UNIVERSITY OF CALIFORNIA,
IRVINE

Learning and Asset Pricing

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Michael Shin

Dissertation Committee:
Professor William Branch, Chair
Professor John Duffy
Professor Eric Swanson
Professor Cars Hommes

2019

DEDICATION

To my parents.

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF TABLES

ACKNOWLEDGMENTS

This dissertation would not have been possible without the support of some very important people. First and foremost I want to thank my advisor Bill Branch. Without his patience, generosity, and constant mentorship, I would not have been able to complete this nor be the researcher I am today. No amount of words can express my gratitude. I next want to thank John Duffy for his continuing support and guidance throughout my graduate career. I thank Eric Swanson for his valuable feedback both during my advancement and as a dissertation committee member. Next I thank Cars Hommes for his hospitality during my visit to the University of Amsterdam and his advice throughout the different stages of my job market paper. I also thank Chong Huang, Fabio Milani, and Guillaume Rocheteau. I thank Pat Testa for his friendship throughout one of the most important part of our lives. Finally I thank my parents for their constant financial and moral support throughout the brightest and darkest times during my graduate career.

CURRICULUM VITAE

Michael Shin

EDUCATION

Doctor of Philosophy in Economics **2019**
University of California, Irvine *Irvine, CA*

Master of Arts in Economics **2015**
University of California, Irvine *Irvine, CA*

Bachelor of Arts in Business Economics **2014**
University of California, Riverside *Riverside, CA*

TEACHING EXPERIENCE

Teaching Assistant **2014–2019**
University of California, Irvine *Irvine, CA*

ABSTRACT OF THE DISSERTATION

Learning and Asset Pricing

By

Michael Shin

Doctor of Philosophy in Economics

University of California, Irvine, 2019

Professor William Branch, Chair

The rational expectations (RE) hypothesis although elegant and useful requires demanding assumptions on part of the agent. A key outcome of the RE hypothesis is that beliefs disappear as an independent force in the model. A branch of the literature focuses on relaxations of the RE hypothesis to allow agents to instead learn the data-generating process (DGP) over time. With adaptive learning, beliefs re-emerge as a key element that influences the DGP. I argue that this relaxation is important for asset pricing settings, which are complex environments where individual beliefs play a key role in decision making. My dissertation will explore instances where learning can improve our understanding of asset pricing.

Chapter 1 presents a simple asset pricing model with endogenous participation that can match key volatility moments when agents adaptively learn about both the risk and the return of stocks. With learning about risk, excess volatility of prices is driven by fluctuations in the participation rate that arise because agents' risk estimates vary with prices. I find that learning about risk is quantitatively more important than learning about returns. A calibrated model can jointly match the mean participation rate, the volatility of participation rates, and explain 25% of the excess volatility of stock prices observed in U.S. data.

Chapter 2 presents a simplified version of the model in Chapter 1 and tests the model in a laboratory setting. Recent evidence suggests subjective returns play a key role in stock market participation. Furthermore, there is strong evidence that stock market experiences,

i.e. realized returns, impact subjective returns. I bring a model into the laboratory and find that learning-driven subjective returns can explain limited participation. Stock market participation is increasing in both subjective returns and past realized returns. I find direct evidence that "learning from experience" generates heterogeneity in subjective returns, where subjects who experience low returns have lower subjective returns than subjects who experience high returns. In particular, subjects over-weigh price trends when they experience high returns and under-weigh it when they experience low returns.

Chapter 3 presents an asset pricing model where agents test the specification of their models, while adaptively updating the parameters and find that restricted-perceptions equilibria (RPE) naturally arise. I extend upon recently developed model specification techniques to a multi-agent framework. Multiple agents are endowed with different models which they update and test the specification in real time. When a model is rejected, agents draw a new model from a distribution. I find that the rational expectations equilibrium (REE) is not locally stable with respect to hypothesis testing under reasonable parameterization. With constant-gain learning, the model spends most of its time in a subset of the RPE and in particular, the dominant model used is not the fully-specified model, but a misspecified one.

Chapter 1

Endogenous Participation, Risk, and Learning in the Stock Market

1.1 Introduction

In a seminal paper, Mankiw and Zeldes (1991) report that in 1984 only 27.6% of households in the PSID participated in the stock market. This "participation puzzle" is at odds with standard assumptions in asset pricing models. Subsequent studies demonstrate that limited participation is robust across time periods, asset classes, direct/indirect holdings, and countries (Bertaut and Starr-McCluer 2002, Guiso and Jappelli 2002, Campbell 2006).

This paper focuses on the dynamic relationship between participation and asset prices. Table 1.1, taken from the Survey of Income and Program Participation (SIPP), documents fluctuations in participation rates over time with a low of 19.6% and a high of 29.4%. More recently, Arrondel et al (2014) provide structural econometric evidence of a causal relationship between expected returns and participation rates. I propose a theory of endogenous fluctuations in participation rates and demonstrate that it can be an important driver of stock price volatility.

Year	Participation Rate (%)
1995	20.7
1998	27.1
2000	27.1
2002	29.4
2004	26.4
2005	25.1
2009	21.8
2010	20.4
2011	19.6
2013	20.0

Table 1.1: Stock Market Participation Rates from 1995 - 2013. Data was extracted from the Survey of Income Program and Participation.

I present a mean-variance asset pricing model with two key departures: a costly participation margin and imperfect knowledge about the stochastic processes driving prices. Participation is costly and agents choose to participate in the stock market by balancing entry costs against the risk-adjusted expected return from participating. I relax the rational expectations (RE) assumption and instead assume that agents behave like good econometricians who formulate and estimate a well-specified forecasting model for future stock prices. A key assumption is that agents have to also estimate the risk, i.e. the conditional variance of returns. Learning about the risk and return provides two different feedback mechanisms that contribute to price fluctuations with learning about risk being quantitatively more important. I find that with learning, changes in agents' risk estimates lead to large fluctuations in the participation rate which in turn lead to large fluctuations in the price.

