclude the following:




Figure 12: CDFs of liquidation points

Result 1. *In the aggregate there is evidence of a disposition effect in terms of liquidation points: the median lower liquidation point is significantly smaller than b^* and the median upper liquidation point is significantly smaller than B^* . This difference is also economically significant as it implies a probability of realizing a gain conditional on a sale 56% larger than optimal.*

The disposition effect I find in the aggregate data is robust to learning over rounds. To 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 12 plots the empirical distributions of liquidation points. It is possible to observe that the deviation of the median liquidation 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 liquidation points may be a biased estimate of the underlying process, as random termination of play implies that liquidation points farther from s_0 are less frequently observed. Thus the empirical distribution may be biased towards liquidation points closer to the initial value. This has different implications for winners and losers. The result that the average liquidation 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 liquidation points could in principle be the main driver of the finding that the average liquidation point for winners is




Figure 13: CDFs of liquidation points

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 liquidation 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 liquidation points from rounds in which, at some time before expiration, the asset price reaches the optimal lower threshold. In these restricted 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 13 plots the empirical distributions for the original sample (labeled “All”) and the restricted samples. The original and restricted samples of losers liquidations have very similar distributions. There is some evidence of censoring bias in the distribution of sales of winners, as the median liquidation point in the restricted sample is larger than in the original sample. However, even after correcting for censoring the median liquidation point in the restricted sample is lower than the optimal threshold. It is still possible to reject the hypothesis that sales of winners are clustered around the optimal level, with a Wilcoxon signed-rank p-value of nearly zero.

I have shown that the median aggregate liquidation points are consistent with the disposition effect condition (11), 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 liquidation 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 sell at that point. This is a very different behavior from selling 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 therefore seem to suggest early liquidations may in fact represent liquidation 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.¹⁵ For each round in which a liquidation 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 liquidation 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 liquidation decisions that seemed to occur too late based on liquidation points, i.e. $s_t > B_c^*$ for winners and $s_t < b^*$ for losers. Figure 14 illustrate the results (for each fraction I also show a two-standard error bar). For both criteria I show the fraction of late liquidations 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 liquidation decisions for winners occur too early, while the majority of the liquidation decisions for losers occur too late.

In the aggregate 80% of sales of losers occur later than optimal and similarly the majority of sales of 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. Again I consider both a liquidation-point criterion and a stopping-time criterion. I thus define b as the median from the sample of a subject’s lower liquidation points (and similarly for B). I classify subjects according to the following taxonomy. A subject’s behavior is consistent with the disposition effect (“DE”) if: $b < b^* \wedge B < B^*$. The label “Anti-DE” applies to subjects whose median liquidation points satisfy the inequalities: $b > b^* \wedge B > B^*$. The other two cases are a wider band, labelled “Risk-seeking”, and a narrower band relative to the optimal, labelled “Risk-averse”. I adapt this taxonomy to the stopping-time criterion. For each subject I compute the fraction of liquidations decisions that occurred too late relative to the optimal stopping times for both winners and losers, F_w and F_l respectively. I then classify the subject as “DE” if $F_w <$

¹⁵In terms of design, subjects’ choices could be restricted to threshold strategies as in Oprea (2012).

