

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

Boundedly rational behavior in dynamic stochastic settings: three essays

Permalink

<https://escholarship.org/uc/item/75k9f9sc>

Author

Magnani, Jacopo

Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
UNIVERSITY OF CALIFORNIA SANTA CRUZ

**BOUNDEDLY RATIONAL BEHAVIOR IN DYNAMIC STOCHASTIC
SETTINGS:
THREE ESSAYS**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

Jacopo Magnani

June 2014

The Dissertation of Jacopo Magnani
is approved:

Professor Daniel Friedman, Chair

Professor Carl Walsh

Professor Ryan Oprea

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Copyright © by

Jacopo Magnani

2014

Table of Contents

List of Figures	v
Abstract	vi
Dedication	vii
1 The Disposition Effect and Realization Preferences: a Direct Test	1
1.1 Introduction	1
1.2 Theoretical Framework	7
1.2.1 Optimal Asset Liquidation with a Constant Salvage Value	7
1.2.2 The Disposition Effect: Definition and Measurement	11
1.2.3 Risk Preferences	13
1.2.4 Realization Preferences	14
1.3 The Experiment	21
1.3.1 Implementing the Model in the Lab	21
1.3.2 Experiment Details	23
1.4 Results	26
1.4.1 Tests of the Disposition Effect using Liquidation Data	26
1.4.2 Structural Analysis	32
1.5 Discussion	38
1.6 Appendix: Model Details	41
1.6.1 The Value Function Inside the Inaction Region	41
1.6.2 Comparative Statics	44
1.6.3 Discrete Approximation	45
1.7 Details of the Empirical Analysis	46
1.7.1 Purchasing Decisions	46
2 Time and State Dependence in an Ss Decision Experiment	47
2.1 Introduction	47
2.2 Theoretical Model, Ss Predictions and Bounded Rationality	55
2.2.1 Model and State Dependent Optima	56
2.2.2 State Dependence, Time Dependence and Bounded Rationality	58

2.3	Experimental Implementation and Treatment Design	63
2.3.1	Binomial Conversion	63
2.3.2	Implementation	66
2.3.3	Treatments and Motivation	68
2.4	Results	70
2.4.1	Reduced Form Results	71
2.4.2	Structural Results	75
2.4.3	Robustness: The CORR treatment	78
2.5	Discussion	80
2.6	Appendix: Theoretical Details	87
2.6.1	Full State Dependent Model and Solution Derivation	87
2.6.2	Payoff Hills	93
2.6.3	Solution to the discretized problem	93
2.6.4	Derivation of HJB equation (2.9)	95
2.6.5	Discrete Time Cognitive Costs Model	96
2.7	Appendix: Instructions to Subjects	98
3	Evolutionary Games, Climate and the Generation of Diversity	105
3.1	Introduction	105
3.2	Results	107
3.3	Discussion	112
3.4	Methods	115

List of Figures

1.1	The value function of the liquidation problem	10
1.2	Disposition effects	12
1.3	The effect of risk aversion	14
1.4	The shape of the realization utility function	16
1.5	Realization preferences and the lower threshold	18
1.6	Realization utility and the disposition effect	19
1.7	The intensity of the disposition effect	19
1.8	Timeline of a tick.	22
1.9	Buying stage.	24
1.10	Selling stage.	25
1.11	Liquidation points	27
1.12	Liquidation points in early and late rounds	28
1.13	Liquidation points in the restricted samples	29
1.14	Fraction of liquidations decisions that occur too late	31
1.15	Estimated hazard functions $\eta(s)$	36
1.16	Comparative statics.	44
1.17	The distribution of the number of shares bought	46
2.1	Histogram of AC Nielsen data	49
2.2	Histograms of simulated adjustments	58
2.3	Software screen shot	65
2.4	Histograms of adjustments	71
2.5	Estimates	74
2.6	The CORR treatment	78
2.7	Payoff hills	94
2.8	Value functions	96
2.9	Screen example	100
2.10	Earnings in a tick	103
3.1	Data series	107
3.2	Data and Nash equilibria	110

Abstract

Boundedly rational behavior in dynamic stochastic settings:

three essays

by

Jacopo Magnani

In this work, I study the behavior of boundedly rational agents in dynamic stochastic settings. The first chapter develops a new laboratory test of the hypothesis that individual investors have an irrational preference for selling winning stocks vs. selling losing stocks. Subjects tend to delay selling losers beyond the optimal point and to sell winners before reaching the optimal liquidation point. Such behavior is shown to be consistent with a model of reference-dependent preferences. The second chapter presents the results of a lab experiment on a menu cost model of price adjustments. Subjects' average adjustment points show significant time dependent components in adjustment rules. The data suggest subjects experience substantial cognitive costs from responding to the state, accounting for these patterns. The third chapter presents techniques for studying the dynamics of evolutionary games that are subject to stochastic disturbances and applies them to morph frequency data from a population of side-blotched lizards. The estimated payoff matrix implies a powerful interior attractor and recover the observed oscillations.

I would like to thank Dan Friedman, Ryan Oprea, Carl Walsh, Aspen Gorry and
Barry Sinervo.

Chapter 1

The Disposition Effect and Realization Preferences: a Direct Test

1.1 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

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.

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 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

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

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 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

jump in selling at zero profits.

⁴The laboratory environment that is closest to my experiment is that of Oprea et al. [2009b]. While Oprea et al. [2009b] 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.

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

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

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 [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.

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 et al. [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

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.

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 1.3 describes the experimental design. Section 1.4 presents the results and section 1.5 concludes.

1.2 Theoretical Framework

1.2.1 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:

$$ds_t = \mu s_t dt + \sigma s_t dw_t \quad (1.1)$$

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:

$$y_t = \delta s_t \quad (1.2)$$

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

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.

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 (1.2). Let

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

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

$$v(s_t) = \max_{\tau_o, \tau_c} E_t \left\{ e^{-\rho(\tau-t)} [\mathbb{1}_{\{\tau=\tau_o\}}x + \mathbb{1}_{\{\tau=\tau_c\}}s_\tau + \mathbb{1}_{\{\tau=\tau_\lambda\}}y_\tau] \right\} \quad (1.3)$$

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:

$$v(s) = s, \forall s \geq B \quad (1.4)$$

$$v(s) = x, \forall s \leq b \quad (1.5)$$

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:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}s + C_1s^{R_1} + C_2s^{R_2}, s \in (b, B) \quad (1.6)$$

where R_1 and R_2 are algebraic functions of the parameters (given in Appendix 1.6.1), 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 (1.3) 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:

$$\lim_{s \searrow b^*} v(s) = x \quad (1.7)$$

$$\lim_{s \searrow b^*} v'(s) = 0 \quad (1.8)$$

$$\lim_{s \nearrow B^*} v(s) = B^* \quad (1.9)$$

$$\lim_{s \nearrow B^*} v'(s) = 1 \quad (1.10)$$

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:

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

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




Figure 1.1: Value function and inaction region of the problem

The two optimal thresholds can then be found by standard numerical methods. The value function of the investor's problem is illustrated in Figure 1.1.

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 1.6.2.

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

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 (1.7),(1.8),(1.9),(1.10).

1.2.2 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 1.2, where I denote by B (b) the actual upper (lower) liquidation threshold used by the investor. Panels 1.2a and 1.2b 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 1.2c illustrates the notion that the disposition effect leads to realizing gains more frequently and losses more rarely than optimal. The situation illustrated in Figure 1.2 is summarized in the following:

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

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.





Figure 1.2: Disposition effects

counterpart). $\{b, B\}$ satisfy the disposition effect condition if:

$$b < b^* \wedge B < B^* \quad (1.11)$$

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 (1.11).

1.2.3 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:

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




Figure 1.3: The effect of risk aversion

where $u(m) = \frac{m^{1-\gamma}}{1-\gamma}$. In figure 1.3 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 (1.11) 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.

1.2.4 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 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




Figure 1.4: The shape of the realization utility function

realization utility function, with diminishing sensitivity:

$$ru(s_t) = \begin{cases} \beta \left(\frac{s_t - s_0}{s_0} \right)^\alpha & \text{if } \frac{s_t - s_0}{s_0} \geq 0 \\ -\beta \left(-\frac{s_t - s_0}{s_0} \right)^\alpha & \text{if } \frac{s_t - s_0}{s_0} < 0 \end{cases} \quad (1.12)$$

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 1.4 plots the shape of the realization 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:

$$v(s_t) = \max_{\tau_o, \tau_c} E_t \left\{ e^{-\rho(\tau-t)} \left[\mathbb{1}_{\{\tau=\tau_o\}}(x + ru(s_\tau)) + \mathbb{1}_{\{\tau=\tau_c\}}(s_\tau + ru(s_\tau)) + \mathbb{1}_{\{\tau=\tau_\lambda\}}y_\tau \right] \right\} \quad (1.13)$$

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 (1.13) does not admit a lower threshold, that is the agent chooses to never realize losses. This is illustrated in Figure 1.5, 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 threshold is set to zero. When the lower threshold is zero, no losers are ever realized.

In Figure 1.6 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

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 (1.6) and the only changes are in the boundary conditions.




Figure 1.5: Realization preferences and the lower threshold

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 1.7 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:

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




Figure 1.6: Realization utility and the disposition effect




Figure 1.7: The intensity of the disposition effect

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 voluntary 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 1.6a 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 1.6 and 1.7: 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.

1.3 The Experiment

1.3.1 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:

$$s_{t+\Delta t} = \begin{cases} s_t(1+h) & \text{with probability } p \\ s_t(1-h) & \text{with probability } 1-p \end{cases} \quad (1.14)$$

As in other dynamic experiments (such as Oprea et al. [2009b]), 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




Figure 1.8: Timeline of a tick.

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 1.6.3. Figure 1.8 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 2.1 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 $\Pi^* = 34\%$.

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 1.1: Parameter values and theoretical predictions.

1.3.2 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. $v(s_0) > s_0$).




Figure 1.9: Buying stage.

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 1.9 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 1.10 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 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



Figure 1.10: Selling stage.

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.

1.4 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 1.7.1, 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 1.17).

1.4.1 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 1.11 illustrates the distribution of liquidation points pooling all subjects in the study. Letting S represent the random variable that generates liquidation points, Panel 1.13b plots $Prob(S \leq s | S \geq s_0)$, while panel 1.13a 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:

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




Figure 1.11: CDFs of liquidation points

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 1.12 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




Figure 1.12: CDFs of liquidation points

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 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




Figure 1.13: CDFs of liquidation points

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 1.13 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 (1.11), 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 1.14 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.

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




Figure 1.14: Fraction of liquidations decisions that 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 $F_w < 0.5$ 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

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

Table 1.2: Subject classification: percentage of subjects in each category

$F_l > 0.5$, “Risk-averse” if $F_w < 0.5$ and $F_l < 0.5$. The resulting classification of subjects is summarized in table 1.2. 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 1.2 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 informaton 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.

1.4.2 Structural Analysis

In light of the discussion of the disposition effect presented in Section 1.2.2, 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 (1.13): the environment is that described by the diagram in Figure 1.8, where the driving exogenous stochastic process is given in expression 1.14. 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 :

$$\ln \frac{\eta(s)}{1 - \eta(s)} = a_0 + a_1 V(s) \quad (1.15)$$

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 ($W(s)$):

$$V(s) = \max(s, x) + ru(s) - W(s) \quad (1.16)$$

where the realization utility function $ru(\cdot)$ has the form specified in equation 1.12. Finally the value of holding on to the asset is given by the following Bellman equation:

$$W(s) = E_{s'} \{ \eta(s') [\max(s', x) + ru(s')] + [1 - \eta(s')] [l\delta s' + (1 - l - r)W(s')] | s \} \quad (1.17)$$

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 (1.14). Equation (2.30) is solved numerically by spline polynomial approximations.

I estimate the model given by equation (2.30) 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:

$$\ell = \prod_{i=1}^N \eta(s_i)^{d_i} [1 - \eta(s_i)]^{1-d_i} \quad (1.18)$$

Estimates of the parameters $\{\alpha, \beta, a_0, a_1\}$ are obtained by maximizing (1.18) numerically. I use the simulated annealing method, a global optimization algorithm, with frequent re-annealing and an agnostic starting value. Table 1.3 summarizes the estimation results.

	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 1.3: Parameter estimates.

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 2.32 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

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 1.15: Estimated hazard functions $\eta(s)$

data is that they are 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.

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 et al. [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]).

1.5 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 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.

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.

1.6 Appendix: Model Details

1.6.1 The Value Function Inside the Inaction Region

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

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

Taking $\Delta t \rightarrow 0$:

$$\rho v(s) = \lambda [y - v(s)] + \frac{1}{dt} E_t dv(z(t)), s \in (\underline{s}, \bar{s}) \quad (1.19)$$

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:

$$(\rho + \lambda)v(s) = \lambda y + \mu s v'(s) + \frac{1}{2} \sigma^2 s^2 v''(s), s \in (\underline{s}, \bar{s}) \quad (1.20)$$

I use (1.2) to express the dividend in terms of the current price and obtain:

$$(\rho + \lambda)v(s) = \lambda\delta s + \mu s v'(s) + \frac{1}{2}\sigma^2 s^2 v''(s), s \in (\underline{s}, \bar{s}) \quad (1.21)$$

The general solution of this second order differential equation is:

$$v(s) = \frac{-2\lambda\delta}{\sigma^2(1 - R_1)(1 - R_2)}s + C_1 s^{R_1} + C_2 s^{R_2} \quad (1.22)$$

where:

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

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

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:

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}b^* + C_1b^{*R_1} + C_2b^{*R_2} - x = 0 \quad (1.23)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1R_1b^{*R_1-1} + C_2R_2b^{*R_2-1} = 0 \quad (1.24)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)}B^* + C_1B^{*R_1} + C_2B^{*R_2} - B^* = 0 \quad (1.25)$$

$$\frac{-2\lambda\delta}{\sigma^2(1-R_1)(1-R_2)} + C_1R_1B^{*R_1-1} + C_2R_2B^{*R_2-1} - 1 = 0 \quad (1.26)$$

1.6.2 Comparative Statics

The panels in Figure 1.16 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 1.16: Comparative statics.

1.6.3 Discrete Approximation

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

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

1.7 Details of the Empirical Analysis

1.7.1 Purchasing Decisions

Figure 1.17 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 1.17: The distribution of the number of shares bought

Chapter 2

Time and State Dependence in an Ss Decision Experiment¹

2.1 Introduction

When the world changes, firms must change with it to continue optimizing. But when it is costly to make changes, it pays to be judicious about timing. Indeed, Ss models show that when the future is uncertain and adjustment is costly, it is best to ignore minor deviations from optimality and delay re-optimizing until failing to do is sufficiently costly. Such decision problems are canonical, arising throughout economics. Indeed, Ss models have become a workhorse in the field in the last three decades, helping economists understand sticky prices, inventory dynamics, cash flows, hiring and firing decisions, and investments.²

¹This is a joint work with Ryan Oprea and Aspen Gorry.

²As summarized in Caplin and Leahy [2010], the development of Ss models grew out of an interest in understanding how to manage inventories. Arrow et al. [1951] showed that Ss policies are cost minimizing in infinite horizon inventory problems while Scarf [1959] demonstrated that optimal decisions in a large class of such inventory problems take an Ss form. Given the pervasiveness of both uncertainty and fixed costs of adjustment in economics, Ss models have since been deployed in a wide range of applications in economics

Ss models point out that agents should choose when to re-optimize solely based on the state of the world (“state dependent” adjustment) and not on any direct function of time itself (“time dependent” adjustment).

Do people use optimal state dependent rules when making adjustments? Among the best studied Ss-type problems empirically is firms’ timing of price changes in the face of “menu costs.” A central ambition of the empirical literature on price adjustment is to discern whether firms tend to employ the sorts of state dependent price changes predicted in Ss models or whether they instead employ “time dependent” rules of the sort described by Taylor [1980] or Calvo [1983], a question of critical importance to monetary policy.³ Although state dependent models are empirically successful in several dimensions, much of the evidence so far suggests that they fail to describe the data in important ways: firms also show some evidence of time dependence in their actual decision making.⁴ To illustrate, Figure 2.1 plots log price adjustments using AC Nielsen data (the data were constructed in Midrigan [2008] and used as an empirical benchmark in Costain and Nakov [2011b,a]). As in state dependent models, adjustments are bimodal with modes above and below zero.

However, as in time dependent models the sizes of price adjustments are diffuse and not

including the interaction of portfolio choice and housing demand (Grossman and Laroque [1990]), the demand for durables (Bertola and Caballero [1990]), portfolio choice and stock ownership (Vissing-Jorgensen [2002]), investment (Pindyck [1988] and Abel and Eberly [1994]), the decision of countries to exit the Euro (Alvarez and Dixit [2013]) and price adjustment (Barro [1972] and Sheshinski and Weiss [1977]).

³As we discuss in more detail in section 5, while time dependent pricing is not a necessary condition for monetary policy to have significant real effects, different monetary policies can have different effects under time dependent behavior and under state dependent behavior.