To introduce endogenous fluctuations in participation, I implement a cost function which captures features beyond fixed participation costs while keeping the model tractable. This approach is motivated by recent empirical evidence revealing costs to participate in the stock market that go beyond fixed entry costs such as financial awareness, financial literacy, and other cognition costs (Guiso and Sodini 2013). Similar to labor-leisure decisions, I model participation as the result of costly effort. Individuals who exert more effort are more likely to enter the stock market.

It is well known that asset pricing models with RE have difficulty generating excess volatility (Timmermann 1993). RE requires subjective beliefs to align with the objective measured probability distribution that is implied by those beliefs. Therefore with RE, beliefs disappear as an independent force driving prices, volatility, and participation. I argue that belief-driven learning dynamics are key in explaining the interplay between participation and stock price volatility. Hence I take a step down from RE and implement an adaptive learning rule.

I first characterize the steady-state equilibrium and do comparative statics which give insights on participation without learning. I find that limited participation lowers the steady-state price because fewer agents participating in the market corresponds to lower market demand for the asset. In the steady-state, changes in the structural parameters shift both the asset demand and participation decision. Therefore, the participation decision can either shift in the same direction as the asset demand, amplifying the effect on prices, or in the opposite direction and reduce the effect. For instance, a decrease in the risk-free rate increases the demand for the risky asset which increases the price but also increases the participation rate which leads to a further increase in the price.

I then study the learning dynamics while keeping risk constant in order to characterize the learning about returns channel. Along a temporary equilibrium path, agents exert effort to participate in the stock market, where the level of effort depends on return expectations. Additionally, participation has a direct effect on asset prices and returns. When prices increase, expected returns decrease, leading to a decrease in participation which in turn decreases prices. This feedback loop due to learning about returns is an important mechanism in our model for explaining limited participation and excess volatility of stock prices.

The role for learning about risk is motivated by survey responses in Arrondel et al (2014) who find that 20.7% of nonparticipants did not invest in the stock market due to the perceived riskiness of stocks. Since risk influences participation, risk itself is an equilibrium object jointly determined along with prices and returns. I follow the approach in Branch and Evans (2011) by explicitly calculating the conditional variance of returns. Risk affects participation because higher risk lowers returns in certain states and hence lowers the expected utility from participation. An increase in the subjective risk leads to a decrease in the participation rate which leads to a decrease in the price. Furthermore a decrease in prices increases realized returns which leads to an increase in the subjective risk which further decreases the participation rate. This process continues until risk estimates are adjusted and the mechanism moves in the opposite direction. This feedback mechanism due to learning about risk, is key to generating more volatility in prices than the model with exogenous risk.

I also find that learning about risk is quantitatively more important than learning about returns. Learning about risk generates larger volatility in participation rates which directly contributes to larger volatility in prices. Essentially, learning about risk is more important for volatility because changes in risk have a persistent impact on prices. There is a self-fulfilling aspect between prices and risk which is amplified by the participation margin. As agents learn about the risk and subjective risk increases, participation decreases and prices decrease as well. In this sense, higher risk leads to persistently lower prices leading to higher price volatility. In contrast, prices and expected returns have a negative relationship such that higher prices lead to lower expected returns which lowers participation. Hence with learning about returns, higher prices are offset by lower expected returns leading to lower persistence in volatility. A quantitative exercise demonstrates that the model with learning about risk can match the mean participation rate, the volatility of participation rates, and generate 25% of the excess volatility in stock prices.

1.1.1 Literature Review

This paper contributes primarily to two literatures. First, to the literature on limited participation and household finance. There is a large literature on exogenous limited participation such as Guo (2004), Guvenen (2009), and Lansing (2015). The first paper to endogenize limited participation is Allen and Gale (1994) who implement fixed costs in a one-shot asset pricing game. They find that endogenous participation can increase the volatility of asset prices. This paper is most similar to Orosel (1998), who models endogenous participation in an overlapping generations model with fixed costs. My model differs from theirs by implementing a variable cost function, which allows us to tractably analyze the dynamics of the model while also mapping participation rates to the data. Gomes and Michaelides (2005) and Fagereng et al (2017) implement fixed costs in a life-cycle model and calibrate it. Models in this strand of the literature focus on matching the cross-section of asset holdings. In contrast, I focus on aggregate participation and how it jointly impacts asset prices and expected returns in the time-series.

Second, I contribute to the literature on learning. This paper follows a strand of literature put forth by Marcet and Sargent (1989) and Evans and Honkapohja (2001) which relaxes the RE hypothesis and replaces it with an econometric learning rule. The first paper to analyze learning in an asset pricing model is Timmermann (1993) who shows that adaptive learning can generate excess volatility. Our environment is similar to Branch and Evans (2011) who calibrate a mean-variance asset pricing model where agents also learn about the risk. We differ from their approach by adding a participation decision and focus on price volatility rather than asset bubbles. More recently, Nakov and Nuño (2015) calibrate an asset pricing model with learning and Blanchard-Yaari households. Finally Adam et al (2016) formally test a consumption asset pricing model with learning. As far as I know this is the first paper to combine an asset pricing model with endogenous participation and learning.