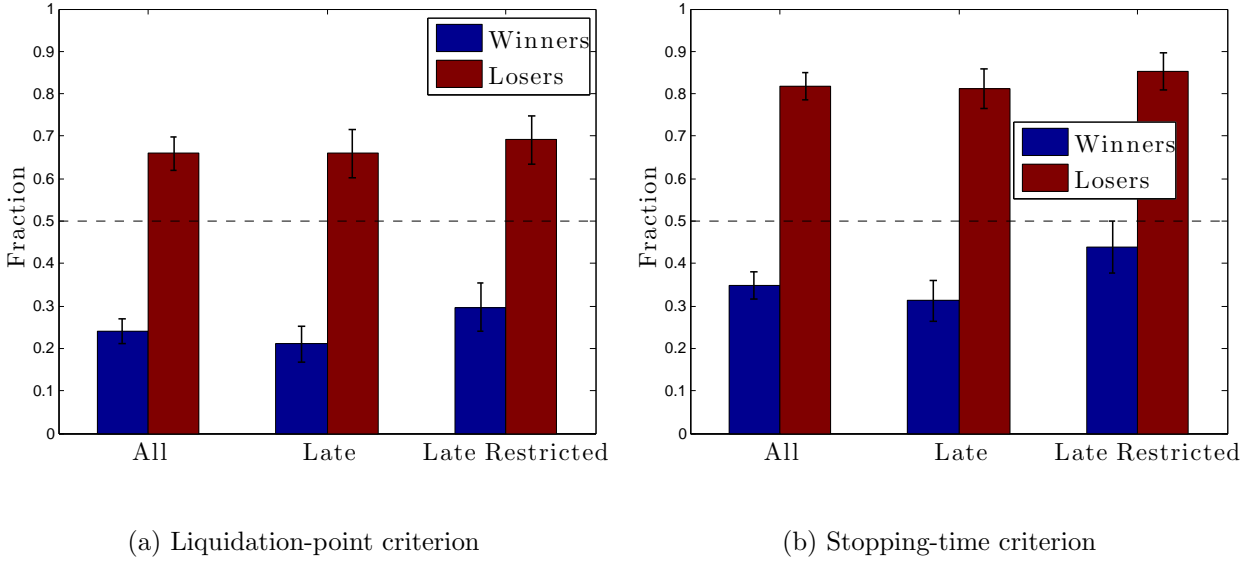


Figure 14: Fraction of liquidations decisions that occur too late

	DE	Anti-DE	Risk-seeking	Risk-averse
Liquidation-point	66%	0	13%	21%
Stopping-time	68%	0	27%	5%

Table 2: Subject classification: percentage of subjects in each category

0.5 and $F_l > 0.5$, “Anti-DE” if $F_w > 0.5$ and $F_l < 0.5$, “Risk-seeking” if $F_w > 0.5$ and $F_l > 0.5$, “Risk-averse” if $F_w < 0.5$ and $F_l < 0.5$. The resulting classification of subjects is summarized in table 2. For each category, in the first row I show the fraction of subjects according to the liquidation-point criterion and in the second line using the stopping-time criterion. Interestingly, according to both criteria around two thirds of the subject pool are classified as disposition-effect investors. The liquidation-point criterion underestimates the number of subjects that tend to liquidate both winners and losers too late (risk-seekers) and overestimates the number of subjects that tend to liquidate both winners and losers too early (risk-averse). The subject-level analysis summarized in table 2 leads to the following:

Result 2. *The behavior of more than half of the subjects is consistent with the disposition effect: they tend to sell winners too early and ride losers too long.*

Clearly the analysis of liquidation points ignores all the information on subjects’ inaction spells, i.e. when a subject chooses to wait. Instead of trying to incorporate this information in the otherwise straightforward approach of this subsection, I turn to structural estimation of a dynamic model that uses tick-level data on both liquidations and inaction.

4.2 Structural Analysis

In light of the discussion of the disposition effect presented in Section 2.2, the evidence on liquidation points suggests that the behavior of most subjects can be described by a model of realization preferences. Here I provide structural estimates of the realization utility model based on tick-level data. In the first place, structural estimation is an alternative way of testing for the disposition effect. More importantly, parameter estimates of realization utility models are an important input for ongoing theoretical work in behavioral finance.