⁴Indeed, the recent empirical literature has produced mixed support for both pure state and time dependent models. While Midrigan [2010] finds significant evidence of selection effects in firms’ price changing decisions, consistent with state dependence, Nakamura and Steinsson [2008] find evidence of seasonality in price changes and no evidence for the upward sloping hazard functions of price adjustment that are predicted by standard state dependent models. Klenow and Kryvtsov [2008] find that purely state dependent models are unable to generate as many small and large price changes as are observed in the data. On the other hand, they also find that the correlation between the time between adjustments and the size of adjustments is weaker than would be expected from a purely time dependent model.




Figure 2.1: Histogram of AC Nielsen data of log price adjustment data were constructed by Midrigan [2008] and used as an empirical benchmark in Costain and Nakov [2011b,a].

tightly clustered around modes; adjustment data provides little evidence that firms establish the inaction regions predicted by Ss models.⁵

There are two broad classes of explanations available for the predictive failures of Ss models. One class is institutional: the environments studied in the field may include institutional features that rationalize time dependence but are not included in the structure of simple Ss models.⁶ Another is cognitive: the models have the description of the decision

⁵Of course underlying parameters of the economy and institutional features of the decision task are unobservable in naturally occurring data such as this, making it difficult to make direct and easily interpretable comparisons to theoretical models. This observability problem is a central motive for conducting direct tests using controlled and sterile experiments as a complement to conventional empirics.

⁶In order to explain the empirical distribution of price changes, “second-generation” state dependent pricing models have been developed that rationalize the empirical evidence by including additional institutional details in the model. Some of these competing explanations are: variable menu costs as in Dotsey et al. [1999], infrequent shocks under small menu costs as in Gertler and Leahy [2008] and economies of scope for changing the prices of multiple products as in Midrigan [2011] and Alvarez and Lippi [2012]. For a survey of the literature see Klenow and Malin [2010].

problem right but limits in cognitive capacity or attention prevent firms from fully assessing or responding to the current state.⁷ Such boundedly rational agents' adjustment rules include properties of time dependent models that can generate patterns like those in Figure 2.1.

Because the underlying decision problem faced by firms is only imperfectly observed in naturally occurring data, it is difficult to directly judge between these two classes of explanations using conventional empirics. In this paper we use laboratory experiments to help shed some new light on this question. Experiments have, for our purposes, one important advantage over the field: they give the analyst perfect information regarding and perfect control over the exact decision problem faced by subjects. If we precisely implement the simple environment of a standard Ss pricing model in the controlled and sterile setting of the lab and observe deviations in behavior matching the distinctive ones observed in the field, the cause can only be bounded rationality; alternative explanations are unavailable in our laboratory environment.

In our experiment, subjects are shown a graphical visualization of their static price and are also shown the evolution of the optimal relative price which fluctuates according to a geometric Brownian motion. The further the optimal price strays from the current price (either above or below), the lower their flow earnings. A subject's only decision is

⁷Woodford [2008] builds a model based on the concept of rational inattention (costly information gathering) developed by Sims [1998, 2003, 2006] where the standard Ss model is generated as a limiting case of no information costs. More recently Costain and Nakov [2011b,a] develop a similar theory of smoothly state dependent pricing that assumes agents face cognitive costs in processing and using information to make optimal decisions. Finally, Alvarez et al. [2011] show that costly information processing (of a sort modeled in Caballero [1989], Reis [2006] and Alvarez et al. [2012b]) in menu cost Ss model can also generate an optimal blend of time and state dependent behavior. In this paper we focus on cognitive costs, rather than information frictions. However, plausible information frictions could be easily added to the experiment and we leave this interesting lab design for future investigations.

when to pay a fixed cost to adjust her price to the current optimum. Ss theory gives us clear predictions: subjects should set thresholds above and below the optimum, exercising their option to adjust only when the relative price passes one of these thresholds. Because we have access to the exact parameters of the game, we can precisely evaluate the theory's predictive power and compare it to time dependent alternatives – something that cannot be done using naturally occurring data.

Our first contribution is to show that subjects' adjustment patterns in the experiment match those observed in the field strikingly well. As in the field, adjustments in our data are bimodal, conforming with the two sided adjustment rules predicted by the theory. Moreover, because we have access to the induced parameters of the economy we are able to show, for the first time, that these adjustment points are, on average, nearly identical to those predicted in Ss models and obey Ss predicted comparative statics to a remarkable degree. However, as in the field, adjustments are extremely diffuse: subjects often make price adjustments far from optimal threshold points and do not establish true inaction regions. Subjects therefore show strong evidence of the same sort of time dependence measured in the field, though here only bounded rationality seems to be a plausible explanation.

Our second contribution is to use the unique precision of laboratory data to structurally estimate cognitive costs using a model of bounded rationality similar to ones recently proposed in the literature (Woodford [2008], Costain and Nakov [2011b,a]). Laboratory data gives us access to all of the relevant parameters of the game and therefore allows us to precisely evaluate the degree to which subjects deviate from optimal adjustment rules. Because of this, we can collect unusually credible structural estimates of the cognitive costs

of implementing state dependent decision rules for each subject in our sample.⁸ Maximum likelihood estimates suggest that sizable cognitive costs are widespread in our sample and lead to significant time dependence in the average subject’s adjustment decisions.

Our final contribution is to use experimental treatment variation to study whether these cognitive costs are exogenous as is often implicitly assumed in the literature or if they are, instead, influenced by other fundamentals of the economy. We focus attention on volatility, a variable of recent interest in the Ss empirical literature (e.g. Bloom et al. [2007], Vavra [2013]). Volatility has a distinctive predicted effect in an Ss setting: it should substantially increase the average size of price changes by pushing out thresholds for adjustment.⁹ Though recent work lends indirect support to such predictions, our access to the underlying parameters of the economic environment allow us to perform a precise and direct test with our data. We find that volatility does indeed push out average adjustment thresholds (almost to the exact degree predicted). However we also find that volatility has an important second effect not predicted by the orthodox model: it decreases the degree of state dependence in adjustment behavior, reducing individuals’ sensitivity to shocks. To confirm this finding, we run an additional diagnostic treatment in which the predicted effect of higher volatility is fully offset by a simultaneous reduction in the fixed cost of adjustment, keeping the predicted inaction region identical to the baseline treatment. Even after controlling for optimal adjustment rules in this way, subjects still show a lower responsiveness to shocks, confirming that factors not captured in orthodox models contribute to the effect

⁸Alvarez et al. [2012a] comes closest to our paper on this by directly estimating costs of observation of the state variable in an investment problem using a novel household survey dataset. While the mechanism they study is related to ours, the estimate is from a very different setting than firms’ pricing decisions.

⁹To be precise, volatility has two effects. First, and directly, for a given adjustment threshold, a rise in volatility triggers more frequent adjustment. Second, and in response, an optimizing agent should expand her inaction region, pushing these thresholds further away from the optimal state.

of volatility on adjustment behavior. Structural estimates reveal that cognitive costs of attending to the state are (i) much larger in high volatility than in low volatility treatments and (ii) are similar across our two high volatility treatments. Volatility therefore directly (and robustly) increases the cost of optimizing and decreases the degree of state dependence in our data.

This final result suggests that cognitive costs are powerfully shaped by the economic environment and cannot be treated as invariant to other parameters of the economy. We argue in our concluding discussion that this finding has both methodological and policy implications. Methodologically, it suggests that caution must be taken in conducting counterfactual exercises using cognitive cost estimates, especially without attending to their dependence on environmental variables like volatility. Policy-wise, it suggests a potentially severe behavioral cost to volatility: volatility influences agents' ability to optimize. This has potential implications for understanding the effects of monetary policy shocks on the economy and contributes to a growing literature on the effects of uncertainty on macroeconomic processes.

Although our experiment uses a price adjustment setting, our findings have a broader scope: costly adjustment against uncertain fundamentals arises throughout economics. As a result, our findings shed light on an important general component of economic decision making. Indeed, given the importance of Ss models in economics, it is surprising that (to the best of our knowledge) no previous experimental studies have examined behavior in these common settings. The closest study to ours is perhaps Oprea et al. [2009a] which studies how subjects learn to value real options in a related dynamic-stochastic environ-

ment, though the economic forces studied in that paper are quite different.¹⁰ More distantly related are several experiments that study pricing frictions in macroeconomic-inspired settings. Davis and Korenok [2011] study a multi player market in which agents are directly induced to use time dependent pricing rules and examines how these timing friction effect the speed of price convergence over a number of trading periods. Wilson [1998] studies the behavior of a monopolist facing nominal price shocks and shows that even when faced with menu costs, prices converge quickly to monopolistic levels in dynamic settings. Finally, Noussair et al. [2013] implement a multi-agent DSGE economy and show that menu cost treatments reduce the variability of inflation relative to other treatments. None of these papers, however, studies an Ss-type setting and none study the contrast between state and time dependence that animates our study. Duffy [2008] provides a comprehensive overview of the experimental macroeconomics literature.

Our paper also relates to a broader literature on bounded rationality and cognitive costs, usefully surveyed in Conlisk [1996], who note that “deliberation about an economic decision is a costly activity, and good economics requires that we entertain all costs.” A recent example by Caplin et al. [2011] develops and experimentally tests a model of satisficing where cognitive fixed costs shape sequential search over a choice set. They find that these costs increase with complexity, a result that seems related to our finding that volatility increases such costs in Ss decisions.

The rest of the paper is organized as follows. In section 2.2, we describe a simple

¹⁰In Oprea et al. [2009a], subjects choose when to make a one-time fixed investment in a project whose value evolves according to a geometric Brownian motion with positive drift. The reward function, action space and predictions are very different from those in the Ss models studied here. Moreover, while Oprea et al. [2009a] is focused on understanding learning rules driving aggregate long-run time series, ours is focused on understanding cross sectional bounded rationality and time dependent departures from threshold adjustment rules.

Ss model inspired by the menu cost literature and compare its predictions to time dependent models and hybrid bounded rationality models recently studied in the literature. In section 2.3, we discuss our strategy for implementing the model in the lab and motivate our experimental design. In section 2.4, we present our results. Finally, we conclude with a discussion in section 2.5. Theoretical details and instructions to subjects are collected in Online Appendices.

2.2 Theoretical Model, Ss Predictions and Bounded Rationality

In this section we discuss the standard Ss pricing model that we take to the lab, and compare its optimal state dependent predictions to time dependent and boundedly rational alternatives. Section 2.2.1 describes optimal state dependent predictions for a textbook version of a standard, continuous time Ss pricing model as described by Stokey [2008] (these environments were first studied by Barro [1972] and Sheshinski and Weiss [1977]). Full details on the motivating economic problem and technical details on the solution are provided in Online Appendix 2.6. Section 2.2.2 compares optimal state dependent adjustment to time dependent adjustment (specifically the Calvo [1983] adjustment rule). It then develops a model of bounded rationality, driven by costly cognition, that features purely state and time dependent adjustment behaviors as special cases.

2.2.1 Model and State Dependent Optima

Consider a firm that produces a single good at a unit cost c over an infinite time horizon with profits discounted at rate r . The firm's price relative to the market price index, $z(t) \equiv \frac{p(t)}{P(t)}$,¹¹ exogenously evolves according to the following geometric Brownian motion:¹²

$$dz(t) = (\sigma^2 - \mu)z(t)dt - \sigma z(t)dw(t) \quad (2.2)$$

generating flow profits equal to

$$g(z(t)) = (z(t))^{-\theta}(z(t) - c), \quad (2.3)$$

given a demand parameter θ .¹³

The firm's task, at each moment in time, is to decide whether to exert control over its relative price by discontinuously adjusting it to the optimal level z^* . The firm faces

¹¹ $p(t)$ is the firm's own price and $P(t)$ is the aggregate price index. The entire problem can be described entirely in terms of relative price, $z(t)$, because the profit function is homogeneous of degree one in both $p(t)$ and $P(t)$ and profits can be written entirely in terms of relative price. By reducing variables in this way, we will considerably simplify the problem for experimental subjects and focus the design on the central issue of state versus time dependence in adjustment.

¹²The exogenous motion of the firm's relative price is driven by underlying Brownian evolution of the aggregate price index which follows the process

$$dP(t) = \mu P(t)dt + \sigma P(t)dw(t) \quad (2.1)$$

with drift μ , volatility σ and standard Wiener process $w(t)$. Equation 2.2 converts the process to relative price space through a straightforward application of Ito's lemma.

¹³The firm is assumed to face a standard Spence-Dixit-Stiglitz demand function

$$q = \left[\frac{p}{P} \right]^{-\theta}. \quad (2.4)$$

where θ is the elasticity of substitution across differentiated product varieties.

Flow profits in the underlying model are

$$\pi(p(t), P(t)) = \left(\frac{p(t)}{P(t)} \right)^{-\theta} \left(\frac{p(t)}{P(t)} - c \right). \quad (2.5)$$

a fixed cost of m for each such adjustment and thus faces a sequence of optimal stopping problems: to choose the optimal times, $\{T_i\}$, to adjust relative prices to the optimal level, z^* , to maximize its value:

$$v(z_0) = \max_{\{z_i^*\}, \{T_i\}} E_0 \left\{ \int_0^{T_1} e^{-rt} g(z(t)) dt + \sum_i \left[-e^{-rt} m + \int_{T_i}^{T_{i+1}} e^{-rt} g(z(t)) dt \right] \right\}. \quad (2.6)$$

The optimal solution to this problem is to establish an “inactivity band” for $z(t)$, $(\underline{z}, \bar{z})$: whenever $z(t)$ lies within the inactivity band the firm makes no adjustment and whenever $z(t)$ reaches or crosses these bounds, the firm pays m to adjust $z(t)$ to z^* . Note that this optimal decision rule is unrelated to calendar time: switches are entirely dictated by the current state of the world, $z(t)$. In contrast to “time dependent” rules discussed in the literature (i.e. Calvo [1983] or Taylor [1980]), the optimal rule is purely “state dependent.”

Inside the inaction region the value function is characterized by the Hamilton-Jacobi-Bellman equation and outside the inaction region it is equal to the value at the optimal return point net of the menu cost. Thus the value of the firm over the entire state space is defined by:

$$rv(z) = g(z) + (\sigma^2 - \mu)zv'(z) + \frac{1}{2}\sigma^2 zv''(z), \quad z \in (\underline{z}, \bar{z}) \quad (2.7)$$

$$v(z) = v(z^*) - m, \quad z \notin (\underline{z}, \bar{z}) \quad (2.8)$$

Using these value functions, we can calculate the optimal bands (and the optimal adjustment point z^*) numerically using standard methods that exploit optimal return, value




Figure 2.2: Histogram of simulated adjustments for (a) state dependent agents, (b) time dependent agents (following Calvo rules) and (c) hybrid, boundedly rational agents.

matching, and smooth pasting conditions (full details are provided in Online Appendix 2.6.1).

Our experimental design, described below, will leverage two comparative static predictions made by the model. First, an increase in volatility (an increase in σ) widens the optimal inaction band $(\underline{z}, \bar{z})$. Second, a reduction in the cost of adjustment, m , narrows the optimal band. See Stokey [2008] for a fuller discussion of the comparative statics of the model.

2.2.2 State Dependence, Time Dependence and Bounded Rationality

What alternative decision rules might agents use in the setting of the model described in the previous section? Consider the following simple descriptive framework: agents set a probability λ of adjusting to z^* at each moment and can condition this λ on the current state z if they choose.

The optimal solution described in the previous section has a sharp prediction in

this framework: agents will set $\lambda = 0$ when z is in the inactivity band $(\underline{z}, \bar{z})$ and $\lambda = 1$ upon reaching thresholds $\underline{z}$ or $\bar{z}$. This is a strongly state dependent prediction as adjustment is entirely governed by (and is perfectly responsive to) the current value of $z(t)$. Panel (a) of Figure 2.2 plots a histogram of the distribution of switching points from a simulation populated by purely state dependent agents.¹⁴ A vertical dotted red line marks z^* and solid red lines mark the boundaries $\bar{z}$ and $\underline{z}$. The observable pattern generated by this decision rule is clear: we will observe adjustments near $\bar{z}$ and $\underline{z}$ but very little adjustment away from these boundaries.¹⁵

By contrast, classical time dependent models (ala Taylor [1980] and Calvo [1983]) predict that subjects will adjust based on calendar time rather than state. The Calvo model in particular predicts that agents will simply set $\lambda = \tilde{\lambda}$ (for some fixed $\tilde{\lambda}$) at each date, ignoring the state, z , altogether. Panel (b) of Figure 2.2 shows a histogram from a simulation using $\tilde{\lambda} = 0.02$. Unlike the state dependent behavior in panel (a), the distribution is strongly unimodal in state and quite diffuse, with adjustments occurring at many different relative prices. In the setting described in the previous section, such time dependent rules constitutes a substantial deviation from profit maximization.

What could cause agents to make time dependent instead of optimal state dependent decisions? One hypothesis recently advanced in the literature is that subjects may experience costs from implementing decision rules that require complex thought about and attention to state. While a time dependent rule requires virtually no thought to imple-

¹⁴The simulation (and others in this section) uses a discretization of the model parameterized by the vector $(\alpha, \sigma, r, m, \theta, c) = (0.5, 0.0268, 0.0012, 5, 5, 0.4)$ (the same parameters we use in our BASE treatment in the experiment). All simulations in this section use the discretized binomial approximation we use in the experiment, described in section 2.3.1 below.