1.2 Model

Time is measured in discrete periods $t = 1, 2, ...$ and there are overlapping generations of agents who live for 2 periods. All agents have CARA utility functions of the form: $u(c) = -e^{-\rho c}$, where $\rho > 0$ is the coefficient of absolute risk aversion. There is one non-storable consumption good which is taken as the numeraire. There are two assets traded in perfectly competitive markets: a risky Lucas tree and a riskless one-period bond. Like Lucas (1978), shares underlie firms that produce exogenous stochastic output of the consumption good. Participation in the risky market requires effort and none is required in the riskless market. The riskless one-period bond as an analogue to a savings account or a storage technology. In reality, participation in the bond market also requires effort but the cost is presumably lower. I assume that the riskless asset gives an exogenous gross return $R = 1 + r > 1$ of the consumption good and the supply is infinitely elastic.

The initial old are endowed with $S > 0$ shares, where each share pays at the beginning of the period a dividend D_t. D_t follows an exogenous process:

$$D_t = \mu + \epsilon_t^D$$

where $\mu > 0$ and ϵ_t^D is white noise with distribution $N(0, \sigma_D^2)$. The dividend process is simplistic for technical convenience and to clearly focus on the participation channel.[1] After

[1] In order to focus on the interactions between learning and participation, I abstract from seriously modeling dividends and asset supply, both of which are better approximated by persistent or non-stationary processes.

the initial old is endowed with the shares, subsequent S follow an exogenous process:

$$S_t = S + \epsilon_t^S \tag{1.1}$$

where ϵ_t^S is white noise with distribution $N(0, \sigma_S^2)$. The stochastic supply is a proxy for volatility in asset float where firms create new issues and provide options that are periodically exercised changing the available supply at a given time. Furthermore, the impact of asset float is well documented in the literature (Baker and Wurgler 2000).

I follow Branch and Evans (2011) who show that in a similar model, stochastic variation in the population of young agents can produce shocks in per capita asset supply. At the beginning of each period, a new generation n_t enters the economy, where n_t is an iid random process with an inverse mean of one. Because n_t is random, the per capita asset supply S_t is also random, and follows the stochastic process in Equation (1.1). Each agent lives for two periods, has initial endowment w normalized to 1, and consumes only in the second period. This is to abstract away from savings decisions in order to focus entirely on the lifetime portfolio choice and the stock market entry decision of the young households.

There are costs to participate in the stock market beyond fixed entry costs such as investing in financial literacy, financial awareness, and other cognition costs (Guiso and Sodini 2013). I implement a cost function that captures these features while also keeping the model tractable. Agents can exert up to one unit of effort e. Similar to labor-leisure decisions, exerting effort is assumed to be costly in terms of utility. Agents face a variable cost function $\Phi(e)$ that is increasing in their effort at a decreasing rate with $\Phi(0) = 0$, and $\Phi'(0) = 0$.

An iid random variable χ which takes on values 0 and 1 determines the young's ability to participate. When $\chi = 1$, the young can participate in the stock market, else they are unable to enter. Furthermore, the young can influence the likelihood of χ by exerting effort. If the agent exerts $e = 1$, then he enters the market with certainty. Similarly, if the agent exerts $e = \frac{1}{2}$, then he enters the market with probability $\frac{1}{2}$. Implicitly, agents who exert more effort are more likely to increase their financial awareness or invest in financial literacy and hence are more likely to enter the stock market.

My modeling approach is similar to employment lotteries in labor models following a technique pioneered by Rogerson (1988). Since entering the stock market is an indivisible choice, households can improve their welfare by drawing lotteries amongst themselves and enter the market probabilistically. A natural interpretation, following Ljungqvist and Sargent (2011), is that this formulation is equivalent to choosing a portion of your lifetime in which to enter the stock market. Hence e can alternatively be interpreted as the fraction of an agent's life in which they would like to participate in the stock market.[2] Because of the Law of Large Numbers, e also corresponds to the aggregate participation rate.

1.3 Equilibrium

1.3.1 Portfolio Choice

Consumption depends on whether the household is a stock market participant. Hence $c_t = c_{\chi t}$, where $c_{\chi t}$ is state-contingent consumption. Let c_{0t} be risk-free consumption and c_{1t} be risky consumption. Then agents maximize the following program:

$$\underset{x_t(\chi), e_t}{\text{maximize}} \quad (1 - e_t)u(c_{0t}) + e_t E_t u(c_{1t}) - \Phi(e_t)$$

$$\text{subject to} \quad c_{\chi t} = \begin{cases} R + x_t(p_{t+1} + D_{t+1} - Rp_t) & \text{if } \chi = 1 \\ R & \text{if } \chi = 0 \end{cases} \tag{1.2}$$

[2]Their exact interpretation is in terms of the labor market in which agents choose their career lengths. Alternatively, one can imagine agents having a distribution of fixed entry costs and the representative agent being a stand-in for the heterogeneity. This interpretation is similar in spirit to Orosel (1998).

where x_t is the asset holding decision and p_t is the price of the risky asset.

I make a timing assumption on the portfolio and participation decision. In particular, I assume that the decisions are made sequentially, that is, agents make the participation decision before the portfolio decision. Equation (1.2) is the agent's budget constraint. Agents allocate their endowment between the risky asset and the one-period bond. Agents choose some portfolio x_t and effort level e_t to maximize their lifetime utility. Furthermore, agents also assume that the payoffs, $p_{t+1} + D_{t+1}$, are normally distributed, which implies c_{1t} is also normally distributed. Since the utility is CARA, I arrive at the following first-order conditions:

$$x_t = \frac{E_t(p_{t+1} + D_{t+1}) - Rp_t}{\rho\sigma_p^2} \tag{1.3}$$

$$\Phi'(e_t) = \max\{E_t u(c_{1t}) - u(c_{0t}),\ 0\} \tag{1.4}$$

where $\sigma_p^2 \equiv Var_t(p_{t+1}+D_{t+1})$ is the conditional variance of returns, i.e. the agents' subjective measure of risk. For now σ_p^2 is treated as a constant but will be made endogenous in subsequent sections. The inverse function is:

$$e_t = \min\{\Phi'^{-1}[E_t u(c_{1t}) - u(c_{0t})],\ 1\} \tag{1.5}$$

Equation (1.3) is the standard mean-variance asset demand function which is downward sloping in the price and Equation (1.5) is the participation decision. Hence, the agent's optimal effort level depends on equating the expected utility difference of entering and not entering with the marginal cost of entry.