In order to fit tick-level data I adopt a discretized version of the realization utility model summarized in equation (13): the environment is that described by the diagram in Figure 8, where the driving exogenous stochastic process is given in expression 14. The major obstacle in taking the theoretical framework to the data is that optimal behavior in the model takes the form of a (constant) threshold rule, while actual behavior is more noisy. In order to account for this, I assume that an agent liquidates in each tick with some positive probability and thus define $\eta(s)$ as the hazard of liquidating at a price s . The main additional assumption I make is that the log-odds of liquidating is an affine function of the value gain from liquidating ($V(s)$), with intercept a_0 and slope a_1 :

$$\ln \frac{\eta(s)}{1 - \eta(s)} = a_0 + a_1 V(s) \quad (15)$$

with $a_1 > 0$. This is a simple formulation that allows noise in actual behavior, while requiring the agent to be more (less) likely to liquidate when it is more (less) profitable.

The value gain from liquidating is given by the difference between the (material and psychological) payoff from liquidating and the option value of holding on to the asset ($W(s)$):

$$V(s) = \max(s, x) + ru(s) - W(s) \quad (16)$$

where the realization utility function $ru(\cdot)$ has the form specified in equation 12. Finally the value of holding on to the asset is given by the following Bellman equation:

$$W(s) = E_{s'} \{ \eta(s') [\max(s', x) + ru(s')] + [1 - \eta(s')] [l\delta s' + (1 - l - r)W(s')] | s \} \quad (17)$$

In words, with probability $\eta(s')$ in the next tick there is a liquidation and the agent obtains the (material and psychological) payoff: $\max(s', x) + ru(s')$. With probability $1 - \eta(s')$ the agent waits. Then she obtains the dividend with probability l , the stock expires with probability r and with probability $1 - l - r$ the agent receives the continuation value. The expectation $E_{s'}$ is computed using the process defined in (14). Equation (17) is solved numerically by spline polynomial approximations.

	β	α	a_0	a_1
Median	2.6	0.4	-6.2	1.8
Min	0.3	0.05	-10.0	0.1
Max	4.5	0.98	-3.4	2.6

Table 3: Parameter estimates.

I estimate the model given by equation (17) at the individual level by maximum likelihood. For each subject I have N tick-level observations: a set of dummies $\{d_i\}$, with $d_i = 1$ if there is a liquidation in tick i and 0 otherwise, and a sample path of the price process: $\{s_i\}$. The contribution to the likelihood of a liquidation point s is the liquidation probability $\eta(s)$, while a tick from an inaction spell contributes the complement $1 - \eta(s)$. Thus I form the following likelihood function:

$$\ell = \prod_{i=1}^N \eta(s_i)^{d_i} [1 - \eta(s_i)]^{1-d_i} \quad (18)$$

Estimates of the parameters $\{\alpha, \beta, a_0, a_1\}$ are obtained by maximizing (18) numerically. I use the simulated annealing method, a global optimization algorithm, with frequent reannealing and an agnostic starting value. Table 3 summarizes the estimation results.

4.2.1 Preferences for Realizing Winners

There are a number of conclusions that can be drawn from the results. First, the parameter estimates support the hypothesis that subjects have an irrational disposition to sell winners and hold on to losers. The estimates of the realization utility weight, β , are large, with a median at 2.6.¹⁶ Thus subjects behave as if they experience significant realization utility from liquidation. That these estimates imply a considerable disposition effect can be gauged also by inspecting the shape of the liquidation hazard functions $\eta(s)$. Figure 15 plots the hazard for each subject and the unweighted average hazard (the thick dashed line).

For all subjects the liquidation hazard is near zero when the current stock price is close to the purchase price (10) and it increases as the current price deviates from the purchase price. This feature is consistent with some recent evidence from field data: for example, using data on stock trading by retail investors, Ben-David and Hirshleifer (2012) find that the graph of the probability of selling has a V-shape centered around zero profits. The other, more important feature of the hazard functions I estimate from experimental data is that they are

¹⁶A median estimated β of 2.6 is large in the sense that the threshold model delivers an intensity of the disposition effect at least as high as $\Psi = 9\%$. This minimum is reached when $\alpha = 1$ and the effect is larger for lower values of α . For $\alpha = 0.5$, $\beta = 2.6$ implies an increase in the expected fraction of winners in total sales relative to optimal of $\Psi = 64\%$.