¹⁵Because of the discretization, adjustments sometimes occur just beyond the predicted thresholds.

ment, conditioning adjustment on the state can require considerable attention and calculation. This is true in a specific and measurable sense: under the time dependent rule the (Bernoulli distributed) adjustment random variable and z have zero mutual information¹⁶, while the state dependent rule incorporates a great deal of information encoded in the state, resulting in very high mutual information. As Woodford [2008] and Costain and Nakov [2011b,a] point out, if the additional cognition necessary to account for state when adjusting is subjectively costly, agents may avoid using purely state dependent rules for economically sensible reasons.

To model this idea, suppose agents experience a psychic cost from implementing a decision rule that deviates in terms of relative entropy from the very simple Calvo rule. While using a mechanical Calvo rule generates no psychic costs, switching to a different adjustment probability (as the state changes) causes such costs to rise proportionally to the change in entropy. We measure the change in entropy caused by state dependence by calculating the Kullback-Leibler divergence $D_{KL}(\lambda, \tilde{\lambda})$ between the Calvo benchmark ($\tilde{\lambda}$) and the rule actually utilized by the agent (λ). The model is parameterized by a marginal cognitive cost κ .¹⁷

This cost of accounting for state can be added to the standard problem in a straightforward way. When an agent chooses to adjust to z^* , she still gets her expected discounted value at z^* (net of the menu cost, m), $[v(z^*) - m - v(z)]$, where $v(\cdot)$ is the value

¹⁶See Cover and Thomas [2012] for a formal definition of mutual information and its relation to entropy and relative entropy.

¹⁷The idea that decision control produces a disutility proportional to the reduction in relative entropy of choices follows Stahl [1990] and Mattsson and Weibull [2002]. This type of assumption provides the microfoundations for many static models of probabilistic choice with important empirical applications, such as quantal response equilibrium (McKelvey and Palfrey [1995]). For a recent and related application of this assumption, see Costain and Nakov [2013] who model near-rational retail price adjustments and study its macroeconomic implications.

function¹⁸. However, she also suffers the cost of control just described if she accounts for the state at all in her decisions. Adding in the remaining ingredients – the instantaneous profit flow, $g(z)$, and the capital gain-like term $\frac{1}{dt}E_t dv(z(t)) = [v'(z(t))(\sigma^2 - \mu)z(t) + \frac{1}{2}v''(z(t))\sigma^2 z^2(t)]$ – the Hamilton-Jacobi-Bellman equation of the problem becomes:

$$rv(z) = g(z) + v'(z)(\sigma^2 - \mu)z + \frac{1}{2}v''(z)\sigma^2 z^2 + \max_{\lambda} \left\{ \lambda [v(z^*) - m - v(z)] - \kappa D_{KL}(\lambda, \tilde{\lambda}) \right\} \quad (2.9)$$

where the reduction in relative entropy achieved by decision rule λ is given by:

$$D_{KL}(\lambda, \tilde{\lambda}) = \lambda \ln \left(\frac{\lambda}{\tilde{\lambda}} \right) + (1 - \lambda) \ln \left(\frac{1 - \lambda}{1 - \tilde{\lambda}} \right) \quad (2.10)$$

It is then straightforward to show that the optimal choice of the adjustment probability, λ , is given by:

$$\lambda(v(z)) = \frac{1}{1 + \frac{1 - \tilde{\lambda}}{\tilde{\lambda}} \exp \left(- \frac{v(z^*) - m - v(z)}{\kappa} \right)} \quad (2.11)$$

This expression reveals that purely state and time dependent rules are special cases of this model, governed by the cost, κ , of responding to the current state. When κ is zero, and cognition is costless, $\lambda = 1$ in the adjustment region and zero in the inaction region. The agent is capable of precisely adjusting as a function of state as visualized in Figure 2.2 (a). In the limit, when κ is infinite, the agent ignores the state, setting a fixed $\lambda = \tilde{\lambda}$ and

¹⁸Note that this model treats z^* as parametric, matching the experimental design in which subjects automatically move to z^* when adjusting.

implementing the Calvo rule pictured in Figure 2.2 (b).

At intermediate values of κ , agents will display a hybrid pattern between these two extremes. Panel (c) of Figure 2.2 plots switching points from a simulation using the exact same $\tilde{\lambda}$ parameter employed in panel (b) but with an intermediate $\kappa = 2$.¹⁹ The results contain ingredients of both the state dependent behavior in panel (a) and the time dependent behavior in panel (b). As in panel (a), behavior is bimodal and peaked near optimal switching points, $\underline{z}$ and $\bar{z}$. As in panel (b), adjustments are diffuse, with frequent adjustments at states far from the optimum (many very near z^*).

The distinctive pattern of “smooth state dependence” plotted in Figure 2.2(c) bears a strong resemblance to patterns documented in field data, as illustrated in Figure 2.1.²⁰ However, key variables (and even the form of the decision problem itself) are generally unobservable in naturally occurring data, making it difficult to directly evaluate such explanations for time dependence. Do decision makers avoid setting purely state dependent adjustment rules even when we know they are acting in a standard Ss environment? Is bounded rationality sufficient to generate the sorts of time dependent patterns observed in the field? In order to directly answer such questions, we turn to the controlled setting of a laboratory experiment.

¹⁹Simulations use the discretization described in Online Appendix 2.6.5.

²⁰The major difference between the simulation shown in Figure 2.2 (c) and the field data plotted in Figure 2.1 is that Figure 2.2 (c) features asymmetric concentrations around the two modes: the left hand mode is much taller than the right hand mode. This is a straightforward consequence of the specific parameters used in the simulations (and in the experiment). Briefly, because $\bar{z}$ is considerably closer to z^* than is $\underline{z}$ in this parameterization, agents adjust back to z^* much more often below z^* than above, creating a much larger concentration of upward adjustments in price than downward adjustments and generating the asymmetric modes observed in the figure (this asymmetry can also be seen in 2.2 (a)). A different choice of parameters would generate symmetric modes (or asymmetries in the opposite direction) in both the state dependent and bounded rationality models.

2.3 Experimental Implementation and Treatment Design

In this section we describe our strategy for taking the model described in section 2.2.1 to the laboratory. In section 2.3.1, we show how we convert the continuous time, infinite horizon model into a closely related near-continuous time, indefinite horizon game. In section 2.3.2, we describe the visual software we built to test the theory and provide details regarding the implementation of the experiment. Finally, in section 2.3.3, we describe and motivate our experimental treatments.

2.3.1 Binomial Conversion

There are two barriers to directly implementing the model described in section 2.2.1 in the laboratory. First, true Brownian motion evolves in infinitesimal time increments and cannot therefore be generated in simulation. Second, the model has an infinite horizon and experimental sessions are obviously of finite duration.

To cope with the first problem, we turn to a well known binomial approximation standardly used to simulate Brownian motion. We divide the game into a number of ticks, each of length Δt . In each tick, the series $P(t)$ iterates up with probability p and down with probability $1 - p$. If the tick moves up, it rises to $(1 + h)P(t)$ for a step size parameter h and, if down, to $(1 - h)P(t)$.

As Δt approaches zero, the parameters converge to the two key Brownian parameters, drift

$$\alpha = \lim_{\Delta t \rightarrow 0} \frac{(2p-1)h}{\Delta t}. \quad (2.12)$$

and volatility

$$\sigma^2 = \lim_{\Delta t \rightarrow 0} \frac{4p(1-p)h^2}{\Delta t}. \quad (2.13)$$

Thus, by setting a very low value of Δt (1/5 of a second in our experiment), we can generate random series that are nearly continuous and closely mimic the Brownian motion described in the theory. See Dixit [1993] for details.

To deal with the second problem, we implement a per-tick probability q of the game ending in lieu of a discount rate. As shown in Oprea et al. [2009a], this converts to the Brownian discount rate in the limit according to the following formula.

$$r = \frac{-\ln(1-q)}{\Delta t}. \quad (2.14)$$

Appendix 2.6.3 presents evidence that our choice of Δt is small enough to generate optima in the binomial approximation that are virtually identical to those predicted in the continuous version of the model.




Figure 2.3: Screen shot of custom experimental software.

2.3.2 Implementation

We implemented the experiment using a custom piece of software programmed in a new Javascript environment called Redwood. The subject display, shown in Figure 2.3, consists of three panels, each visualizing a different part of the decision problem.

First, on the top panel, we show subjects their current price, $p(t)$ and the optimal price $z^*P(t)$ (labeled the “Ideal price” on the screen). Subjects see $p(t)$ as a stable blue line and $z^*P(t)$ as a point fluctuating over time with previous values drawn as a trailing yellow line. This slight change of framing relative to the theoretical discussion above has no effect on the theoretical predictions but makes the problem considerably easier to explain to subjects.

The bottom two panels display detailed information on the real time earnings consequences of subjects’ decisions. The middle panel shows subjects their flow profits. Positive flows are shown as regions shaded in green, negative flows as regions shaded in pink. The bottom line charts the total profits accumulated so far.

A subject’s only decision is when to press a button labeled “Move Current Price.” When this button is pressed, the blue line immediately jumps to join the current Ideal price, maximizing instantaneous earnings, and a cost of m is incurred. The experiment is divided into a sixteen periods (each period is a complete play of an indefinitely repeated game), with the current period ending with a fixed probability, q , in each tick. When the period ends, a red vertical line is shown on the screen and $p(t)$ is automatically and costlessly moved to $z^*P(t)$. By pressing the space bar, the subject begins the new period, restarting the evolution of the Ideal price.

<i>Treatment</i>	Binomial			Brownian			Other			Predictions		
	p	h	q	α	σ	r	m	θ	c	z^*	$\underline{z}$	$\bar{z}$
BASE	0.5	0.0268	0.0012	0	0.06	0.006	5	5	0.4	0.51	0.41	0.68
VOL	0.5	0.0492	0.0012	0	0.11	0.006	5	5	0.4	0.52	0.39	0.80
CORR	0.5	0.0492	0.0012	0	0.11	0.006	1.6	5	0.4	0.51	0.41	0.68

Table 2.1: Parameter values and theoretical predictions for each treatment. In all treatments, the time step in minutes is $\Delta t = 0.0033$.

Data was collected in the LEEPS laboratory at the University of California, Santa Cruz from April to December 2012. A total of 64 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 2.7 , were read aloud prior to the beginning of the experiment. Subjects were told the exact binomial parameters used in the experiment and were given a graph showing how flow profits are impacted by deviations from z^* . Of course, they were not told optimal threshold points (or even that a threshold rule is optimal).

Each period's length was randomly pre-drawn prior to data collection such that all subjects experienced identical period lengths (though of course subjects did not know the duration of any given period until the period ended). Subjects were paid a \$5 showup fee and \$0.02 for each point of profit earned from one randomly selected period. Sessions lasted roughly 1 hour and 30 minutes including instructions and average subject earnings were \$13.60.

2.3.3 Treatments and Motivation

Table 2.1 summarizes our treatment design, displaying binomial parameters used, corresponding approximate Brownian values, other parameters used in the experiment, and both the thresholds of the optimal inactivity bands ($\underline{z}, \bar{z}$) and the optimal price ratio, z^* for each treatment. Our design focuses on two main treatments (BASE and VOL) and one diagnostic treatment (CORR). Each subject was observed under only one treatment (we employed a completely between-subjects design). The treatments were parameterized to achieve several inferential goals.

First, in order to give us sufficient data on adjustments both above and below z^* (critical for fully assessing the predictions of the model) we focused the design on a non-inflationary setting. The probability that relative prices rise and fall were set to $p = 0.5$ generating a drift in the price index of $\alpha = 0$. This means prices will stray below the optimum as often as it strays above, allowing us to evaluate the organizing power of both optimal switching points, $\bar{z}$ and $\underline{z}$, on behavior. The ability to observe adjustments on both sides of z^* is also important for structural estimation because the flow payoff function is highly asymmetric.²¹

Second, we carefully selected our parameters to minimize flat payoffs around the optimum, a perennial concern in implementing dynamic stochastic problems in the lab. We considered dozens of parameterizations before selecting the ones we, in the end, took to the lab to ensure that subjects face substantial costs by deviating from optimal decisions.

²¹While for $z < z^*$ the payoff function is steep and large deviations can easily produce losses, for $z > z^*$ the payoff function is less steep and tends to flatten out above zero. Thus observing subjects' behavior on both sides of z^* allows us to obtain a more powerful estimate of the relation between the value of adjusting at a given state and the adjustment probability. This relation is central to the structural cognitive cost model estimated in Section 2.4.2.

Payoff hills from the parameterizations we selected are shown in Online Appendix 2.6.

Third, we used parameterizations that feature asymmetric switching points below and above z^* in all of our treatments: $\underline{z}$ is substantially closer in each case to z^* than is $\bar{z}$. This feature of the design allows us an additional distinctive theoretical prediction to test: state dependence should cause subjects to adjust sooner and more often below z^* than above.

Fourth, while the central aim of the experiment is simply to assess the degree of state dependence that arises in a controlled setting, we also wanted to study how the degree of state dependence is influenced by the economic environment. We chose to focus the treatment design on volatility, a variable of recent interest in the Ss literature (e.g. Vavra [2013] and Bloom et al. [2007]). Volatility is interesting both because it has a subtle predicted comparative static effect in orthodox models but also because it is a variable with a plausible influence on time dependence-generating bounded rationality. In the high volatility (VOL) treatment we nearly doubled the value of σ relative to the baseline treatment (BASE). This is predicted, in the standard model, to reduce the sensitivity of adjustment decisions to shocks in relative prices in a straightforward way: as shown in Table 2.1, increasing volatility expands the size of the optimal inactivity band $(\underline{z}, \bar{z})$.

Fifth, in order to better understand the effects of volatility observed in the VOL treatment, we added a diagnostic treatment called CORR (volatility, corrected) that features the same high volatility as the VOL treatment, but has the exact same optimal inactivity band as the BASE treatment.²² By comparing CORR to BASE, we can examine

²²This is accomplished by simultaneously lowering the cost of adjustment, m , relative to the BASE treatment.

the effects of volatility in an environment where standard Ss reasoning predicts no effect on the optimal threshold.

2.4 Results

We present the results in three parts. Section 2.4.1 presents reduced form evidence on switching behavior in our main BASE and VOL treatments. We show that while average decisions are meaningfully organized by and responsive to the underlying economic environment (suggesting state dependence), subjects do not choose sharp switching points or set true inaction regions (suggesting time dependence). Reduced form estimates of the adjustment hazard function confirm significant time dependence in our data and suggest that this time dependence is substantially higher in the VOL than in the BASE treatment.

Section 2.4.2 digs deeper into the roots of time dependence by estimating a structural model of bounded rationality, discussed in section 2.2.2 and discretized for estimation in Online Appendix 2.6.5. The results show that subjects experience substantial cognitive costs from responding to the state when timing their adjustments. These costs are estimated to be, on average, three times larger than the monetary switching costs induced in the BASE treatment; in the higher volatility VOL treatment average cognitive costs rise by over 50% to 5 times the induced switching cost.

Finally, in section 2.4.3 we report the results of a robustness treatment, CORR, that features identical volatility to the VOL treatment but identical Ss predictions to the BASE treatment. We show that both reduced form estimates of time dependence and structural estimates of bounded rationality are statistically indistinguishable from those in




Figure 2.4: Weighted histograms of states at which adjustments are made. Solid vertical lines mark $\underline{z}$ and $\bar{z}$ while a dotted vertical line marks z^* .

the VOL treatment but significantly higher than in the BASE treatment, confirming that volatility has a meaningful and fairly stable impact on time dependence.

In all of our analysis, we focus on behavior from the last half of the session (after tick 6000). Thus all behavior analyzed in this section features subjects with considerable experience with the decision making environment. Results are qualitatively similar (though noisier) if we instead use the full sample.

2.4.1 Reduced Form Results

Figure 2.4 shows histograms of the state, $z(t)$, at adjustment across subjects for each of our two main treatments.²³ For reference, vertical red dotted lines show z^* while vertical solid lines show optimal switching points $\underline{z}$ and $\bar{z}$. Several observations are worth making.

²³These histograms are weighted so each subject makes an equal contribution to the plot. Without such weighting, subjects with extreme time dependence will tend to be significantly overweighted in the histogram, giving an unrepresentative portrait of the data.

First, the distributions of adjustment points are bimodal, with modes close to optimal switching points in each case. Second, and relatedly, the modes are substantially further apart in the VOL treatment than in the BASE treatment, as predicted by the Ss model. Third, as predicted by the Ss model in our parameterization, the modes to the left of z^* are considerably higher than to the right: subjects adjust sooner and more often to price decreases than to price increases.

Fourth, and most importantly for our research questions, switching points are not tightly clustered around the optimal thresholds but are instead quite diffuse, with a relatively large share of adjustments occurring far from optimal levels; there is little evidence that subjects establish true inaction regions as predicted by the Ss model. These often sizable deviations are a sign of time dependence, suggesting that subjects are using other (possibly random) factors than the state to trigger adjustment. Indeed, the histograms look strikingly like both the patterns generated by the bounded rationality model, pictured in panel (c) of Figure 2.2, and the field evidence plotted in Figure 2.1.

Finally, adjustments are significantly less tightly clustered around the mode (and the optima, particularly $\bar{z}$) in the VOL treatment than the BASE treatment, suggesting that the degree of state dependence is significantly impacted by the volatility of the underlying economy.