1.3.2 Steady-State

To better understand the participation decision, it is illustrative to analyze the steady-state. I assume a particular form for the cost function:

$$\Phi(e_t) = \frac{1}{2A} e_t^2, \text{ where } A > 0 \text{ is some technology or efficiency parameter.}$$

Then the inverse of the derivative is:

$$e_t = \Phi'^{-1}(y) = Ay, \text{ where } y \geq 0 \text{ is some input.}$$

Taking the first-order condition I now get:

$$e_t = \min\{A\Gamma(E_t p_t, E_t p_{t+1}),\ 1\} \tag{1.6}$$

where:

$$\Gamma(E_t p_t, E_t p_{t+1}) = e^{-\rho R} - e^{-\rho R - \frac{[E_t(p_{t+1}+D_{t+1}) - RE_t p_t]^2}{2\sigma_p^2}} \tag{1.7}$$

Equation (1.7) follows from the transformation of an exponential function with respect to normal random variables and is derived in Appendix A. $\Gamma(p)$ is the expected utility difference between the two states which can be interpreted as the expected excess utility return of entering the stock market. Since the participation decision is made prior to the portfolio decision, there is an E_t on the price p_t. Thus agents care about the expected price $E_t p_t$ when participating. The market-clearing condition is as follows:

$$e_t x_t = S_t$$

Once I impose market-clearing, I get the following pricing equation:

$$p_t = R^{-1}\left[E_t(p_{t+1} + D_{t+1}) - \frac{S_t}{e_t}\rho\sigma_p^2\right] \tag{1.8}$$

This is the same as the standard mean-variance pricing equation except now the price also depends on e_t, where again, e_t is the participation rate. When $e_t = 1$, the model collapses to the standard mean-variance case. Otherwise, when $e_t < 1$, the limited participation steady-state price will be lower than the corresponding full participation price. Since market-clearing implies prices must be positive, the participation rate e_t will always be positive in equilibrium and hence Equation (1.8) is well-defined.

There are two propagation mechanisms with the addition of the participation decision. The first mechanism is through $E_t p_{t+1}$. In the standard model, $E_t p_{t+1}$ affects the price directly, but in our model it also impacts it indirectly through e_t since the participation decision now depends on expected prices. Second, as e_t increases, p_t increases. In particular, lower participation rates lead to lower prices and higher participation rates lead to higher prices. This means that increases in expected prices further increase the price through the participation channel. Thus I can view e_t as an amplification mechanism, where changes in participation rates are demand shocks.

These two effects interact nonlinearly. In order to build intuition about the participation channel, I look at the steady-state equilibrium. I find the participation channel can act as both an amplification and dampening mechanism. For instance, an increase in R decreases both the price through the asset demand and through the participation channel. In contrast, an increase in the risk σ_p^2 decreases the price through the asset demand but increases it through the participation channel.

I characterize the steady-state equilibrium where $S_t = S$ and $p_{t+1} = p_t = \bar{p}$. Once I

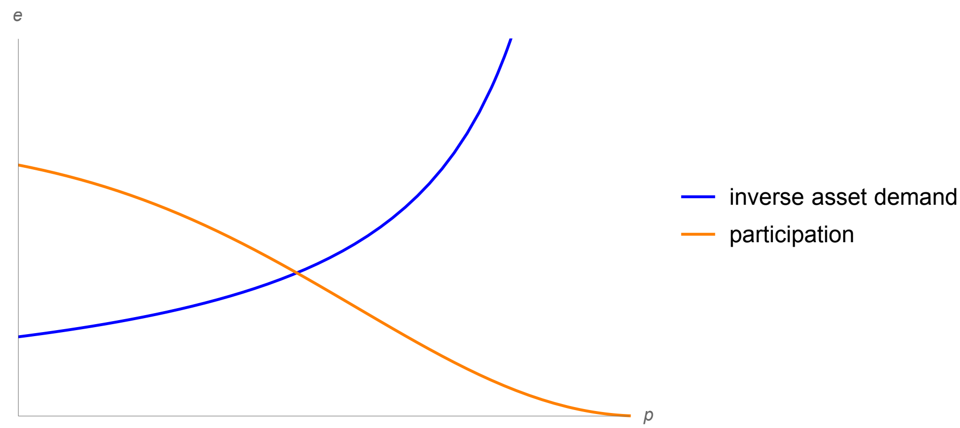

Figure 1.1: Steady-state Equilibrium.

solve for the steady-state I get the following form:

$$\bar{x} = \frac{\mu - (R-1)\bar{p}}{\rho\sigma_p^2}$$

Plugging in for the cost function I get the following participation equation:

$$\bar{e} = \min\{A\Gamma(\bar{p}),\ 1\}$$

where:

$$\Gamma(\bar{p}) = e^{-\rho R} - e^{-\rho R - \frac{[\mu-(R-1)\bar{p}]^2}{2\sigma_p^2}}$$