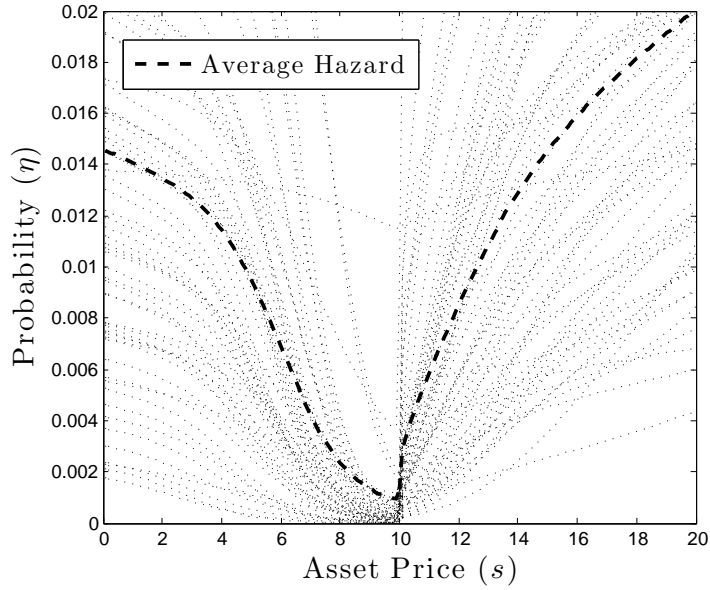


Figure 15: Estimated Hazard Functions $\eta(s)$

asymmetric around zero and tend to be higher for gains than for losses. For example, the average liquidation hazard at +2 points from the purchase price is 0.9%, while it is 0.2% at -2 points. At +4 points the average hazard is 1.29% while at -4 points it is 0.69%. This implies that subjects tend to sell stocks that have experienced gains in their price sooner relatively to stocks that have fallen in price since purchase. This feature of the estimated hazard functions is remarkably similar to what recent studies have found. For example, using data from both a US large discount broker and a Finnish dataset, 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.

Result 3. *The structural estimation results support the hypothesis that subjects have an irrational disposition to sell winners and hold on to losers: the estimates of the realization utility weight (β) are large and estimated liquidation hazard functions are higher for winners than for losers.*

4.2.2 Diminishing Sensitivity

The second set of conclusions that emerge from the structural estimation concerns the shape of the realization utility function, as parameterized by α . The median estimates of α is 0.4. It is interesting to compare this result to existing estimates of the sensitivity parameter for S-shaped utility functions. Existing estimates are derived from static lottery-choice experiments in the context of estimation of prospect-theoretic decision models. 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, but most recent studies have found larger values. Abdellaoui (2000) obtains an average α around 0.9. Bruhin, Fehr-Duda and Epper (2010) find that most subject-level estimates of α are around 1. The present estimate of α implies that the sensitivity of realization utility decreases much faster in the size of realized gains and losses than what is implied by canonical estimates of prospect-theoretic models.

As discussed above, the realization utility model generates a stronger disposition effect jointly with voluntary realization of losses when the sensitivity to gains and losses diminishes faster (i.e. α is lower). Therefore a small estimated α (together with a significant β) is consistent with the evidence of a strong disposition effect in terms of liquidation points. The former result is not a simple relabeling of the latter, however, as the analysis of liquidation points ignored all the information on inaction spells; rather this provides some evidence that the identification strategy has internal validity. I conclude by summarizing this discussion in the following:

Result 4. *The median estimate of the sensitivity parameter of the realization utility function (α) is 0.4. The estimates are consistently lower than those obtained from static lottery-choice tasks (such as Tversky and Kahneman (1992); Wu and Gonzalez (1996); Abdellaoui (2000); Bruhin et al. (2010)).*

5 Discussion

In this paper I build a model of liquidation decisions specifically tailored for detecting and structurally estimating preferences for realizing winners vs. losers. Many empirical studies have found evidence suggesting the disposition effect in financial markets - the tendency of individual investors to sell assets whose price has increased and keep assets that have dropped in value. However, most of the evidence is indirect and whether the effect actually reflects an irrational disposition of unsophisticated investors is still debated. On the theoretical front, the search for a solid microfoundation of the disposition effect has produced a new model of investor behavior, realization utility, that needs to be empirically validated. Providing precise evidence on the shape of the realization utility function is particularly important as even qualitative results, such as whether investors voluntarily realize losses or not, depend on it.