Together, these observations give us a first result:

Result 5. Adjustment distributions are bimodal but diffuse with many switches occurring far from modal levels, closely resembling evidence from the field.

Although our main interest in the data is the degree of state dependence governing

subjects' decisions, it is useful to note that subjects' behavior appears to be responsive to the economic environment in a meaningful way – their decisions are far from random. Figure 2.4 shows adjustments clustered (though very diffusely) around modes located near optimal switch points. To study this more formally, we calculate the median adjustment state for each subject (i) conditional on adjusting below z^* and (ii) conditional on adjusting above z^* .

Below z^* the median adjustment state (across subject-wise medians) is 0.417 in BASE and 0.402 in VOL, strikingly close to the predictions ($\underline{z}$) of 0.41 and 0.39. Even more striking, above z^* , the median adjustment is 0.662 in BASE and 0.795 in VOL, nearly identical to the predictions ($\underline{z}$) of 0.68 and 0.80. Finally, Mann-Whitney tests confirms that median adjustments are significantly higher above z^* in VOL than BASE ($p < 0.001$) and significantly lower below z^* ($p = 0.01$), supporting the key comparative static prediction of the Ss model for our design.²⁴

The data thus suggests that subject behavior is, on average, quite responsive to the economic forces in the model. There is little evidence of systematic bias relative to the optimum and strong evidence that an important and subtle comparative static prediction organizes cross-treatment behavior. This finding gives us a second result:

Result 6. Average switching behavior is not biased relative to optimal state dependent benchmarks, and on average, obeys comparative static predictions.

Although there is little evidence of systematic bias in switching behavior, Figure

²⁴Another predicted consequence of a rise in volatility is an increased rate of adjustment. Simulations suggest we should expect adjustment rates of 0.01 in BASE and 0.017 in VOL. The median adjustment rate (across subjects) is 0.011 in BASE and 0.020 in VOL, indicating a slightly larger than predicted increase in adjustments in the median subject. This difference is only marginally significant ($p = 0.091$, Mann-Whitney test) but is more decisively significant by the final quarter of the session ($p = 0.038$).




Figure 2.5: Estimates of (a) switching hazards and (b) structural cognitive costs.

2.4 shows that subjects do not adjust tightly around thresholds as a state dependent model would suggest. Our motivating hypothesis is that this lack of tightness is due to subjects failing to use consistent switching points (and failing to establish inaction regions) because of time dependence in their adjustment rules. However another possibility is a simple aggregation problem: subjects tend to employ different switching points and this generates apparent variation in adjustments when decisions are aggregated.

In order to sort out whether our results are due to time dependence or aggregation issues, we estimate the parameters of the following simple reduced form hazard function for each subject in our sample.

$$h(\tilde{z}; \alpha, \beta) = \frac{1}{1 + \alpha \exp(-\frac{\tilde{z}}{\beta})} \quad (2.15)$$

where $\tilde{z} \equiv \left| \ln \frac{z}{z^*} \right|$. Here α is a simple shifter of the hazard function (essentially an intercept term) while β measures the slope of the hazard function, a simple measure of time depen-

dence. Under the hypothesis of perfect state dependence, this parameter will be close to zero while a large β suggests insensitivity to state, indicating time dependence. We estimate 2.15 by maximum likelihood for each subject independently and plot CDFs of the critical parameter β for each treatment in panel (a) of Figure 2.5.

Two observations are important. First, values of β are far from zero in each treatment for most subjects, suggesting that hazard functions are flat relative to a state dependent threshold rule for individual subjects. Most individual subjects, therefore, exhibit considerable time dependence in their decision making: the diffuse adjustment pattern observed in 2.4 are not an artefact of aggregation of heterogeneous subject decisions. Second, values of β tend to be far higher in the VOL treatment than in the BASE treatment. Indeed, the median VOL subject's β is roughly double that of the median BASE subject and this difference is highly significant ($p < 0.001$, Wilcoxon test). Estimates thus confirm the pattern observed in Figure 2.4: volatility significantly increases time dependence in subjects' decision making.

Result 7. Hazard functions are gradually sloped across the sample in both treatments, evincing significant time dependence in most subjects' decision making. This time dependence roughly doubles when we increase volatility in the underlying economy.

2.4.2 Structural Results

Subjects show evidence of time dependence in their decision making, generating patterns that resemble those generated by the bounded rationality model discussed in section 2.2.2. In order to better interpret our results, we structurally estimate the parameters

of that model for each subject in our dataset. Our primary interest is to recover estimates of κ , subjects' cognitive costs of implementing state dependent rules. Online Appendix 2.6.5 discretizes the model for purposes of estimation.

Our sample consists of a set of observations $\{(z_t, \delta_t), t = 1, \dots, T\}$, for each subject, where z_t is the relative price at t and δ_t is a dummy variable defined as follows:

$$\delta_t \equiv \begin{cases} 1 & \text{if there is an adjustment at } t \\ 0 & \text{if there is no adjustment at } t \end{cases} \quad (2.16)$$

Thus the hazard function can be defined as $\Lambda(z; \tilde{\lambda}, \kappa)$, the implicit function that assigns to each state z the adjustment probability computed using equation (2.32) and the optimal value function obtained from the Bellman equation (2.33). We estimate the parameters $\tilde{\lambda}, \kappa$ of the hazard function $\Lambda(z; \tilde{\lambda}, \kappa)$ by maximum likelihood:

$$\max_{\kappa, \tilde{\lambda}} \prod_{t=1}^T \Lambda(z_t; \tilde{\lambda}, \kappa)^{\delta_t} \left[1 - \Lambda(z_t; \tilde{\lambda}, \kappa)\right]^{1-\delta_t}. \quad (2.17)$$

We plot the distribution of κ estimates across subjects for each treatment in Figure 2.5. Estimates in each case are bounded away from zero, indicating significant costs of optimizing. As with our reduced form time dependence measure, β , the distributions of κ differ substantially across treatment; at 5.84, the median κ in the VOL treatment is over 50% larger than the median of 3.84 estimated in the BASE treatment. A Wilcoxon test allows us to reject the hypothesis that κ estimates are equal in BASE and VOL subjects ($p = 0.015$). Costs of optimizing are thus economically and statistically significantly higher in the high volatility treatment.

How should we interpret the magnitude of these estimated cognitive costs? The parameter κ is the marginal cost of deviating from a purely time dependent rule, measured in payoff units per nat (a unit of entropy). As such, the scale of κ is not very intuitive. However, our model can be used to produce a more meaningful description of average cognitive costs.

The cognitive cost of deviating from a purely time dependent rule $\tilde{\lambda}$ by choosing to adjust with probability λ is equal to $\kappa \cdot D_{KL}(\lambda, \tilde{\lambda})$. The median estimate of $\tilde{\lambda}$ is very small at 0.019,²⁵ meaning the average subject incurs a considerable cognitive costs when he wants to put a high probability on adjusting at a specific state. In order to adjust with probability $\lambda = 1$ (i.e. to follow a perfectly state dependent rule), the agent thus incurs the following cost:

$$\Omega(\kappa) \equiv \kappa \cdot D_{KL}(1, \tilde{\lambda}) = \kappa \ln \left(\frac{1}{\tilde{\lambda}} \right) \quad (2.18)$$

Intuitively, $\Omega(\kappa)$ is the cost the agent must incur in order to guarantee an adjustment when desired (or to avoid any “slip of the hand” in her decision to adjust at any point in time). That is, it measures the additional costs for a boundedly rational agent to follow a state dependent rule on top of the menu cost of m induced in the design. We evaluate the cost $\Omega(\kappa)$ at the median estimates of κ and $\tilde{\lambda}$ for each treatment and obtain: $\Omega_{\text{BASE}} = 15$ and $\Omega_{\text{VOL}} = 22$. These values are in model payoff units and thus can be directly compared with the menu cost, $m = 5$ in BASE and VOL. The exercise suggests that the costs of perfect adjustment are several times larger than the ones induced in the model. Moreover, these

²⁵That is, the baseline time dependent rule that the median subject adopts when she expends no cognitive costs on the task involves a conditional probability of adjustment of around 2% in each tick.




Figure 2.6: (a) Histogram of switches in CORR and estimates of (b) switching hazards and (c) structural cognitive costs.

costs increase by roughly 50% when we increase volatility.

Result 8. Structural cognitive costs are estimated to be sizable, generating a fixed cost of action equivalent to three times the menu cost on average in the BASE treatment. These costs are highly sensitive to volatility: increasing volatility in the VOL treatment raises cognitive costs to almost five times the menu cost.

2.4.3 Robustness: The CORR treatment

Results so far suggest that time dependence (and cognitive costs) are heavily influenced by volatility in the price index. However, volatility has simultaneous *predicted* effects on behavior, stemming from orthodox Ss forces. In order to get a cleaner view of the relationship between volatility and time dependence, we designed a robustness treatment called CORR (Volatility, Corrected). In this treatment, volatility levels are identical to those in the VOL treatment but we compensate for its effect by lowering m , eliminating the effect of volatility on optimal thresholds relative to the BASE treatment. Under the

hypothesis that volatility itself directly influences time dependence, we expect measures of time dependence in CORR to resemble those in VOL but to be substantially higher than in BASE.

Panel (a) of Figure 2.6 shows a histogram of adjustment states in the CORR treatment, with optima plotted as solid lines. The treatment generates a pattern similar to those in the BASE and VOL treatments: adjustments are bimodal and peaked near optimal levels. However, relative to the BASE treatment CORR switch points are quite diffuse, with many more adjustments clumped inside the predicted inactivity band and a much longer tail above z^* .

Does the increase in volatility in the CORR treatment lead to an increase in time dependence in individual subjects relative to BASE? To find out we estimate the reduced form hazard function (2.15) on the CORR data and plot a CDF of subject-wise β estimates in panel (b) of Figure 2.6. For reference we re-plot the CDFs from BASE and VOL in the same panel. The figure shows that hazards tend to be substantially larger in CORR than in BASE ($p < 0.001$, Mann Whitney test) but are statistically indistinguishable from VOL ($p = 0.678$, Mann Whitney test). Volatility therefore increases time dependence even when its predicted effect on switching behavior is zero. Moreover the effect of volatility on time dependence is quite stable.

We next estimate our structural model on the CORR data and plot a CDF of the cognitive cost parameter, κ , across subjects alongside estimates from BASE and VOL. Cognitive costs increase substantially relative to the BASE treatment ($p = 0.026$) but are indistinguishable from those in the VOL treatment ($p = 0.862$). The data thus also suggest

that the effect of volatility on cognitive costs is relatively stable and robust, even when volatility has no impact on optimal behavior.

Finally, using κ estimates, we find that the median subject would have to expend cognitive costs of $\Omega_{\text{CORR}} = 22$ on top of the actual menu cost to use a purely state dependent rule. This figure is identical to the $\Omega_{\text{VOL}} = 22$ estimated in the VOL treatment and, again, considerably larger than the $\Omega_{\text{BASE}} = 15$ estimated for the BASE treatment. We conclude that

Result 9. Time dependence and estimated cognitive costs in the CORR treatment are similar to those measured in the VOL treatment but considerably larger than in the BASE treatment. This suggests that the effect of volatility on time dependence is relatively stable and is independent of the effects of volatility on theoretical predictions.

2.5 Discussion

As Caplin and Leahy [2010] point out, Ss models have become a “touchstone” of economics, used to explain phenomena ranging from investment dynamics to sticky pricing, hiring decisions to inventory problems. However the predictive and explanatory power of these models is difficult to fully evaluate using conventional empirics. Perhaps the best developed empirical Ss literature so far focuses on price adjustment in the face of menu costs (e.g. Klenow and Kryvtsov [2008], Nakamura and Steinsson [2008], Midrigan [2010]) and findings so far provide a mixed picture of the descriptive accuracy of simple Ss state dependent predictions. On the one hand price adjustments tend to be bimodal and responsive to shocks as state dependent models predict. On the other, adjustments are not tightly

clustered around modal adjustment points, with many changes too small or large to fit the stark predictions of the basic Ss model. Field data instead resembles a hybrid of state dependent adjustment of the sort described by Ss models and the sort of time dependent adjustment described by the famous Calvo [1983] rule.

We report the results of a laboratory experiment designed to not only help interpret findings from this active branch of the empirical literature but also to dig deeper into the fundamental relationship between behavior and Ss theory than has been possible using naturally occurring data. In our experiment subjects repeatedly make adjustment decisions in a precise implementation of a standard menu cost pricing problem (Stokey [2008]). By controlling the parameters and form of the experimental task, we can directly evaluate Ss models and significantly narrow down the set of available explanations for observed deviations from their predictions.

Price adjustments in our data are highly bimodal and diffuse, strongly resembling patterns documented in the field; we are able to “grow” data in the lab with the distinctive characteristics of naturally occurring data. Because we know the precise optima of the laboratory task *ex ante*, we can go further, testing for the first time whether average switching points below and above optimal prices are biased relative to Ss threshold predictions. We show that in fact adjustments are almost exactly equal to optimal Ss thresholds *on average* and almost perfectly obey comparative static predictions: exogenously increasing volatility across treatments widens the range of average adjustments below and above the optimum to almost exactly the degree prescribed by Ss models.

Though it is clear that Ss-like forces powerfully shape average outcomes, it is

equally clear that subjects do not use stable threshold rules and do not establish true inaction regions. Indeed, subjects make numerous and often sizable errors, frequently adjusting far from average (and optimal) adjustment points. As in the field, this pattern suggests that decision rules are not tightly tied to the current state, *contra* Ss models. Unlike in the field, the set of explanations capable of accounting for this “time dependence” is quite small in the lab.

One avenue pursued in the literature for explaining these patterns in the field is to replace the setup of simple Ss models with more elaborate “second generation” models. However, our results suggest that such elaborations may not be necessary: strong evidence of time dependence emerges in our experiment even though we *know* subjects are acting in a simple, first generation Ss environment. An alternative, recently pursued in the literature (Woodford [2008], Costain and Nakov [2011b], Costain and Nakov [2011a], Alvarez et al. [2011]), is to posit that agents experience cognitive or attention costs that can also generate time dependent patterns of adjustment. Such bounded rationality explanations seem to be the only remaining explanation for time dependence emerging in the sterile, controlled setting of our experiment. Though our results cannot, of course, rule out that second generation Ss models apply in some field environments, they do suggest that cognitive costs are very real and are at least a plausible contributing factor to time dependent patterns of adjustment documented in the world.

Using laboratory data, we can generate particularly credible structural estimates of these cognitive costs and explore how they are influenced by the environment. Our estimates suggest that the costs of exercising perfect state dependence are, on average, large relative

to the standard forces in the model (3 times larger than nominal switching costs in the baseline treatment) and evident in most subjects in our sample. More importantly, our results suggest that costs are fundamentally shaped by the economic environment. We find that estimated cognitive costs are over 50% larger in high volatility than in low volatility environments. In our view, this effect of volatility is intuitive: volatility substantially increases the complexity of the optimization task (and demands more attention of the observer), making it considerably more costly to implement a threshold rule. The effect of volatility is also quite stable: estimated cognitive costs in our two high volatility treatments are statistically indistinguishable even though the two treatments generate very different S_s predictions.

Though caution must be taken in projecting results from stylized laboratory experiments onto the field, there are several reasons to suspect our main results are useful for interpreting naturally occurring behavior. First, by the time they enter our analysis, our subjects have a great deal of experience with the decision making environment²⁶ and have developed considerable sophistication, narrowing the gap between our undergraduate subject pool and decision makers in field environments. Indeed, such sophistication is evident in the strong organizing power of theoretical predictions on average behavior. Second, the deviations from theoretical predictions we do observe – time dependent diffusion of adjustment points – are very similar to those observed in naturally occurring data. This match between field and experiment suggests that we are likely measuring phenomena that are not artifacts of the lab. Finally, our ability to crisply interpret these deviations is uniquely

²⁶As we point out above, our analysis focuses entirely on the last half of experimental sessions, meaning all subjects are observed with experience.

enabled by our laboratory environment: precise information on parameters and structure in the decision making environment are generally unobservable in the field. Though laboratory experiments are unlikely to (and shouldn't) supplant traditional empirics, observability of key variables makes laboratory evidence a powerful complement to the ongoing conversation between theory and empirics in this literature.

Still, it is important to bear in mind that the magnitudes of cognitive costs faced by subjects in our experiment are likely to be different from those experienced by the average decision maker in the field. After all, our experiment was designed to minimize the costs of making optimal choices: we do not require subjects to calculate optimal price ratios or current values of the price index and all information is neatly presented to subjects using a highly intuitive graphical interface. That we observe strong evidence of bounded rationality even here is striking. In the field, the problem of choosing optimal adjustment times is likely orders of magnitude more difficult, as agents must gather and interpret noisy information from multiple sources and estimate optimal prices themselves. As a result, our estimates probably substantially understate cognitive costs most decision makers actually face. On the other hand, incentives to optimize are also considerably larger for decision makers in the field, counterbalancing this cost difference to a great degree. Unless cognitive costs and optimization incentives scale identically in the move from lab to field, the degree of time dependence generated by bounded rationality may be larger (or smaller) than that observed in our experiment. However, to the degree bounded rationality is a key driver of time dependence in the field (as it certainly is in the lab), the considerable time dependence documented in recent work suggests that cognitive costs are probably large relative to

incentives as in the lab.