Proposition 1. *There exists a unique steady-state equilibrium.*

Proofs are provided in Appendix A. Figure 1.1 depicts Proposition 1 graphically for a set of parameters. Given that the steady-state equilibrium exists and is unique, I derive the expression for the steady-state price. I also compare it to the standard mean-variance case. The steady-state equation for the price in the standard mean-variance model is as follows:

$$\bar{p} = \frac{\mu - S\rho\sigma_p^2}{R-1} \tag{1.9}$$

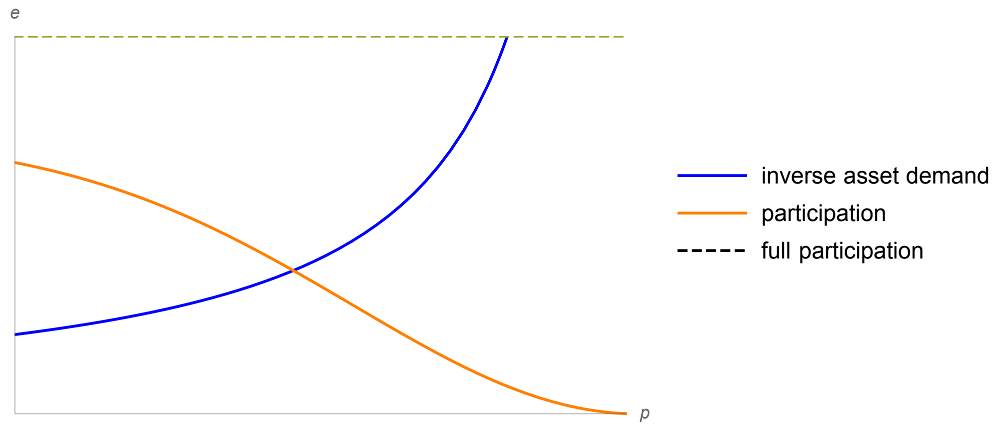

Figure 1.2: Steady-state Equilibrium: Limited and Full Participation.

The steady-state equation for my model is:

$$\bar{p} = \frac{\mu - \frac{S}{\bar{e}}\rho\sigma_p^2}{R - 1} \tag{1.10}$$

where:

$$\bar{e} = \min\{A\Gamma(\bar{p}),\ 1\}$$

When $\bar{e} = 1$, the model again collapses to the full participation case. Since $\bar{e}$ is decreasing in $\bar{p}$, our steady-state price will be lower than the benchmark. I graph Equations (1.9) and (1.10) in Figure 1.2 to describe the relationship between the two models.

In Figure 1.2, we see that the full participation model has a higher steady-state price than with limited participation. Another thing to note is that changes in the structural parameters shift both functions so the magnitude of the change is different than the benchmark. Moreover, I can plug Equation (1.10) into the steady-state participation function to find $\bar{e}$ as an implicit function of the fundamentals:

$$\bar{e} = \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{S^2\rho^2\sigma_p^2}{2\bar{e}^2}},\ 1\}$$

I now sign the derivatives for the steady-state participation and pricing functions.

Proposition 2. *For $\bar{e} < 1$, the derivative signs for steady-state participation are as follows:*

$\frac{\partial \bar{e}}{\partial R} < 0$, $\frac{\partial \bar{e}}{\partial \mu} = 0$, $\frac{\partial \bar{e}}{\partial A} > 0$, $\frac{\partial \bar{e}}{\partial \sigma_p^2} > 0$, and $\frac{\partial \bar{e}}{\partial \rho}$ is indeterminate.

Proposition 3. *The derivative signs for steady-state price are as follows: $\frac{\partial \bar{p}}{\partial R} < 0$, $\frac{\partial \bar{p}}{\partial \mu} > 0$, $\frac{\partial \bar{p}}{\partial A} > 0$, $\frac{\partial \bar{p}}{\partial \sigma_p^2} < 0$, and $\frac{\partial \bar{p}}{\partial \rho}$ is indeterminate.*

With endogenous participation, the participation and asset demand functions need not move in the same direction. For instance, when dividends μ increase, prices increase because agents increase their asset demand but the steady-state participation rate is unchanged. In Equation (1.10), we see before the substitution that steady-state participation is a function of μ. Nevertheless the increase in μ increases participation but this effect is exactly offset by the increase in prices. When the interest rate R increases, agents lower their asset holdings which decreases the price. They also decrease participation since the risk-free rate now gives a higher return which decreases their expected utility gain from investing, further decreasing the price. Next, an increase in the cost parameter A lowers the cost of participating, which increases participation and increases the price.

Furthermore, when the risk σ_p^2 increases, agents lower their asset holdings which lowers the price but their participation rate increases. Similar to the change in μ, there are counter-balancing effects and the intuition is as follows. For the individual agent, participation is decreasing in σ_p^2 because it decreases their expected utility gain from investing. Participation is also increasing as steady-state price goes down. In equilibrium, the price effect dominates and steady-state participation is increasing in σ_p^2. Finally, when agents become more risk averse, they decrease their asset holdings and price decreases. The participation decision now has a u-shaped relationship with respect to ρ. Participation is increasing in ρ up to some threshold value, and then decreasing afterwards. This threshold depends on the risk-free rate being sufficiently high. If the risk-free rate is high enough, then participation is increasing in ρ. This is because in the steady-state, participation is decreasing in prices because higher prices lower returns. Hence, a change in the price due to the asset demand can be partially dampened by the participation effect, but the change in prices is indeterminate.