To address these issues, I develop and conduct a laboratory test based on a dynamic stochastic model of liquidation decisions. I design a particular security that bundles a risky asset, whose price follows a stochastic process in continuous time, with a perpetual put option. I analyze an impatient investor who makes a decision about when, if ever, to liquidate the investment. Optimal behavior entails maintaining the current position in the security until the asset price reaches either an upper threshold or a lower threshold. Deviations from

this benchmark matching the disposition effect have a very distinctive character and cannot be confounded by standard preferences, such as risk aversion.

Results from my experiment strongly support the hypothesis that individual investors have an irrational preference for selling winning stocks vis-à-vis selling losing stocks. The majority of subjects tend to delay selling losers beyond the optimal point and to sell winners before reaching the optimal liquidation point. This departure from rational behavior is economically significant as it implies a probability of realizing a gain conditional on a sale 56% larger than optimal. The disposition effect hypothesis is also supported by structural analysis that uses both data on inaction and liquidation decisions. I find that for most subjects the estimated utility weight on gains and losses is large and significant. Estimated liquidation hazards are higher for gains than for losses, illustrating another dimension of the disposition effect.

Finally, I find that the estimated realization utility functions satisfy the property of diminishing sensitivity. Moreover, the sensitivity of realization utility to gains and losses decreases faster than what is implied by canonical estimates of prospect-theoretic value functions. This finding has important theoretical implications. In order to obtain quantitative predictions about trading and other market statistics, researchers in theoretical behavioral finance usually calibrate their models to parameters estimated from experiments on static risky-choice tasks. For example, Ingersoll and Jin (2013) calibrate the sensitivity parameter of their realization utility model to $\alpha = 0.88$ as estimated in the classical study of Tversky and Kahneman (1992). However, Ingersoll and Jin (2013) find that for $\alpha = 0.88$ investors never voluntarily realize losses in their model, and that a good fit to market statistics can be achieved only using values of α that are extreme when interpreted as measures of risk attitudes. Here I provide direct evidence that the elasticity of realization utility to gains and losses is far from canonical estimates obtained from static risky-choice tasks. A broader conclusion is that important properties of reference-dependent preferences may vary significantly with the context.

While this paper shows that unsophisticated investors prefer to realize gains over losses, this experiment was not designed to test whether such realization preferences matter for markets.¹⁷ Two issues seem particularly worth exploring. First, I have shown through descriptive statistics and structural results that the subjects of my experiments are heterogeneous, although most of them behave consistently with the disposition effect. Future work could investigate whether market interaction between investors with realization preferences and

¹⁷There is still very little work on the market implications of realization preferences. Ingersoll and Jin (2013) argue that realization utility can explain a flatter security market line and the negative pricing of idiosyncratic risk. Hartzmark and Solomon (2012) find mispricing in a sport prediction market that is consistent with the disposition effect.

rational investors dampens or magnifies the disposition effect. Second, in this experiment some subjects are willing to deviate from the strategy that maximizes expected consumption utility. A related question is whether realization preferences dominate market incentives to trade optimally, such as the returns to using private information in an asset market.