Though the precise magnitudes of our cognitive cost estimates cannot be readily exported to the field, the economic forces shaping them likely can. Perhaps our most empirically relevant result is that state dependence itself is severely and consistently eroded by volatility, a result driven by sharply increasing costs of optimizing in more complex environments. This finding suggests that caution should be taken in interpreting cognitive costs (and perhaps other behavioral measures) as structural parameters, independent of exogenous features of the economic environment. Counterfactual exercises projecting estimates from low volatility environments into high volatility environments are likely to yield misleading policy advice.

The idea that optimization itself can be strengthened or weakened by fundamentals of an economy has potentially important implications for the interpretation of theory in a variety of policy settings. The effectiveness of monetary policy, for instance, depends critically on the degree to which firms price according to state dependent rules. While time dependent pricing is not a necessary condition for monetary policy to have significant real effects, different monetary policies can have different effects under time dependent behavior and under state dependent behavior.²⁷

While much of the empirical research on price adjustment, motivated by this policy

²⁷In time dependent models such as Taylor [1980] and Calvo [1983] monetary shocks are persistent as firms take longer to adjust while in state dependent models monetary shocks have a significantly attenuated role. Caplin and Spulber [1987] construct an example where nominal shocks to the economy do not create disturbances to aggregate output despite the fact that individual firms only respond sporadically to shocks. Danziger [1999] extends this result to show that the stark neutrality can hold when there are both firm level and aggregate productivity shocks. While both of these theoretical results are special cases with reduced form solutions, in a model that matches the frequency of firm level price adjustments Golosov and Lucas [2007] find that the real effects of price disturbances are small. Caballero and Engel [2007] provide a general model to compare the results from the Ss literature with time dependent models and find that variation in the “extensive margin” of which firms adjust in response to shocks is the important feature that reduces the effect of shocks from the purely time dependent model of Calvo [1983].

concern, has investigated *whether* decision makers follow state dependent rules, our study suggests this may not be quite the right question: a better question might be ‘*when* do decision maker follow state dependent rules?’ Treatment-driven comparative statics results suggest that in exceptionally high volatility environments agents may be largely time dependent while in very low volatility environments they may be largely state dependent. Similar such optimization effects may arise in other settings in which the efficacy of policy depends on the degree to which agents reliably optimize. Our results suggest that investigating the relationship between optimization and environment in such settings may be a fertile ground for future research.

2.6 Appendix: Theoretical Details

2.6.1 Full State Dependent Model and Solution Derivation

Following the literature, we assume firms act in a standard Spence-Dixit-Stiglitz demand structure in which a continuum of firms produce a differentiated good. A firm's variety is indexed by j . Consumers have a constant elasticity of substitution utility over consumption of the different varieties:

$$Q \equiv \left[\int_{j \in \Omega} q(j)^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}} \quad (2.19)$$

The consumer's problem is:

$$\begin{aligned} & \max_{\{q(j)\}} Q \\ & \text{s.t.} \\ & \int_{j \in \Omega} p(j)q(j)dj \leq E \end{aligned}$$

Let P be the aggregate price index, defined by $PQ \equiv \int_{j \in \Omega} p(j)q(j)dj$. The optimal solution to the consumer's problem yields demands for each variety:

$$q(j) = \left[\frac{p(j)}{P} \right]^{-\theta} e$$

where $e \equiv \frac{E}{P}$ is the expenditure in real terms, taken as given (following Stokey 2008). The aggregate price index can then be expressed as:

$$P = \left[\int_{j \in \Omega} p(j)^{1-\theta} dj \right]^{\frac{1}{1-\theta}} \quad (2.20)$$

We focus on the problem of a single firm and thus drop the firm index j and treat P as exogenous. We normalize $e = 1$.

Time is continuous and $P(t)$ by assumption follows a geometric Brownian motion with infinitesimal parameters μ and σ :

$$dP(t) = \mu P(t)dt + \sigma P(t)dw(t)$$

where $w(t)$ is a Wiener process. We assume the firm has a constant marginal cost c , in real terms. The firm's instantaneous profits in real terms are then given by:

$$\Pi(p(t), q(t), P(t)) = \frac{1}{P(t)} p(t)q(t) - cq(t)$$

Substituting the demand function, it is possible to rewrite the firm's real profits as:

$$\pi(p(t), P(t)) = \left(\frac{p(t)}{P(t)} \right)^{-\theta} \left(\frac{p(t)}{P(t)} - c \right)$$

Then we use the firm's relative price, $z(t) \equiv \frac{p(t)}{P(t)}$ and rewrite real profits with a change of variable as follows:

$$g(z(t)) = (z(t))^{-\theta} (z(t) - c) \quad (2.21)$$

The value of the firm in the frictionless problem is given by:

$$v^f(z_0) = \max_{\{z(t)\}} E_0 \int_0^\infty e^{-rt} g(z(t)) dt$$

We assume that when the firm adjusts its price it must pay a menu cost of m . The problem of the firm is then to choose the optimal stopping times $\{T_i\}$ and relative prices $\{z_i^*\}$ that maximize its value:

$$\begin{aligned} v(z_0) = & \max_{\{z_i^*\}, \{T_i\}} E_0 \left\{ \int_0^{T_1} e^{-rt} g(z(t)) dt + \right. \\ & \left. + \sum_i \left[-e^{-rt} m + \int_{T_i}^{T_{i+1}} e^{-rt} g(z(t)) dt \right] \right\} \end{aligned} \quad (2.22)$$

Between adjustments, the relative price $z(t)$ follows a geometric Brownian motion. Using the process for $P(t)$ and applying Ito's lemma it is possible to obtain the following:

$$\begin{aligned} dz(t) &= \left[-\frac{z(t)}{P(t)} \mu P(t) + \frac{z(t)}{P^2(t)} \sigma^2 P^2(t) \right] dt - \frac{z(t)}{P(t)} \sigma P(t) dw(t); \\ dz(t) &= (\sigma^2 - \mu) z(t) dt - \sigma z(t) dw(t) \end{aligned} \quad (2.23)$$

The optimal policy is characterized by an inaction region $(\underline{z}, \bar{z})$ and a return point z^* such that if $z(t)$ does not belong to $(\underline{z}, \bar{z})$ then the firm adjusts to z^* and otherwise does not adjust. This implies that outside the inaction region the value function must satisfy:

$$v(z) \leq v(z^*) - m \quad (2.24)$$

Inside the inaction region the value function can be characterized in the following way. Over a small time interval Δt , the following approximation holds:

$$\begin{aligned} v(z(t)) &= g(z(t))\Delta t + \frac{1}{1+r\Delta t} E_t v(z(t+\Delta t)); \\ r\Delta t v(z(t)) &= g(z(t))(1+r\Delta t)\Delta t + E_t v(z(t+\Delta t)) - v(z(t)); \\ rv(z(t)) &= g(z(t))(1+r\Delta t) + E_t \frac{v(z(t+\Delta t)) - v(z(t))}{\Delta t} \end{aligned}$$

Taking the limit $\Delta t \rightarrow 0$ gives:

$$rv(z(t)) = g(z(t)) + \frac{1}{dt} E_t dv(z(t))$$

Using Ito's lemma it is possible to evaluate the expectation term:

$$\begin{aligned} dv(z(t)) &= \left[v'(z(t))(\sigma^2 - \mu)z(t) + \frac{1}{2}v''(z(t))\sigma^2 z^2(t) \right] dt + v'(z(t))\sigma z(t)dw(t); \\ E_t dv(z(t)) &= \left[v'(z(t))(\sigma^2 - \mu)z(t) + \frac{1}{2}v''(z(t))\sigma^2 z^2(t) \right] dt \end{aligned}$$

This leads to the Hamilton-Jacobi-Bellman equation:

$$rv(z(t)) = g(z(t)) + (\sigma^2 - \mu)z(t)v'(z(t)) + \frac{1}{2}\sigma^2 z^2(t)v''(z(t)) \quad (2.25)$$

The Hamilton-Jacobi-Bellman equation for this problem is a second order differential equation with variable coefficients. First I obtain the general solution of the associated homo-

geneous equation:

$$rv(z(t)) = (\sigma^2 - \mu)z(t)v'(z(t)) + \frac{1}{2}\sigma^2 z^2(t)v''(z(t))$$

A particular solution of the homogeneous equation has the form $f_i(z) = z^{R_i}$. Substituting this guess in the equation and solving for R_i gives:

$$\begin{aligned} rz^{R_i} &= (\sigma^2 - \mu)zR_i z^{R_i-1} + \frac{1}{2}\sigma^2 z^2 R_i(R_i - 1)z^{R_i-2}; \\ r &= (\sigma^2 - \mu)R_i + \frac{1}{2}\sigma^2 R_i(R_i - 1); \\ \frac{1}{2}\sigma^2 R_i^2 + \left(\frac{1}{2}\sigma^2 - \mu\right)R_i - r &= 0; \\ R_i &= \frac{-\frac{1}{2}\sigma^2 + \mu \pm \sqrt{\frac{1}{4}\sigma^4 - \sigma^2\mu + \mu^2 + 2\sigma^2 r}}{\sigma^2} \end{aligned}$$

Then, in order to find a solution to the Hamilton-Jacobi-Bellman equation, we use the variation of parameters method. A solution has the form:

$$v(z) = A_1(z)f_1(z) + A_2(z)f_2(z)$$

The functions $A_1(z)$ and $A_2(z)$ are given by:

$$\begin{aligned} A_1(z) &= \int -\frac{1}{\frac{1}{2}\sigma^2 z^2} \left(-\frac{g(z)f_2(z)}{W(z)} \right) dz + B_1 \\ A_2(z) &= \int -\frac{1}{\frac{1}{2}\sigma^2 z^2} \left(-\frac{g(z)f_1(z)}{W(z)} \right) dz + B_2 \end{aligned}$$

where B_1 and B_2 are constants of integration and $W(z)$ is the Wronskian:

$$W(z) = f_1(z)f_2'(z) - f_1'(z)f_2(z)$$

The value function is then given by:

$$\begin{aligned} v(z) = & \frac{2}{\sigma^2(R_2 - R_1)} \left(\frac{z^{1-\theta}}{1-\theta-R_1} + \frac{z^{-\theta}}{\theta+R_1}c \right) + B_1 z^{R_1} + \\ & + \frac{-2}{\sigma^2(R_2 - R_1)} \left(\frac{z^{1-\theta}}{1-\theta-R_2} + \frac{z^{-\theta}}{\theta+R_2}c \right) + B_2 z^{R_2} \end{aligned}$$

where R_1 and R_2 are given above and B_1 and B_2 are constants to be determined.

The value of the firm over the entire state space is then defined by:

$$v(z) = \begin{cases} g(z) + (\sigma^2 - \mu)zv'(z) + \frac{1}{2}\sigma^2zv''(z) & z \in (\underline{z}, \bar{z}) \\ v(z^*) - m & z \notin (\underline{z}, \bar{z}) \end{cases} \quad (2.26)$$

The solution for the value function is obtained requiring that the choice of z^* is optimal and $v(z)$ and $v'(z)$ are continuous at $\underline{z}$ and $\bar{z}$. These conditions are the following:

1. optimal return

$$v'(z^*) = 0$$

2. value matching

$$v(\underline{z}) - v(z^*) + m = 0$$

$$v(\bar{z}) - v(z^*) + m = 0$$

3. smooth pasting

$$v'(\underline{z}) = 0$$

$$v'(\bar{z}) = 0$$

This system of five equation determines the two constants, B_1 and B_2 and the three numbers describing the optimal inaction region: z^* , $\underline{z}$ and $\bar{z}$.

2.6.2 Payoff Hills

Here we plot the payoff hill for each treatment parametrization. The payoff hill plots the value of the policy that uses a pair of thresholds $(\underline{z}; \bar{z})$ (not necessarily the optimal one) for adjusting, evaluated at the starting point z^* . For every pair $(\underline{z}; \bar{z})$ this value function is computed fixing z^* to its optimal level and using the value matching conditions to determine the two constants, B_1 and B_2 .

2.6.3 Solution to the discretized problem

In this appendix we check that the solution to the continuous time model is a good approximation of the optimal rule in the discretized model that subjects actually face in the lab. The discretized problem is characterized by the following Bellman equation:

$$V(z) = \max\{V^n(z); V^a(z) - m\} \tag{2.27}$$

$$V^n(z) = \pi(z)\Delta t + qE_{z'}[V(z')|z] \tag{2.28}$$

$$V^a(z) = \pi(z^*)\Delta t + qE_{z'}[V(z')|z^*] \tag{2.29}$$




Figure 2.7: Payoff hills

where $\pi(z) = z^{-\theta}(z - c)$ and the approximations described in section 2.3.1 apply. We solve this Bellman equation numerically using a spline approximation of the value function. Figure 2.8 plots the value function in the three parametric settings we used in the lab. In each subplot, the two red continuous lines mark the adjustment thresholds $(\underline{z}; \bar{z})$ of the continuous time model. The dotted red line marks z^* . The value function of the discrete model, $V(z)$, has a maximum at z^* . Thresholds in the discrete model are located at the points where the curve becomes flat, i.e. where $V^n(z) = V^a(z) - m$. The adjustment thresholds are virtually identical to the ones obtained from the continuous time model, as can be seen from visual inspection.

2.6.4 Derivation of HJB equation (2.9)

In this appendix we provide a derivation of the Hamilton-Jacobi-Bellman equation 2.9. Consider a small time interval Δt . At each point in time, the value of the firm when the relative price is z , $v(z)$, is given by the flow profit $\pi(z)$ accruing over the time interval Δt and the discounted expected value in the next state. Future payoffs are discounted at rate $\frac{1}{1+r\Delta t}$. Within this time interval the agent adjusts with probability $\lambda\Delta t$. If an adjustment occurs by the end of the time interval the agent gets the value at the return point net of the menu cost, $[v(z^*) - m]$. With probability $1 - \lambda\Delta t$ the adjustment does not occur and the agent gets $v(z')$. When setting the adjustment probability, over the time interval Δt the agent incurs a cost proportional to the Kullback-Leibler divergence $D_{KL}(\lambda, \tilde{\lambda})$ with marginal cost of attention given by κ . Then the agent's value satisfies:

$$\begin{aligned} rv(z) \simeq & g(z)(1 + r\Delta t) + E_{z'} \max_{\lambda} \left\{ \lambda [v(z^*) - m - v(z)] \right. \\ & \left. + (1 - \lambda\Delta t) \frac{[v(z') - v(z)]}{\Delta t} - \kappa D_{KL}(\lambda, \tilde{\lambda}) \right\} \end{aligned}$$

Taking the limit $\Delta t \rightarrow 0$ gives:

$$\begin{aligned} rv(z) &= g(z) + E_t \max_{\lambda} \left\{ \lambda [v(z^*) - m - v(z)] + \frac{1}{dt} dv(z(t)) - \kappa D_{KL}(\lambda, \tilde{\lambda}) \right\}; \\ rv(z) &= g(z) + \frac{1}{dt} E_t dv(z(t)) + \max_{\lambda} \left\{ \lambda [v(z^*) - m - v(z)] - \kappa D_{KL}(\lambda, \tilde{\lambda}) \right\} \end{aligned}$$




Figure 2.8: Value functions

2.6.5 Discrete Time Cognitive Costs Model

In this appendix we describe a discrete time version of the cognitive cost model. At each date, the value of the firm when the relative price is z , $v(z)$, is given by the flow profit $g(z)\Delta t$ and the discounted expected value in the next state. Future payoffs are discounted at rate q . Conditional on the next state, z' , the agent adjusts with probability λ . If an adjustment occurs the agent gets the value at the return point net of the menu cost, $[v(z^*) - m]$. With probability $1 - \lambda$ the adjustment does not occur and the agent gets $v(z')$. When setting the adjustment probability, the agent incurs a cost proportional to the reduction in the entropy of her decision relative to a benchmark (Calvo-like) adjustment probability $\tilde{\lambda}$, measured by the Kullback-Leibler divergence $D_{KL}(\lambda, \tilde{\lambda})$. The marginal cost of attention is κ . The Bellman equation of the agent is:

$$\begin{aligned}
v(z) &= g(z)\Delta t + qE_{z'} \max_{\lambda} \left\{ \lambda [v(z^*) - m] + (1 - \lambda) v(z') - \kappa D_{KL}(\lambda, \tilde{\lambda}) \right\} \\
&= g(z)\Delta t + qE_{z'} \max_{\lambda} \left\{ \lambda [v(z^*) - m - v(z')] - \kappa D_{KL}(\lambda, \tilde{\lambda}) + v(z') \right\}
\end{aligned} \tag{2.30}$$

where the reduction in relative entropy is given by:

$$D_{KL}(\lambda, \tilde{\lambda}) = \lambda \ln \left(\frac{\lambda}{\tilde{\lambda}} \right) + (1 - \lambda) \ln \left(\frac{1 - \lambda}{1 - \tilde{\lambda}} \right) \tag{2.31}$$

The optimal choice of the adjustment probability is given by:

$$\lambda(v(z)) = \frac{1}{1 + \frac{1 - \tilde{\lambda}}{\tilde{\lambda}} \exp \left(-\frac{v(z^*) - m - v(z)}{\kappa} \right)} \tag{2.32}$$

Substituting the optimized values in (2.30) gives the following Bellman equation:

$$v(z) = g(z)\Delta t + qE_{z'} \left\{ \kappa \ln \left[1 - \tilde{\lambda} + \tilde{\lambda} \exp \left(\frac{v(z^*) - m - v(z')}{\kappa} \right) \right] + v(z') \right\} \tag{2.33}$$

2.7 Appendix: Instructions to Subjects

You are about to participate in an experiment in the economics of decision-making. The National Science Foundation and other agencies have provided the funding for this project. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

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

In the experiment you will make decisions over several periods.