I now elaborate on the intuition behind the risk σ_p^2 comparative statics since it plays a key role in our model. In the steady-state, an increase in σ_p^2 makes the asset riskier to hold, but prices become low enough such that the equilibrium level of participation will be higher. Out of steady-state, the price effect only dominates when $E_t p_{t+1}$ approaches $\bar{p}$. With learning, the effect of an increase in σ_p^2 will decrease participation which will be the main driver of volatility in prices. Hence to understand the dynamic relationship between risk and participation, it is important to analyze the learning dynamics.

1.4 Asset Pricing Dynamics with Learning

Because the stochastic model is a complicated non-linear rational expectations equation, it is not possible to characterize the full set of rational expectations equilibria (REE). However, since the unique steady-state is locally determinate, I am able to solve for one type of REE, the noisy steady-state REE. The noisy steady-state REE is a non-linear REE where the equilibrium path is a sequence of noisy deviations around the steady-state. I characterize the noisy steady-state REE and then analyze its stability under learning. I do this by first taking the risk σ_p^2 as exogenous to clearly understand the dynamic properties of the participation decision. I then analyze the numerical properties when σ_p^2 is endogenous.

1.4.1 Rational Expectations Equilibrium

I start by characterizing the noisy steady-state REE with exogenous risk. The key equation in the model is the following expectational difference equation:

$$p_t = R^{-1}\left[E_t p_{t+1} + \mu - \frac{S_t}{e_t}\rho\sigma_p^2\right] \tag{1.11}$$

Definition 1. *A noisy steady-state REE is a sequence* $\{p_t\}_{t=0}^{\infty}$ *and* $\{\epsilon_t^S\}_{t=0}^{\infty}$ *such that the* *sequences solve the equation:*

$$p(\epsilon_t^S) = R^{-1}\left[E_t p(\epsilon_{t+1}^S) + \mu - \frac{S}{e(E_t p(\epsilon_{t+1}^S), p(\epsilon_t^S))}\rho\sigma_p^2 - \frac{\epsilon_t^S}{e(E_t p(\epsilon_{t+1}^S), p(\epsilon_t^S))}\rho\sigma_p^2\right] \tag{1.12}$$

where:

$$e = \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{\left[E_t p(\epsilon_{t+1}^S) + \mu - R E_t p(\epsilon_t^S)\right]^2}{2\sigma_p^2}},\ 1\}$$

Hence the noisy steady-state is characterized by a function $p(\epsilon_t^S)$ that solves Equation (1.12). Since $E_t p(\epsilon_{t+1}^S) = E_t p(\epsilon_t^S)$, given that ϵ_t^S is iid:

$$e = \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{\left[\mu - (R-1)E_t p(\epsilon_{t+1}^S)\right]^2}{2\sigma_p^2}},\ 1\}$$

Then:

$$p(\epsilon_t^S) = R^{-1}\left[E_t p(\epsilon_{t+1}^S) + \mu - \frac{S}{e(\epsilon_t^S)}\rho\sigma_p^2 - \frac{\epsilon_t^S}{e(\epsilon_t^S)}\right]$$

Since ϵ_t^S is white noise, $E_t p$ is a constant and coincides with the nonstochastic steady-state $\bar{p}$. Then e becomes:

$$e = \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{[\mu - (R-1)\bar{p}]^2}{2\sigma_p^2}},\ 1\}$$

which is just $e = \bar{e}$. Then

$$p_t = \bar{p} + \eta_t \tag{1.13}$$

where $\eta_t \equiv -R^{-1}\left[\frac{\epsilon_t^S}{\bar{e}}\rho\sigma_p^2\right]$. I use a proposition by Evans and Honkapohja (1995) that proves that the noisy steady-state REE exists and is unique.

Proposition 4. *If the sequence of shocks $\{\epsilon_t^S\}_{t=0}^{\infty}$ are such that $|\epsilon_t^S| < \alpha$ with probability 1 for all t and $\alpha > 0$ is sufficiently small, then there exists a unique noisy steady-state REE.*

Proposition 4 states that Equation (1.13) is the unique noisy steady-state REE solution to Equation (1.12). Here α characterizes the support of the distribution of shocks. Essentially, the idea of a noisy steady-state REE is that when shocks are iid with compact support, there exists a stochastic equilibrium in a neighborhood around the steady-state. Hence I have fully characterized the noisy steady-state REE of our model and now I implement learning.

In practice, the α parameter which characterizes the support of the distribution is difficult to pin down. Although Proposition 4 states that the distribution exists, it provides no analytical solution for α. Thus, when doing my numerical simulation I use empirical moments and robustness checks to insure that the system is locally stable.

1.4.2 Endogenous Risk

I have treated the risk σ_p^2 as a constant. Importantly, σ_p^2 is an equilibrium object and having the agents learn about the risk has important implications. In asset markets with agents who learn over time, risk plays an important role because the perceived riskiness of an asset can lead to a lower asset demand that leads to lower prices in future periods. I argue that endogenizing σ_p^2 is crucial for understanding asset markets because we otherwise omit an important feedback mechanism that influences prices and expectations.