References

- Abdellaoui, M. (2000). Parameter-free elicitation of utility and probability weighting functions, *Management Science* **46**(11): 1497–1512.
- Barber, B. M. and Odean, T. (2011). The behavior of individual investors, in G. M. Constantinides, M. Harris and R. M. Stulz (eds), *Handbook of the economics of finance: Financial Market and Asset Pricing*, Vol. 2B, Elsevier, Amsterdam.
- Barberis, N. and Xiong, W. (2009). What drives the disposition effect? an analysis of a long-standing preference-based explanation, *the Journal of Finance* **64**(2): 751–784.
- Barberis, N. and Xiong, W. (2012). Realization utility, *Journal of Financial Economics* **104**(2): 251–271.
- Ben-David, I. and Hirshleifer, D. (2012). 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.
- Bruhin, A., Fehr-Duda, H. and Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion, *Econometrica* **78**(4): 1375–1412.
- Dixit, A. (1993). *The Art of Smooth Pasting*, Harwood Academic Publishers.
- Frydman, C., Barberis, N., Camerer, C., Bossaerts, P. and Rangel, A. (2012). Using neural data to test a theory of investor behavior: An application to realization utility, *Working paper*, National Bureau of Economic Research.
- Genesove, D. and Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market, *The Quarterly Journal of Economics* **116**(4): 1233–1260.
- Greiner, B. (2004). The online recruitment system orsee 2.0 - a guide for the organization of experiments in economics, *Working Paper Series in Economics 10*, University of Cologne, Department of Economics.
- Hartzmark, S. M. and Solomon, D. H. (2012). Efficiency and the disposition effect in nfl prediction markets, *The Quarterly Journal of Finance* **2**(03).
- Ingersoll, J. E. and Jin, L. J. (2013). Realization utility with reference-dependent preferences, *Review of Financial Studies* **26**(3): 723–767.

- Odean, T. (1998). Are investors reluctant to realize their losses?, *The Journal of finance* **53**(5): 1775–1798.
- Oprea, R. (2012). Survival versus profit maximization in a dynamic stochastic experiment, *Technical report*, Working Paper, University of British Columbia 2012.
- Oprea, R., Friedman, D. and Anderson, S. T. (2009). Learning to wait: A laboratory investigation, *The Review of Economic Studies* **76**(3): 1103–1124.
- Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence, *The Journal of finance* **40**(3): 777–790.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and uncertainty* **5**(4): 297–323.
- Weber, M. and Camerer, C. F. (1998). The disposition effect in securities trading: An experimental analysis, *Journal of Economic Behavior & Organization* **33**(2): 167–184.
- Wu, G. and Gonzalez, R. (1996). Curvature of the probability weighting function, *Management science* **42**(12): 1676–1690.

Online Appendices

A Model Details

A.1 The Value Function Inside the Inaction Region

Inside the inaction region the value function can be characterized by the following equation:

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

Taking $\Delta t \rightarrow 0$:

$$\rho v(s) = \lambda[y - v(s)] + \frac{1}{dt}E_t dv(z(t)), s \in (\underline{s}, \bar{s}) \quad (19)$$

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 investor 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:

$$(\rho + \lambda)v(s) = \lambda y + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (\underline{s}, \bar{s}) \quad (20)$$

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

$$(\rho + \lambda)v(s) = \lambda \delta s + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (\underline{s}, \bar{s}) \quad (21)$$

The general solution of this second order differential equation is:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}s + C_1 s^{R_1} + C_2 s^{R_2} \quad (22)$$

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 are given by:

$$\lim_{s \searrow \theta^*} v(s) = x$$

$$\lim_{s \searrow \theta^*} v'(s) = 0$$

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

$$\lim_{s \nearrow \Theta^*} v'(s) = 1$$

Substituting the expression for the value function gives:

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}b^* + C_1b^{*R_1} + C_2b^{*R_2} - x = 0 \quad (23)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1R_1b^{*R_1-1} + C_2R_2b^{*R_2-1} = 0 \quad (24)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}B^* + C_1B^{*R_1} + C_2B^{*R_2} - B^* = 0 \quad (25)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1R_1B^{*R_1-1} + C_2R_2B^{*R_2-1} - 1 = 0 \quad (26)$$

A.2 Comparative Statics

The panels in Figure 16 illustrate how the key model parameters affect the upper and lower liquidation thresholds. As in standard stopping time models, a higher volatility widens the inaction region around x . Both the upper and lower thresholds are monotonic in the strike price x . Finally higher impatience, induced either by a larger discount factor or a larger intensity of dividend arrival, shrinks the inaction region around x .