At the end of the last period, you will be paid \$5.00, plus earnings based on your profits in a randomly selected period.

You are in control of a company that sells its product at a **Current Price**. Your job is to decide when and how often to adjust the **Current Price**. Market conditions are constantly changing and so the Ideal Price – the price at which your current profit would be maximized – is constantly changing. Each period of the experiment is a series of many **ticks** (5 ticks elapse each second). The relationship between your Current Price and the Ideal Price determines how much you earn each tick.

The further your Current Price is away from the Ideal Price (either too high or

too low), the less money you make during the tick. If your Current Price is especially far from the Ideal Price you may even lose earnings during the tick. At any moment (and as often as you like) you will be allowed to pay a fixed **cost, C** (written on the white board), to move your Current Price to the Ideal Price. At that point the Ideal Price will continue to fluctuate just as before, possibly moving away from the Current Price again.

The screen is divided into three panels as shown in Figure 2.9.

The **Top Panel** shows you the **Current Price in blue** and the **Ideal Price in yellow**. The rightmost tip (the leading edge) of each shows what is happening **right now**.

The **Middle Panel** helps you understand the consequences of your Current Price. It shows your **Flow Profit**, the amount you are earning (or losing) during each tick. When this is positive (and shaded in green) you are earning money; when it is negative (and shaded in red) you are actually losing accumulated money. By comparing the height of this line to the distance between Ideal Price and Current Price you can get a feeling for how differences between the two influence earnings.

The **Bottom Panel** shows you your **Total Profits**, the total amount of money you have accumulated so far.

At any moment (and as often as you like) you can move your Current Price to the Ideal Price by either pressing the spacebar or clicking the button at the bottom of the screen that says **Move Current Price**. However remember that you will incur a cost each time you do this (that will show up in the Middle Panel and the Bottom Panel).

The experiment will be divided into a number of **periods**. The length of a period




Figure 2.9: Screen example

(in ticks) will be random – you will never know how long a period will last or when it is about to end (details below). When a period ends, a red vertical line will show up on the screen and a new period will begin. At the beginning of a new period, the Current Price will automatically be realigned with the Ideal Price and your Total Profits will be reset to zero.

At the end of the experiment we will pay you for your Total Profits realized in one randomly selected period. This number will be converted to a payment amount according to the formula shown on the white board.

Here are a few more details on the experiment, in case you want to know:

Here are some details on how the **Ideal Price** unfolds:

- Each phase is a series of many “ticks” (5 ticks per second).
- Each tick the **Ideal Price** moves randomly up or down by a fixed percentage, e.g., 2%.
- Upticks are as likely as downticks, e.g., each tick is up with probability 50% or down with probability 50%.
- The period ends with a small probability each tick, e.g., $\frac{1}{2}$ of 1%.
- The Current Price and the Ideal Price will always be identical at the beginning of each period.
- The actual values (for ticks per second, tick size, uptick probability, and period ending probability) will be written on the board before the experiment begins.


Here are a few details on how the relationship between Current Price and the Ideal Price determines your earnings in each tick:

- Earnings are highest when your Current Price is equal to the Ideal Price and drop the further they are away from one another. To be more specific, all that actually

matters is the ratio between the Current Price and the Ideal Price. If this ratio drops below 1 or rises above money you will earn less than maximum earnings.

- For instance, you will get the same earnings when your Current Price is equal to 400 and the Ideal Price is equal to 500 as when the Current Price is equal to 800 and the Ideal Price is equal to 1000.
- The exact function that describes the relation between earnings in one tick, the Current Price and the Ideal Price is plotted in Figure 2.10. However you will get a good feel for this relationship from your screen over time.

Figure 2.10: Earnings in a tick



Note: the red dashed line marks the Ideal Price. The blue line shows the level of your profits as the Current price drops below or above this price.

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

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

money, and our interest is in seeing how people make investment timing decisions.

Q2. How long does a phase last? Is there a minimum or maximum?

Answer: The length of time is random. In the example, the probability is 0.005 that any tick is the last, and there are 5 ticks per second. In this case, the average length of a round is 200 ticks or 40 seconds. Many rounds will last less than the average, and a few will last much longer. Rounds longer than 7 minutes are so unlikely that you probably will never see one. The minimum length is one tick, but it is unlikely you will ever see a round quite that short!

Q3. How many rounds will there be?

Answer: Lots. We aren't supposed to say the exact number.

Q4. Are there patterns in upticks and downticks?

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

Chapter 3

Evolutionary Games, Climate and the Generation of Diversity¹

3.1 Introduction

Game theory has been applied to diverse biological systems [Smith, 1982, Huisman and Weissing, 2001, Gardner et al., 2004] to understand the role of frequency dependent selection. The basic idea is to summarize the fitness impact of morph j on morph i in a $n \times n$ payoff matrix $W = ((W_{ij}))_{i,j=1,\dots,n}$. Empirical work to date has relied mainly on bivariate analysis, e.g., the siring success of each male morph is regressed on the frequency of the alternative morphs in the immediate neighborhood, and the estimated coefficients are reported as the rows of the payoff matrix [Sinervo and Lively, 1996, Sinervo and Calsbeek, 2006].

Here we present new methods for estimating payoff matrices, using the time se-

¹This is a joint work with Daniel Friedman, Barry Sinervo and Dhanashree Paranjpe.


ries of allele frequencies in the overall population and a broader definition of fitness, and experimentally test for the rock-paper-scissors dynamic by manipulating allele frequencies. We obtain maximum likelihood estimates for a structural model in which changes in the allele frequencies are governed by standard discrete time replicator dynamics, together with stochastic disturbances following a Dirichlet distribution (or alternatively, a logistic distribution.) Natural extensions of the structural model allow us to estimate the impact of exogenous covariates such as heat spells [Sinervo et al., 2010] that might limit progeny recruitment in a morph specific fashion, and endogenous covariates such as a female population density game that is linked to the male RPS game [Sinervo et al., 2000].

Color expression in male side-blotched lizards appears to have a simple Mendelian pattern of inheritance; supporting evidence includes controlled laboratory crosses [Sinervo et al., 2001], field pedigree [Sinervo et al., 2006, Zamudio and Sinervo, 2000], linkage disequilibrium gene mapping studies [Sinervo et al., 2006] and theoretical models [Sinervo, 2001]. Males with at least one orange allele ($O=oo, bo, yo$) express an aggressive phenotype and patrol large territories with many females. The aggressive orange strategy is vulnerable to invasion by sneaker males with yellow alleles ($Y= by, yy$), which are female mimics and can cuckold the orange males by visiting females in their harems. The strategy of yellow is beaten by the mate guarding and cooperative strategy of blue males ($B=bb$) that defend adjacent territories, but aggressive orange beats blue to create a rock-paper-scissors evolutionary dynamic.

In addition to a natural population called Main World (MW), our study also includes data from three experimental populations. In 1999 we seeded three populations (rock

outcrops) with progeny from controlled laboratory crosses, and designated them as Orange World (OW), Blue World (BW) and Yellow World (YW) based on the color allele that was made most common [Sinervo et al., 2001]. Experimentally altering morph frequencies to different frequencies in a single year allows us to identify impacts of exogenous factors like weather on the temporal / spatial autocorrelation among populations. Data are plotted in figure 3.1.

Figure 3.1: Fractions of observed OBY alleles observed in four *Uta stansburiana* populations.



3.2 Results

As explained in Methods below, we develop a model for discrete time replicator dynamics with a stochastic structure [Calsbeek et al., 2002] given by the Dirichlet distribution or, alternatively, by the logistic distribution. We derive the likelihood function for that structural model, compute the three morph allele shares (alternatively the morph phe-

notype shares) in each world each year, and find the maximum likelihood (ML) estimates of a 3×3 fitness (or payoff) matrix $W = ((W_{ij}))_{i,j=1,2,3}$ that best fit the share data.

Panel (a) of Table 3.1 presents the basic ML estimates, which assume that the same fitness matrix applies each year in each world. Note that (to a good degree of confidence) the highest entry in column 1 is $W_{13} \approx 0.17$, indicating that strategy 3 (yellow, sneaker) is the best reply to strategy 1 (orange, aggressive). Likewise the data confirm that strategy 2 (blue, bourgeois) is the best reply to 3, and that strategy 1 is the best reply to 2. Thus the estimates confirm the RPS best reply cycle. However, the estimated matrix is what Sinervo and Calsbeek [2006] call apostatic RPS (rather than pure RPS) in that each strategy is its own worst reply.

A more complete model allows covariates to amplify (or diminish) each payoff matrix entry. A priori, the most important covariates appear to be the population density that year (relative to its time average level for that world), and temperature as captured in a measure of excessive heat during the key time period. ML fits of this more complete model appear in panels (b), (c), and (d) of Table 3.1. The apostatic RPS structure remains after including the covariates, while the constraint $W_{oo} \geq 0$ no longer binds, perhaps due to negative effect that higher temperatures have on orange when prevalent. The main density effects are (a) negative for all morphs when orange is prevalent, as would be expected for the despotic orange strategy [Sinervo and Calsbeek, 2003], and (b) positive for orange when yellow is prevalent, possibly indicating a fitness boost for orange when it is rare and conditions are conducive to a large population. The strongest temperature effect is a boost to yellow when it is prevalent. The covariates are jointly highly significant according to the

standard likelihood ratio test.

<table><tr><td>0.0000 (0.0352)</td><td>0.2462 (0.0420)</td><td>0.0207 (0.0472)</td></tr><tr><td>0.1383 (0.0469)</td><td>0.0600 (0.0232)</td><td>0.1941 (0.0456)</td></tr><tr><td>0.1701 (0.0527)</td><td>0.1705 (0.0395)</td><td>0.0000 (0.0319)</td></tr></table>	0.0000 (0.0352)	0.2462 (0.0420)	0.0207 (0.0472)	0.1383 (0.0469)	0.0600 (0.0232)	0.1941 (0.0456)	0.1701 (0.0527)	0.1705 (0.0395)	0.0000 (0.0319)	<table><tr><td>0.0052 (0.0216)</td><td>0.2583 (0.0548)</td><td>0.0273 (0.0477)</td></tr><tr><td>0.0578 (0.1058)</td><td>0.0401 (0.0363)</td><td>0.3304 (0.1064)</td></tr><tr><td>0.0634 (0.0765)</td><td>0.2174 (0.0704)</td><td>0.0000 (0.0380)</td></tr></table>	0.0052 (0.0216)	0.2583 (0.0548)	0.0273 (0.0477)	0.0578 (0.1058)	0.0401 (0.0363)	0.3304 (0.1064)	0.0634 (0.0765)	0.2174 (0.0704)	0.0000 (0.0380)
0.0000 (0.0352)	0.2462 (0.0420)	0.0207 (0.0472)																	
0.1383 (0.0469)	0.0600 (0.0232)	0.1941 (0.0456)																	
0.1701 (0.0527)	0.1705 (0.0395)	0.0000 (0.0319)																	
0.0052 (0.0216)	0.2583 (0.0548)	0.0273 (0.0477)																	
0.0578 (0.1058)	0.0401 (0.0363)	0.3304 (0.1064)																	
0.0634 (0.0765)	0.2174 (0.0704)	0.0000 (0.0380)																	
(a) Estimated W in simple model	(b) Estimated W in full model																		

<table><tr><td>-2.5329 (4.6263)</td><td>0.0456 (2.5650)</td><td>4.7470 (2.7592)</td></tr><tr><td>-3.0065 (4.8190)</td><td>3.4059 (5.5519)</td><td>1.1364 (2.9938)</td></tr><tr><td>-3.3406 (3.0572)</td><td>2.1167 (3.1126)</td><td>-2.5715 (4.5553)</td></tr></table>	-2.5329 (4.6263)	0.0456 (2.5650)	4.7470 (2.7592)	-3.0065 (4.8190)	3.4059 (5.5519)	1.1364 (2.9938)	-3.3406 (3.0572)	2.1167 (3.1126)	-2.5715 (4.5553)	<table><tr><td>3.4187 (4.1768)</td><td>-1.4064 (2.6704)</td><td>-1.9101 (3.8363)</td></tr><tr><td>-0.3363 (3.6359)</td><td>-1.5314 (4.0852)</td><td>-1.8428 (3.3206)</td></tr><tr><td>-1.1370 (3.0220)</td><td>-1.5874 (2.7406)</td><td>6.3326 (4.5957)</td></tr></table>	3.4187 (4.1768)	-1.4064 (2.6704)	-1.9101 (3.8363)	-0.3363 (3.6359)	-1.5314 (4.0852)	-1.8428 (3.3206)	-1.1370 (3.0220)	-1.5874 (2.7406)	6.3326 (4.5957)
-2.5329 (4.6263)	0.0456 (2.5650)	4.7470 (2.7592)																	
-3.0065 (4.8190)	3.4059 (5.5519)	1.1364 (2.9938)																	
-3.3406 (3.0572)	2.1167 (3.1126)	-2.5715 (4.5553)																	
3.4187 (4.1768)	-1.4064 (2.6704)	-1.9101 (3.8363)																	
-0.3363 (3.6359)	-1.5314 (4.0852)	-1.8428 (3.3206)																	
-1.1370 (3.0220)	-1.5874 (2.7406)	6.3326 (4.5957)																	
(c) Estimated β_ν	(d) Estimated β_τ																		


Table 3.1: Estimates of the $(W, \beta_\nu, \beta_\tau)$ model (with bootstrapped standard errors in parentheses). Matrix entries are for strategies $[1, 2, 3] = [o, b, y]$. In Panel (a) the covariates are suppressed, and the estimated effective sample size is $\hat{\delta} = 21.7 \pm 3.4$. The remaining panels show estimates of the full model, for which $\hat{\delta} = 34.2 \pm 1.2$.

The estimated matrices have implications for dynamic behavior. If the temperature and density covariates were always fixed at their average levels, then the estimates in Table 3.1 imply that from any interior initial condition (including the large experimental perturbations to OW, BW and YW) the state would converge in finite time to the equilibrium $S^* = \lambda W^{-1} \mathbf{1} \approx (0.38, 0.40, 0.22)$, as illustrated in figure S2. Even for fairly extreme initial conditions, near a corner of the simplex but with each morph share at least 0.1, it takes at most 5 years for the state to converge precisely to S^* . By contrast, convergence is only asymptotic for a true RPS payoff matrix [Gaunersdorfer and Hofbauer, 1995].

Of course, the covariates are not fixed, and indeed vary stochastically from year to year and across locations. Climate directly influences the morph payoffs via progeny recruitment during hot versus cool summers, and population density has an indirect in-

fluence. To assess the impact, we compute the effective payoff matrix $\hat{\Omega}_t^k$ for each World ($k = M, O, B, Y$) and each year t by amplifying or diminishing each payoff entry by according to the density and temperature variables experienced in that world that year. The blue dots in figure 3.2 show drastic year-to-year changes in the Main World equilibrium, i.e., the Nash equilibrium (here, also the ESS) for the payoff matrix $\hat{\Omega}_t^M$. In many years the equilibrium lies on an edge, indicating under the conditions experienced that year one morph (the one fixated at the corner of the simplex opposite that edge) is in danger of local extinction.

Figure 3.2: Simplex and States in all Worlds. Black points and lines are time series $S(t)$, while blue points are the Nash equilibrium for $\hat{\Omega}_t^{xW}$.



We tested the proposition that climate can fix the game on an edge using published phylogeographic data on morph frequency collected across the range of the side-blotched lizard in the United States and Mexico, in which we have observed 7 independent losses of the yellow allele and two losses of the blue allele [Corl et al., 2010]. We estimated hours of temperature restriction (see Supporting Information), h_r , using maximum daily air temperature, T_{max} , obtained from worldclim.org, the relationship between monthly T_{max} and monthly h_r at Los Banos, published data on body temperature of *Uta*, T_b , across its range, and a published physiological model of climate forcing relating h_r to the difference between T_{max} and T_b [?]. The predicted NE from our climate forcing model of the RPS game, when driven by h_r estimated in the local populations, accurately predicts the observed frequency of o alleles ($F_{1,39}=17.87$, $P < 0.0001$), b alleles ($F_{1,39}=22.78$, $P < 0.0001$) and y alleles ($F_{1,39}=5.48$, $P < 0.03$) (figure 3). The correlations between observed and predicted frequency of o ($P = 0.02$) and b alleles ($P = 0.003$) are also significant using phylogenetic independent contrasts. Finally, the hours of restriction observed across the geographic range is negatively related to the observed frequency of orange ($F_{1,39}=9.98$, $P < 0.004$) but positively related to blue color strategies ($F_{1,39}=5.98$, $P < 0.03$), and yellow strategies ($F_{1,39}=8.21$, $P < 0.007$), with the strategy yellow perilously close to fixation under extremely cool conditions of h_r found in northern populations, where orange is likewise strongly favored.