I now endogenize $\sigma_p^2 \equiv Var_t(p_{t+1} + D_{t+1})$. Then:

$$\sigma_p^2 = E_t(p_{t+1} - E_t p_{t+1} + D_{t+1} - \mu)^2$$

Solving out and plugging in Equation (1.13) I have:

$$\begin{aligned}
\sigma_p^2 &= E_t(-R^{-1}\rho\sigma_p^2 \frac{\epsilon_{t+1}^S}{\bar{e}} + \epsilon_{t+1}^D)^2 \\
&= Var_t(-R^{-1}\rho\sigma_p^2 \frac{\epsilon_{t+1}^S}{\bar{e}} + \epsilon_{t+1}^D) \\
&= \frac{R^{-2}\rho^2(\sigma_p^2)^2}{\bar{e}^2}\sigma_S^2 + \sigma_D^2
\end{aligned}$$

Solving for equilibrium risk σ_p^2 leads to:

$$\sigma_p^2 = \frac{\bar{e}^2 \pm \bar{e}\sqrt{\bar{e}^2 - 4R^{-2}\rho^2\sigma_S^2\sigma_D^2}}{2R^{-2}\rho^2\sigma_S^2} \tag{1.14}$$

Equation (1.14) is identical to Branch and Evans (2011) when $\bar{e} = 1$. We see now σ_p^2 is determined by fundamentals. Importantly, the standard deviation of supply, σ_S^2 now influences the risk since agents consider the effect of the volatility of shares on the volatility of returns. There are also two solutions to Equation (1.14) which correspond to low and high risk steady-states. Branch and Evans (2011) show that the low risk steady-state is unstable under learning. I find a similar result with our numerical analysis and hence focus on the low risk steady-state as well.

Moreover we see that both $\bar{e}$ and σ_p^2 are determined jointly in equilibrium. Unfortunately, because $\bar{e}$ and σ_p^2 have no closed form, I am unable to provide analytical solutions for the case with endogenous risk. Instead, I rely on numerical analysis under learning.

1.4.3 Adaptive Learning

Rational expectations (RE) requires a full understanding of the model as well as beliefs of other agents. In this sense it is a Nash equilibrium, such that coordination between agents requires strong cognitive and informational assumptions. Instead, many applied economists estimate econometric forecasting models and adjust the coefficients in light of new data. Here I adhere to the Cognitive Consistency Principle (Sargent 1993) which requires agents and econometricians to be on equal footing. In this regard I want to understand how an agent's learning mechanism will, in turn, affect the other endogenous variables. With adaptive learning, agents know the form of the REE but not the true parameters. I make a small deviation from RE where agents implement a learning rule and run least-squares regressions on the perceived pricing function.[3]

The REE of the model is a constant plus a noise. Then the agents are regressing prices on a constant and they need to keep track of the regression coefficient each period. I can rewrite the sample average recursively where expectations formation take the following form:

$$p^e_{t+1} = p^e_t + t^{-1}[p_{t-1} - p^e_t]$$

where p^e_t is the subjective expectation of prices formed at time t. This type of learning is called decreasing-gain learning. With decreasing-gain learning, agents estimate the sample average of prices and adjust their expectations as new data becomes available. If agents believe they are in a noisy steady-state and that the REE is a constant then $p_t = a_{t-1} + \nu_t$ where ν_t is the perceived white noise and a_t is updated recursively. Then, evidently $p^e_t = a_{t-1} = p^e_{t+1}$. $\Gamma(p^e_{t+1}, p^e_t)$ then becomes:

$$\Gamma(p^e_{t+1}) = e^{-\rho R} - e^{-\rho R - \frac{[(1-R)p^e_{t+1}+\mu]^2}{2\sigma^2_p}}$$

We say that the REE is locally stable if the model converges to the REE under decreasing-gain learning. I check the properties of the model with decreasing-gain learning and show that the REE is in fact stable under learning.

1.4.4 Stability Under Learning

I show analytically that the REE solution is locally stable under learning. To do this I have to analyze the mapping between the perceived law of motion (PLM) and actual law of motion (ALM). With econometric learning, agents know the form of the REE but not the parameters and hence the PLM is the equation that agents believe generate the observed

[3] I do not assume that agents learn about the dividend process since learning about exogenous processes provides no feedback. This assumption has no impact on the main results.

data. The ALM is the true data-generating process given the beliefs of the agents. Local stability analysis then amounts to understanding the functional relationship between these two objects and determining the conditions for convergence.

Agents believe they are in a noisy steady-state and know the form of the REE. Then the PLM is:

$$p_t = a + \nu_t$$

where the conditional expectation, $E_t^* p_{t+1} = a$. Here the asterisk denotes that the conditional expectation is not fully rational because the agent does not know the true parameter value. The ALM is then:

$$\begin{aligned} p_t &= R^{-1}\left[a + \mu - \frac{S_t}{e_t}\rho\sigma_p^2\right] \\ e_t &= \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{[(1-R)a+\mu]^2}{2\sigma_p^2}},\ 1\} \end{aligned}$$

Plugging into the learning rule, I get:

$$\begin{aligned} a_t &= a_{t-1} + t^{-1}\left[R^{-1}(a_{t-1} + \mu - \frac{S_t}{e_t(a_{t-1})}\rho\sigma_p^2) - a_{t-1}\right] \\ T(a) &= R^{-1}\left[a_{t-1} + \mu - \frac{S_t}{e_t(a_{t-1})}\rho\sigma_p^2\right] \end{aligned}$$

where $T(a)$ is a T-map which is a function that maps the agent's PLM to the ALM. Evans and Honkapohja (2001) show that the T-map can be used to compute local stability using a concept called E-stability. The E-stability principle states that locally stable rest points of the ordinary differential equation (ODE):

$$\frac{da}{dt} = T(a) - a$$

will be attainable under least squares learning. E-stability dictates that the "expectational" stability of a model depends on the signs of the eigenvalues evaluated at the rest point of the ODE. If all the eigenvalues have negative real parts, then the REE is locally stable. The fixed point of the ODE is:

$$a = \frac{\mu - \frac{S}{e}\rho\sigma_p^2}{R-1}$$
$$e = \min\{Ae^{-\rho R} - Ae^{-\rho R - \frac{S^2\rho\sigma_p^2}{2e^2}},\ 1\}$$

where (a, e) correspond to the steady-state values. I now state a proposition showing that the REE is locally stable under decreasing-gain learning.