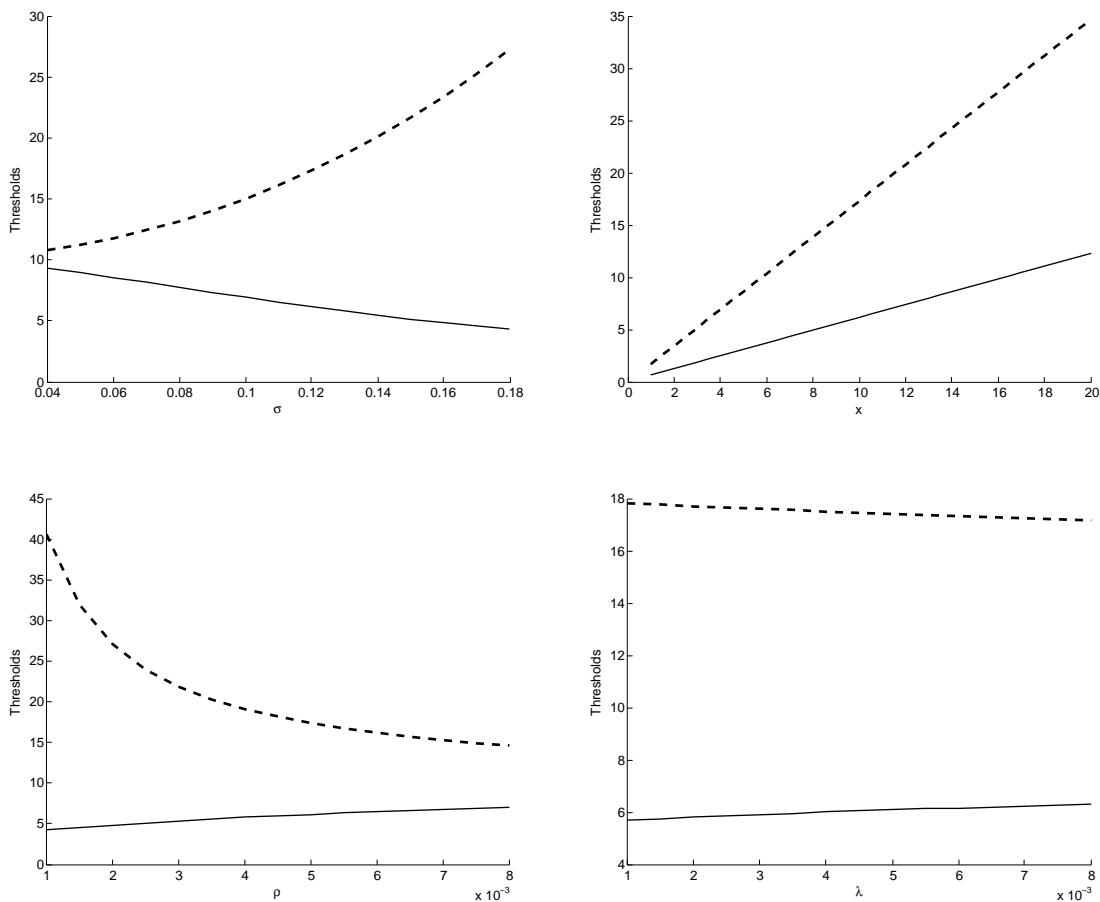


Figure 16: Comparative statics.

Note: the baseline parameter values are those summarized in Table 1.

A.3 Discrete Approximation

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

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

B Details of the Empirical Analysis

B.1 Purchasing Decisions

Figure 17 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 17: The distribution of the number of shares bought