3.3 Discussion

To summarize, our first substantive finding is that the overall payoff matrix estimated for male side-blotched lizards from one generation to the next is of the type called apostatic RPS [Sinervo and Calsbeek, 2006] — both rare strategies have higher fitness than the common strategy, and the best replies are intransitive. Using subscripts $\{o, b, y\}$ instead of $\{1, 2, 3\}$ to denote the orange, blue and yellow color morphs, the estimated fitness orderings are: $W_{yo} > W_{bo} > W_{oo}$, $W_{ob} > W_{yb} > W_{bb}$ and $W_{by} > W_{oy} > W_{yy}$.

Previous estimates of the male lizards' siring success using other methods [Sinervo and Calsbeek, 2006] indicated a pure RPS type of payoff matrix: $W_{yo} > W_{oo} > W_{bo}$, $W_{ob} > W_{bb} > W_{yb}$ and $W_{by} > W_{yy} > W_{oy}$. The discrepancy relative to the previous estimates consists of reduced fitness for common strategies. Indeed, our current estimates put all three common fitnesses (the matrix entries along the main diagonal) near zero.

How might the discrepancy arise? Mindful that previous estimates use a narrower notion of fitness (siring success) than do current estimates (male offspring surviving to adulthood), we can think of two plausible sources of the discrepancy. First, male hatchlings of the common type may have lower first year survival rates, arising from strong selection against self genotypes exerted either by surviving adult males that are still present when progeny emerge or by the progeny themselves as they establish territories [Bleay et al., 2007]. Second, as argued by Alonzo and Sinervo [2001], females may prefer rare male genotypes, and more frequently select their sperm [Zamudio and Sinervo, 2000, Calsbeek et al., 2002], chose them as mates [Bleay et al., 2007, Lancaster et al., 2009], or as suggested by [Alonzo and Sinervo, 2007], females may selectively produce female progeny when the common male

strategy is disadvantaged and the rare strategies are advantaged in the progeny generation.

The dynamic consequences of the new estimates deserve further comment. The picture emerging from previous estimates of a stable pure RPS matrix was of a gradual inward clockwise spiral perturbed a bit by noise. The new estimates of an apostatic RPS matrix instead imply that the unperturbed dynamics are highly stable and only a couple of clockwise cycles would suffice to converge precisely to the stable mix of the three morphs. However, our second substantive finding is that the covariates imply large perturbations in most years (indeed, the equilibrium is often located on an edge of the simplex), so the system is chasing (with a clockwise approach) an erratic moving target.

Besides a better qualitative account of dynamics in each world separately, the new picture can explain the observed approximate entrainment, in less than a decade, of the RPS cycle across four worlds whose cycles were initially set completely out of phase. Temperature, one of the two covariates considered, perturbs the dynamic path similarly in each world each year, and thus suggests entrainment. Consequently the assumption that low spatial correlation sustains diversity (e.g., [Kerr et al., 2002]) may need reconsideration.

The new picture can also explain morph variation across the geographical range of the side-blotched lizard. In northern populations longterm climate variation pushes the NE close to the orange-blue edge of the simplex. Indeed, many northern populations in both the Great Basin and the Coast Range of California exhibit fixation on orange and orange-blue combinations, with yellow being the most susceptible to selective loss [Corl et al., 2010]. This loss likely arises because yellow sires derive most paternity fitness from later season clutches [Zamudio and Sinervo, 2000], but the cooler and shorter season in

northern populations removes the advantage of yellow males, fixing the system on orange and orange-blue combinations. This process generates considerable mating system diversity, but only at the large spatial scale of the entire geographic range of the species.

One can imagine more sophisticated structural models of morph dynamics than considered here. We have treated population density as exogenous, but surely it is partly endogenous and a key aspect of female strategies. One might also like to know more about different fitness components, in particular, those arising from female mate preference, male siring success, progeny survival, and sex ratio modulation [Alonzo and Sinervo, 2007]. One could in principle build a structural model of diploid replicator dynamics with multiplicative components that account for these fitness components while endogenizing density. Indeed, one can imagine structural models that incorporate nonlinearities assumed away when a matrix is used to model fitness payoffs. However, available *Uta* data comprise relatively short time series and a modest number of distinct populations, and thus are clearly inadequate for estimating models with far more parameters than ours. Nevertheless, as a final robustness check, we applied the relatively parsimonious diploid sexual model of Sinervo [2001] to our estimates, and recovered a rapid return to the NE consistent with rapid 6 generation cycles (see Fig S3), in contrast to all previous estimates of male payoffs which require much longer to arrive at the NE [Sinervo, 2001].

The techniques introduced here do seem widely applicable when the concern is overall fitness rather than its components and when frequency dependent selection can be captured in a payoff matrix. Inviting applications include bacterial RPS systems [Kerr et al., 2002] and Y-linked RPS systems in fishes [Bourne et al., 2003]. Our techniques can

easily be specialized to two-morph systems such as the spadefoot toad polyphenism [Pfennig, 1992a,b] when there are experimental or natural perturbations that push the system away from equilibrium. (In general, two strategy one-population games are unlikely to cycle unless they involve density regulation, as with the r-K strategies of *Uta* females noted earlier, or mammalian Chitty cycles [Chitty, 1996].) Extensions of our techniques to more than three morphs (e.g., five as in *Poecilia parae*; see Bourne et al. [2003]) in a single population seem straightforward, and we see no obstacle to applying them to frequency dependent selection in two or more interacting populations such as lynx hare cycles [Krebs et al., 1995], competition cycling dynamics [Huisman and Weissing, 1999], or host parasite cycling dynamics [Hudson et al., 1992]. Given the analogous processes maintaining diversity in mating systems and ecosystems [Sinervo and Calsbeek, 2006, Huisman and Weissing, 1999], our approach highlights the role of climate in potentially reducing ecosystem diversity at small spatial scales, but increasing diversity at large spatial scales.

3.4 Methods

Data on genotypes. The primary raw data are annual counts of individual males, by genotype and location, gathered as part of an annual census of all adult *Uta*. Since 1990 the census has covered the Main World (MW), consisting of connected rock outcroppings covering approximately 0.25 hectares, and since 2000 it has also covered the three separate experimental locales designated Orange, Blue and Yellow Worlds (OW, YW, BW). These censuses record the OBY allele expressed in each mature male. We also estimated color allele frequency for the OW ($f_o = 0.52, f_b = 0.16, f_y = 0.32$) , YW ($f_o = 0.26, f_b = 0.18, f_y =$

0.56), and BW ($f_o = 0.17, f_b = 0.44, f_y = 0.39$) in 1999 from the predicted distribution of color alleles in progeny, because they were produced from controlled laboratory crosses [Sinervo et al., 2001].

Color expression begins in early March as males establish territories and peaks in early April, but the four locales are slightly out of synchrony, with expression of color (and reproduction) initiated first in OW, then MW, BW, and YW. Accordingly, we counted populations in that order, ensuring that males had nearly full color expression on the first census. We also carried out a second and third censuses at monthly intervals to estimate the number males missed on the first pass. Based on 2-3 passes during the reproductive season [Sinervo et al., 2006], we infer that the count includes approximately 98% of male recruits, i.e., surviving males hatched since the previous year's census. There are also a few males that survive for two or more seasons, but *U. stansburiana* is largely an annual at Los Banos.

The 6-vector $(N_{oo}, N_{ob}, \dots, N_{yy})$ of genotype data for a given year and location is reduced to a 3-vector of allele states simply by taking the obvious weighted sums (see Table S1), e.g., $N_o = 2N_{oo} + N_{ob} + N_{oy}$ and, as noted in the body of the paper, the shares are then computed by dividing each component of (N_o, N_b, N_y) by the sum N of the components. Due to diploidy, the population size is $n = N/2$, and the density covariate for each location is defined as n minus its historical mean for that location. Alternatively, one can compute phenotypic state S^P by setting $N_O = N_{oo} + N_{ob} + N_{yo}$, $N_B = N_{bb}$ and $N_Y = N_{yy} + N_{by}$ and dividing each component by the population size, which here is simply the sum $n = N_O + N_B + N_Y$.

The other source of raw data is weather records collected at the California Department of Forestry Station on Gonzaga Road from 1986-2012 (station LBN: 37.053 N, 121.049 W, only 7.33 km from the main population), which provide hourly temperature records throughout the year. However, two years (1989, 1992) were missing records during the summer months, and here we interpolated from an adjacent station. Specifically, we collected two years of data surrounding the missing years from the Los Banos Dam site (LSB: 36.994 N, 120.93 W, and 10.77 km from the main population); regressed the LBN station data on contemporaneous data from the LSB station; and then used the fitted equation to predict the missing data.

For each month, we computed the number h_r of hours when $T_{air} > T_{preferred} = 37.4$ C [Paranjpe et al., 2013]. The idea is to estimate the duration of hatchlings' restricted activity spells due to high air temperatures, which might impact recruitment and survival, and may affect some color morphs more than others. Their small size (≈ 0.5 g) means that hatchlings will be close to thermal equilibrium with the air temperature, particularly because they spend much of their time in shade as the temperature warms, finally retreating entirely at mid-day during the summer when temperatures exceed $T_{preferred}$.

Statistical Procedures. We fit the data to a structural model estimated via maximum likelihood (ML) techniques. The model joins discrete time replicator dynamics with a stochastic structure [Calsbeek et al., 2002] given by the Dirichlet distribution or, alternatively, by the logistic distribution. The goal is to estimate a 3×3 fitness (or payoff) matrix $W = ((W_{ij}))_{i,j=1,2,3}$.

The first step is to take the raw census counts (considering for the moment only

a single locale) of adult males by morph allele N_o, N_b, N_y and $N_{tot} = N_o + N_b + N_y$ (alternatively, by morph phenotype: N_O, N_B, N_Y) and to compute the relative frequencies $S_1 = \frac{N_o}{N_{tot}}$, $S_2 = \frac{N_b}{N_{tot}}$, and $S_3 = \frac{N_y}{N_{tot}}$. Thus the main data input is a $T \times 3$ matrix that lists the state vector $S(t) = (S_1(t), S_2(t), S_3(t))$ each period $t = 1, \dots, T$. Each population state $S(t)$ lies in the *simplex*, the two dimensional triangle $\mathcal{S} = \{(s_1, s_2, s_3) \in \mathbb{R}^3 : s_i \geq 0, \sum_i s_i = 1\}$.

The fitness $W_i(t) = W_i \cdot S(t) \equiv \sum_j W_{ij} S_j(t)$ of morph $i = 1, 2, 3$ at time t is simply its weighted average fitness against the current population frequency. Thus mean fitness at time t is $\bar{W}(t) = S(t)W \cdot S(t) \equiv S(t) \cdot (W_1(t), W_2(t), W_3(t)) \equiv \sum_i \sum_j W_{ij} S_i(t) S_j(t)$. Deterministic *replicator dynamics* (continuous time: ?; discrete time: Smith [1982]) give the next period population state $Z(t+1) \in \mathcal{S}$ that would arise absent stochastic disturbances, and are specified by the difference equation

$$Z_i(t+1) = \frac{W_i(t)}{\bar{W}(t)} S_i(t). \quad (3.1)$$

The most natural way to specify the stochastic part of the model is to use the distribution created for this role, the Dirichlet, so we complete the structural model via $S(t+1) \sim \text{Dir}(\delta Z(t+1))$, where $\delta > -1$ is the effective sample size. This implies that $S_i(t+1)$ is a random variable with mean $Z_i(t+1)$ as specified in the deterministic model, and variance $\frac{Z_i(t+1)[1-Z_i(t+1)]}{\delta+1}$. For further discussion of the Dirichlet distribution and alternative stochastic distributions appropriate for our estimation see the Supplementary Information.

Mathematical and biological principles constrain the parameters we seek to estimate. Each fitness matrix entry $W_{ij} \geq 0$ because, even when most prevalent, a morph j at worst can drive fitness (and thus next period's share) of morph i to zero; negative shares

are not possible. Stochastic disturbances are constrained to keep the state vector within the simplex, but the Dirichlet distribution automatically enforces that constraint. Since matrix entries W appear in both the numerator and denominator of equation (3.1), there is no effect when we multiply all matrix entries by an arbitrary constant $c > 0$. Hence we need to normalize the matrix, and do so via the constraint $\sum_i \sum_j W_{ij} = 1$, i.e., we normalize the original matrix by multiplying by $c = \frac{1}{\sum_i \sum_j W_{ij}}$.

Additionally, we can augment the regression with covariates that capture time- or world-specific shifts in the fitnesses. Probably the most important time-specific covariate is temperature, for which we use $\tau_k(t) = h_r(t) - \bar{h}_r^k$, a measure of temperature restriction the previous summer (when progeny are recruiting) normalized for each world k by subtracting the historical mean value $\bar{h}_r^k$ over the life of that world. For world-specific shifts we take $\nu_k(t)$, the male population size in world k in year t , normalized by subtracting its world-specific mean value. We assume that these covariates shift fitness multiplicatively, with the proportional change in fitness coefficient a linear function of each variable. Thus we will estimate the generously parameterized fitness matrices $\Omega_k = ((\Omega_{ijk}))_{i,j=1,\dots,3}$ where

$$\Omega_{ijk}(t) = W_{ij} \exp(\beta_{ij}^\nu \cdot \nu_k(t) + \beta_{ij}^\tau \cdot \tau_k(t)) \quad (3.2)$$

The normalization of ν and τ ensures that we have $\Omega_k(t) = W$ when both variables are at their average levels. In the most complete specification, therefore, in addition to the 3×3 fitness matrix W we are also estimating two other 3×3 coefficient matrices: β^ν for population density effects and β^τ for temperature effects.

Substituting Ω_k for W in equation (3.1), one sees again that the β matrices need

normalization. The natural choice is that the coefficients sum to zero. Thus estimation of equation (3.2) is subject to the restrictions $W_{ij} \geq 0$, $\sum_{i,j} W_{ij} = 1$, $\sum_{i,j} \beta_{ij}^\nu = 0$, $\sum_{i,j} \beta_{ij}^\tau = 0$.

Analytical Background. The discussion of dynamic implications presumes several properties of equilibria, e.g., that the payoff matrix $\hat{\Omega}_t^k \equiv \Omega_k|_{\tau_k(t), \nu_k(t)}$ has a unique stable interior equilibrium when the temperature and density covariates $(\tau_k(t), \nu_k(t))$ are sufficiently close to zero. That discussion is justified by three known analytic results (see, e.g., Sandholm [2010]) about equilibria of symmetric games specified by an arbitrary $n \times n$ payoff matrix W : (a) If it lies in the interior of the simplex, the vector $S^* = cW^{-1}\mathbf{1}$ is a (mixed) Nash equilibrium, where $\mathbf{1}$ is the vector of 1's and the normalizing constant satisfies $c^{-1} = \mathbf{1}W^{-1} \cdot \mathbf{1}$. Moreover, if the payoff matrix is apostatic RPS, then (b) S^* is the only NE and it indeed lies in the simplex interior, and (c) S^* is also an ESS, hence dynamically stable.

The estimation procedures rely on known facts regarding the Dirichlet distribution, as well as an apparently new derivation of likelihood functions for that stochastic structure applied to replicator dynamics on the 3-simplex $\mathcal{S} = \{(s_1, s_2, s_3) \in \mathbb{R}^3 : s_i \geq 0, \sum_i s_i = 1\}$. A reason for not using the better-known approach, adding a normally-distributed error to deterministic replicator dynamics, is that it leads to nonsensical results, because there is no natural way to ensure that the stochastic dynamics remain in the simplex $\mathcal{S}$.

Given an arbitrary probability vector over 3 outcomes, $p = (p_1, p_2, p_3) \in \mathcal{S}$, the multinomial distribution with δ independent draws gives the probabilities of all possible counts of the three outcomes. The Dirichlet distribution $Dir(\delta z)$ is the conjugate prior,

i.e., given any count vector $\delta z = (\delta z_1, \delta z_2, \delta z_3)$, it gives the probability that the underlying p -vector was $x \in \mathcal{S}$. The Dirichlet density function is

$$f(x_1, x_2, x_3 | \delta, z_1, z_2, z_3) = K(\delta Z) x_1^{\delta z_1 - 1} x_2^{\delta z_2 - 1} x_3^{\delta z_3 - 1} I_{x \in \mathcal{S}}, \quad (3.3)$$

where I is the indicator function which is 0 outside the simplex and 1 on the simplex, and the normalizing constant $K(\delta Z)$ ensures that the density integrates to 1. Textbooks confirm that $x \in \mathcal{S}$ is a random variable with mean z , and that the variance of x_i is $\frac{z_i(1-z_i)}{\delta+1}$. In terms of our structural model, the fitted parameter δ is the effective sample size, which governs the amplitude of stochastic shocks. For large δ , the Dirichlet density (3.3) puts almost all probability mass near the population state $Z(t+1) \in \mathcal{S}$ specified by discrete replicator dynamics; for moderate $\delta > 0$ it picks a random point in the simplex with mean $Z(t+1)$ and its variance is inversely related to δ .

The basic model has parameter vector $\theta = (\delta, W_{11}, W_{12}, \dots, W_{33})$ to be estimated. We assume that the observed state $S(t)$ has the distribution $Dir(\delta Z(t))$, where $Z(t)$ is the row vector of deterministic replicator dynamics with i -th element $Z_i(t) = \frac{W_i(t-1)}{W(t-1)} S_i(t-1)$. Thus $S(t)$ is a random variable supported in the simplex whose i^{th} component has mean $Z_i(t)$ and variance $\frac{Z_i(t)[1-Z_i(t)]}{\delta+1}$. Given the previous state $S(t-1)$, the conditional density of the current state is

$$f(S(t) | S(t-1)) = \frac{\Gamma(\delta)}{\prod_{i=1}^3 \Gamma(\delta Z_i(t))} \prod_{i=1}^3 S_i(t)^{\delta Z_i(t) - 1}. \quad (3.4)$$

This expression spells out the normalizing constant $K(\delta Z)$ in (3.3), using the Gamma

function $\Gamma(z) \equiv \int_0^\infty y^{z-1} e^{-y} dy$.

Summing the log of the right hand side of (3.4) over $t = 1, \dots, T$, the conditional log-likelihood function of the sample is:

$$\ln \ell = T \ln \Gamma(\delta) + \sum_{t=1}^T \sum_{i=1}^3 [-\ln \Gamma(\delta Z_i(t)) + \ln S_i(t)(\delta Z_i(t) - 1)] \quad (3.5)$$

The $t = 1$ terms do not contribute to the sum since, with no previous period to which to apply replicator dynamics, their likelihood is 1. We maximized the right hand side of (3.5) numerically over the set of parameter values $\theta = (\delta, W_{11}, W_{12}, \dots, W_{33})$ satisfying $\delta > -1$, $W_{ij} \geq 0$, and $\sum_{i,j=1}^3 W_{ij} = 1$. The numerical maximization is via Matlab's active set algorithm `fmincon`, using as initial values 1/9 for each entry W_{ij} of the payoff matrix and 20 for δ . The algorithm converged reasonably quickly to the values reported in the Results section for a range of other initial values, e.g., for δ between 15 and 30. The Matlab code is attached at the end of this document.

Extending the basic model to include covariates and panel data is conceptually straightforward. To deal with 4 separate locations ("worlds") we write the conditional log-likelihood function for location k as:

$$\ln \ell_k = T_k \ln \Gamma(\delta) + \sum_{t=1}^{T_k} \sum_{i=1}^3 [-\ln \Gamma(\delta Z_{ik}(t)) + \ln S_{ik}(t)(\delta Z_{ik}(t) - 1)] \quad (3.6)$$

Note that δ and the W_{ij} 's are not location-specific because, to keep the number of free parameters small relative to the overall sample size, we assume that the underlying stochastic structural model is the same in all locations. Of course the stochastic realizations can and

do vary across locations. Adding the covariates ν and τ affects the likelihood function only through the (deterministic) replicator dynamic equation for $Z_{ik}(t)$, where the payoff matrix W is replaced by the augmented fitness matrix $\Omega_k(t)$ defined in equation (2) of the text.

We now assume that stochastic realizations are conditionally independent, i.e., conditional on $S_{ik}(t-1)$ and other covariates, the $S_{ik}(t)$'s are independent across k . Then the augmented model has log-likelihood function

$$\sum_{k \in \{MW, OW, BW, YW\}} \ln \ell_k, \quad (3.7)$$

where $\ln \ell_k$ is defined in (3.6) using the augmented fitness matrices Ω_k in the replicator expression defining $Z_{ik}(t)$. In addition to the coefficient constraints on W mentioned earlier, we also imposed $\sum_{i,j} \beta_{ij}^\nu = 0$, $\sum_{i,j} \beta_{ij}^\tau = 0$, and used the agnostic initial values $\beta_{ij}^\nu = \beta_{ij}^\tau = 0$. To avoid regions of the parameter space where the algorithm failed to converge we also imposed $\beta_{ij}^\nu, \beta_{ij}^\tau \in [-10, 10]$, but these constraints never bind once the algorithm approaches convergence.

Bibliography

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

Andrew B Abel and Janice C Eberly. A unified model of investment under uncertainty. *American Economic Review*, 84(5):1369–84, December 1994.

Suzanne H Alonzo and Barry Sinervo. Mate choice games, context-dependent good genes, and genetic cycles in the side-blotched lizard, *uta stansburiana*. *Behavioral Ecology and Sociobiology*, 49(2-3):176–186, 2001.

Suzanne H Alonzo and Barry Sinervo. The effect of sexually antagonistic selection on adaptive sex ratio allocation. *Evolutionary Ecology Research*, 9(7):1097, 2007.

Fernando Alvarez, Luigi Guiso, and Francesco Lippi. Durable consumption and asset management with transaction and observation costs. *American Economic Review*, 102(5):2272–2300, August 2012a.

Fernando E. Alvarez and Avinash Dixit. A real options perspective on the future of the euro. Working papers, University of Chicago, April 2013.

- Fernando E. Alvarez and Francesco Lippi. Price setting with menu cost for multi-product firms. NBER Working Papers 17923, National Bureau of Economic Research, Inc, March 2012.
- Fernando E. Alvarez, Francesco Lippi, and Luigi Paciello. Optimal price setting with observation and menu costs. *The Quarterly Journal of Economics*, 126(4):1909–1960, 2011.
- Fernando E. Alvarez, Francesco Lippi, and Luigi Paciello. Monetary shocks in a model with inattentive producers. Working papers, University of Chicago, November 2012b.
- Kenneth J. Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, 19(3):205–272, 1951.
- 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.
- Robert J Barro. A theory of monopolistic price adjustment. *Review of Economic Studies*, 39(1):17–26, January 1972.
- 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.
- Giuseppe Bertola and Ricardo J. Caballero. Kinked adjustment costs and aggregate dynamics. In *NBER Macroeconomics Annual 1990, Volume 5*, NBER Chapters, pages 237–296. National Bureau of Economic Research, Inc, January 1990.
- C Bleay, T Comendant, and B Sinervo. An experimental test of frequency-dependent selection on male mating strategy in the field. *Proceedings of the Royal Society B: Biological Sciences*, 274(1621):2019–2025, 2007.
- Nick Bloom, Stephen Bond, and John Van Reenen. Uncertainty and investment dynamics. *Review of Economic Studies*, 74(2):391–415, April 2007.
- Godfrey R Bourne, Felix Breden, and Teresa C Allen. Females prefer carotenoid colored males as mates in the pentamorphic livebearing fish, *poecilia parae*. *Naturwissenschaften*, 90(9):402–405, 2003.
- Adrian Bruhin, Helga Fehr-Duda, and Thomas Epper. Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412, 2010.
- Ricardo J. Caballero. Time dependent rules, aggregate stickiness and information externalities. Discussion papers, Columbia University, Department of Economics, 1989.
- Ricardo J. Caballero and Eduardo M.R.A. Engel. Price stickiness in ss models: New interpretations of old results. *Journal of Monetary Economics*, 54, 2007.
- Ryan Calsbeek, Suzanne H Alonzo, Kelly Zamudio, and Barry Sinervo. Sexual selection and

- alternative mating behaviours generate demographic stochasticity in small populations. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 269(1487):157–164, 2002.
- Guillermo A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 20(3):383–398, September 1983.
- Andrew Caplin and John Leahy. Economic theory and the world of practice: A celebration of the (s, s) model. *Journal of Economic Perspectives*, 24(1):183–202, Winter 2010.
- Andrew Caplin, Mark Dean, and Daniel Martin. Search and satisficing. *American Economic Review*, 101(7):2899–2922, December 2011.
- Andrew S. Caplin and Daniel F. Spulber. Menu costs and the neutrality of money. *The Quarterly Journal of Economics*, 102(4):703–725, November 1987.
- Dennis Chitty. *Do lemmings commit suicide?: beautiful hypotheses and ugly facts*. Oxford University Press, 1996.
- John Conlisk. Why bounded rationality? *Journal of economic literature*, pages 669–700, 1996.
- Ammon Corl, Alison R Davis, Shawn R Kuchta, and Barry Sinervo. Selective loss of polymorphic mating types is associated with rapid phenotypic evolution during morphic speciation. *Proceedings of the National Academy of Sciences*, 107(9):4254–4259, 2010.
- James Costain and Anton Nakov. Price adjustments in a general model of state-dependent pricing. *Journal of Money, Credit and Banking*, 43:385–406, March 2011a.

- James Costain and Anton Nakov. Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics*, 58(6):646–665, 2011b.
- James Costain and Anton Nakov. Logit price dynamics. Technical Report 1301, Banco de Espana, 2013.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Leif Danziger. A dynamic economy with costly price adjustments. *American Economic Review*, 89(4):878–901, September 1999.
- Douglas Davis and Oleg Korenok. Nominal shocks in monopolistically competitive markets: An experiment. *Journal of Monetary Economics*, 58(6):578–589, 2011.
- Avinash Dixit. *The Art of Smooth Pasting*. Harwood Academic Publishers, 1993.
- Michael Dotsey, Robert G. King, and Alexander L. Wolman. State-dependent pricing and the general equilibrium dynamics of money and output. *The Quarterly Journal of Economics*, 114(2):655–690, May 1999.
- John Duffy. Macroeconomics: A survey of laboratory research. Working Papers 334, University of Pittsburgh, Department of Economics, Feb 2008.
- 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.

- Andy Gardner, Stuart A West, and Angus Buckling. Bacteriocins, spite and virulence. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(1547):1529–1535, 2004.
- Andrea Gaunersdorfer and Josef Hofbauer. Fictitious play, shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11(2):279–303, 1995.
- 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.
- Mark Gertler and John Leahy. A phillips curve with an ss foundation. *Journal of Political Economy*, 116(3):533–572, June 2008.
- Mikhail Golosov and Robert E. Lucas. Menu costs and phillips curves. *Journal of Political Economy*, 115:171–199, 2007.
- 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.
- Sanford J Grossman and Guy Laroque. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica*, 58(1):25–51, January 1990.
- Samuel M Hartzmark and David H Solomon. Efficiency and the disposition effect in nfl prediction markets. *The Quarterly Journal of Finance*, 2(03), 2012.
- Peter J Hudson, David Newborn, and Andrew P Dobson. Regulation and stability of a

- free-living host-parasite system: *Trichostrongylus tenuis* in red grouse. i. monitoring and parasite reduction experiments. *Journal of Animal Ecology*, 61(3):477–486, 1992.
- Jef Huisman and Franz J Weissing. Biodiversity of plankton by species oscillations and chaos. *Nature*, 402(6760):407–410, 1999.
- Jef Huisman and Franz J Weissing. Biological conditions for oscillations and chaos generated by multispecies competition. *Ecology*, 82(10):2682–2695, 2001.
- Jonathan E Ingersoll and Lawrence J Jin. Realization utility with reference-dependent preferences. *Review of Financial Studies*, 26(3):723–767, 2013.
- Benjamin Kerr, Margaret A Riley, Marcus W Feldman, and Brendan JM Bohannan. Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature*, 418(6894):171–174, 2002.
- Peter J. Klenow and Oleksiy Kryvtsov. State-dependent or time-dependent pricing: Does it matter for recent u.s. inflation? *The Quarterly Journal of Economics*, 123(3):863–904, August 2008.
- Peter J. Klenow and Benjamin A. Malin. Microeconomic evidence on price-setting. In Benjamin M. Friedman and Michael Woodford, editors, *Handbook of Monetary Economics*, volume 3 of *Handbook of Monetary Economics*, chapter 6, pages 231–284. Elsevier, 2010.
- Charles J Krebs, Stan Boutin, Rudy Boonstra, ARE Sinclair, JNM Smith, Mark RT Dale, K Martin, and R Turkington. Impact of food and predation on the snowshoe hare cycle. *Science*, 269(5227):1112–1115, 1995.

- Lesley T Lancaster, Christy A Hipsley, and Barry Sinervo. Female choice for optimal combinations of multiple male display traits increases offspring survival. *Behavioral Ecology*, 20(5):993–999, 2009.
- Lars-Göran Mattsson and Jörgen W Weibull. Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*, 41(1):61–78, 2002.
- Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.
- Virgiliu Midrigan. Menu costs, multiproduct firms, and aggregate fluctuations. Working paper, New York University, 2008.
- Virgiliu Midrigan. Is firm pricing state or time dependent? evidence from u.s. manufacturing. *The Review of Economics and Statistics*, 92(3):643–656, August 2010.
- Virgiliu Midrigan. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, 79(4):1139–1180, July 2011.
- Emi Nakamura and Jn Steinsson. Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464, November 2008.
- Charles N. Noussair, Damjan Pfajfar, and Janos Zsiros. Pricing decisions in an experimental dynamic stochastic general equilibrium economy. Technical report, University of Tilburg, 2013.
- 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. *Review of Economic Studies*, 76(3):1103–1124, 07 2009a.
- Ryan Oprea, Daniel Friedman, and Steven T Anderson. Learning to wait: A laboratory investigation. *The Review of Economic Studies*, 76(3):1103–1124, 2009b.
- Dhanashree A Paranjpe, Elizabeth Bastiaans, Amy Patten, Robert D Cooper, and Barry Sinervo. Evidence of maternal effects on temperature preference in side-blotched lizards: implications for evolutionary response to climate change. *Ecology and evolution*, 3(7):1977–1991, 2013.
- David W Pfennig. Polyphenism in spadefoot toad tadpoles as a logically adjusted evolutionarily stable strategy. *Evolution*, pages 1408–1420, 1992a.
- DW Pfennig. Proximate and functional causes of polyphenism in an anuran tadpole. *Functional Ecology*, 6(2):167–174, 1992b.
- Robert S Pindyck. Irreversible investment, capacity choice, and the value of the firm. *American Economic Review*, 78(5):969–85, December 1988.
- Ricardo Reis. Inattentive producers. *Review of Economic Studies*, 73(3):793–821, 2006.
- William H Sandholm. *Population games and evolutionary dynamics*. MIT press, 2010.
- Herbert Scarf. Inventory theory. *Operations Research*, 50(1):186–191, 1959.

- 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.
- Eytan Sheshinski and Yoram Weiss. Inflation and the costs of price adjustment. *The Review of Economic Studies*, 44(2):287–303, June 1977.
- Christopher A. Sims. Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49(1):317–356, December 1998.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003.
- Christopher A. Sims. Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163, May 2006.
- Barry Sinervo. Runaway social games, genetic cycles driven by alternative male and female strategies, and the origin of morphs. *Genetica*, 112(1):417–434, 2001.
- Barry Sinervo and Ryan Calsbeek. Physiological epistasis, ontogenetic conflict and natural selection on physiology and life history. *Integrative and Comparative Biology*, 43(3):419–430, 2003.
- Barry Sinervo and Ryan Calsbeek. The developmental, physiological, neural, and genetical causes and consequences of frequency-dependent selection in the wild. *Annu. Rev. Ecol. Evol. Syst.*, 37:581–610, 2006.
- Barry Sinervo and Curt M Lively. The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380(6571):240–243, 1996.

- Barry Sinervo, Erik Svensson, and Tosha Comendant. Density cycles and an offspring quantity and quality game driven by natural selection. *Nature*, 406(6799):985–988, 2000.
- Barry Sinervo, Colin Bleay, and Chloe Adamopoulou. Social causes of correlational selection and the resolution of a heritable throat color polymorphism in a lizard. *Evolution*, 55(10):2040–2052, 2001.
- Barry Sinervo, Alexis Chaine, Jean Clobert, Ryan Calsbeek, Lisa Hazard, Lesley Lancaster, Andrew G McAdam, Suzanne Alonzo, Gwynne Corrigan, and Michael E Hochberg. Self-recognition, color signals, and cycles of greenbeard mutualism and altruism. *Proceedings of the National Academy of Sciences*, 103(19):7372–7377, 2006.
- Barry Sinervo, Fausto Mendez-De-La-Cruz, Donald B Miles, Benoit Heulin, Elizabeth Bastiaans, Maricela Villagrán-Santa Cruz, Rafael Lara-Resendiz, Norberto Martínez-Méndez, Martha Lucía Calderón-Espinosa, Rubi Nelsi Meza-Lázaro, et al. Erosion of lizard diversity by climate change and altered thermal niches. *Science*, 328(5980):894–899, 2010.
- John Maynard Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982.
- Dale O Stahl. Entropy control costs and entropic equilibria. *International Journal of Game Theory*, 19(2):129–138, 1990.
- Nancy L. Stokey. *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press, 2008.
- John B. Taylor. Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):109–132, May 1980.

- Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.
- Joseph Vavra. Inflation dynamics and time-varying uncertainty: New evidence and an ss interpretation. Technical report, University of Chicago, August 2013.
- Annette Vissing-Jorgensen. Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(3):825–853, 2002.
- 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.
- Bart J. Wilson. Menu costs and nominal price friction: An experimental examination. *Journal of Economic Behavior & Organization*, 35(3):371–388, April 1998.
- Michael Woodford. Information-constrained state-dependent pricing. NBER Working Papers 14620, National Bureau of Economic Research, Inc., December 2008.
- George Wu and Richard Gonzalez. Curvature of the probability weighting function. *Management science*, 42(12):1676–1690, 1996.
- Kelly R Zamudio and Barry Sinervo. Polygyny, mate-guarding, and posthumous fertilization as alternative male mating strategies. *Proceedings of the National Academy of Sciences*, 97(26):14427–14432, 2000.