Proposition 5. *If the sequence of shocks* $\{\epsilon_t^S\}_{t=0}^{\infty}$ *are such that* $|\epsilon_t^S| < \alpha$ *with probability 1 for all* t *and* $\alpha > 0$ *is sufficiently small, then the noisy steady-state REE is locally stable under decreasing-gain learning.*

Proposition 5 states that if R^{-1} is less than 1, then the system is E-stable which is satisfied in our model. In my model R^{-1} dictates the strength of the expectational feedback since higher values lead to larger coefficients on the expectations terms. By assumption, R^{-1} is always less than 1 since R is greater than 1. Hence the REE is locally stable under learning.

1.4.5 Constant-Gain

So far I have demonstrated the model properties under decreasing-gain learning. In the simulations I implement constant-gain learning, where agents weigh each observation with geometrically declining weights. This is appropriate because my application is a perpetual learning environment which is best captured by constant-gain learning. Constant-gain learning differs from decreasing-gain learning in the sense that agents are not weighing each observation equally. As γ increases, the agent weighs new evidence higher. I justify this for three reasons. First, constant-gain learning is a robust learning mechanism and is well-represented in the data (Malmendier and Nagel 2011, 2015). Second, when agents are worried

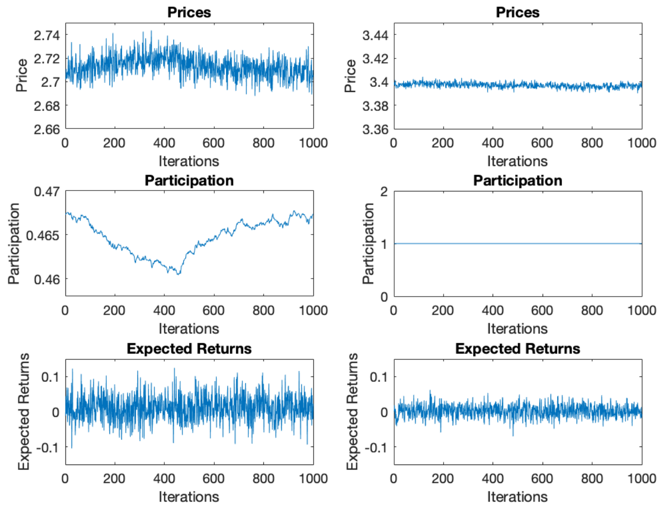

Figure 1.3: Constant-Gain Learning. For 1000 iterations, $\gamma = 0.05$, $\sigma_S^2 = 0.435$, $A = 1.15$, $\sigma_p^2 = 0.46$, $\mu = 1$, $R = 1.007$, $\rho = 0.45$. Left figure is limited participation, right figure is full participation.

about structural changes it is optimal to place higher weights on recent observations. Finally, constant-gain learning converges to a distribution around the REE, so we can still use the REE as a benchmark. The following is the recursive formulation for constant-gain learning:

$$p_{t+1}^e = p_t^e + \gamma[p_{t-1} - p_t^e], \text{ where } \gamma \in [0, 1].$$

Constant-gain learning requires a projection facility to ensure prices remain non-negative and plays a stabilizing role when the risk in endogenized. I implement a projection facility by endogenizing the shares, where the endogenous supply of shares is meant to capture asset float drying up when markets perform poorly. With endogenous supply, shares follow:

$$S_t = \{\min(S, \Phi p_t)\} V_t$$

where $V_t = 1 + \epsilon_t^S$ and $\Phi = \frac{S}{\bar{p}\xi}$, where $\bar{p}$ is the steady-state price and ξ is a fraction between 0 and 1. Here ξ is the fraction of steady-state price at which prices become endogenous.

Figure 1.3 depicts the learning about returns simulation with a constant gain and compares them to the full participation case. As we can see, the model with exogenous risk generates more volatility than the standard model, which is driven mainly by the participation channel. The key mechanism when learning about returns is as follows. When expected returns increase, participation increases. This leads to an increase in the price, which leads to a decrease in the expected returns which leads to a decrease in participation. With constant gain learning, this process leads to persistent fluctuations and adjustments in the learning process which generates more volatility than the standard case.

1.4.6 Learning about Risk

I now implement a learning rule where agents also have to learn about the risk σ_p^2. The most natural learning rule for σ_p^2 is one similar to the rule for prices, where agents regress the risk on a constant.[4] Then the learning rule for σ_p^2 is:

$$\sigma_{p,t+1}^2 = \sigma_{p,t}^2 + \delta[(p_t - p_{t-1}^e + \epsilon_t^D)^2 - \sigma_{p,t}^2] \text{ where } \delta \in [0,1].$$

where ϵ_t^D is the dividend shock. As before, there are 2 steady-state solutions for σ_p^2. Although Branch and Evans (2011) show the high risk steady-state is unstable under learning, it is not obvious if their results follow with the addition of a participation decision. Since $\bar{e}$ and σ_p^2 have no closed-form expression, a complete analytical solution is unavailable. Nevertheless, I find that the low risk steady-state is numerically stable under learning while the high risk steady-state is not. Figure 1.4 depicts the simulation with learning about risk.

I find that there is an increase in volatility in this simulation and in particular there is substantially more fluctuation in the participation rate. The main feedback mechanism with learning about risk is as follows. An increase in the subjective risk estimate σ_p^2 leads

[4] Alternatively, one could use different types of learning rules such as an autoregressive conditional heteroskedasticity (ARCH) model. Branch and Evans (2013) analyze this case and the qualitative results are similar.

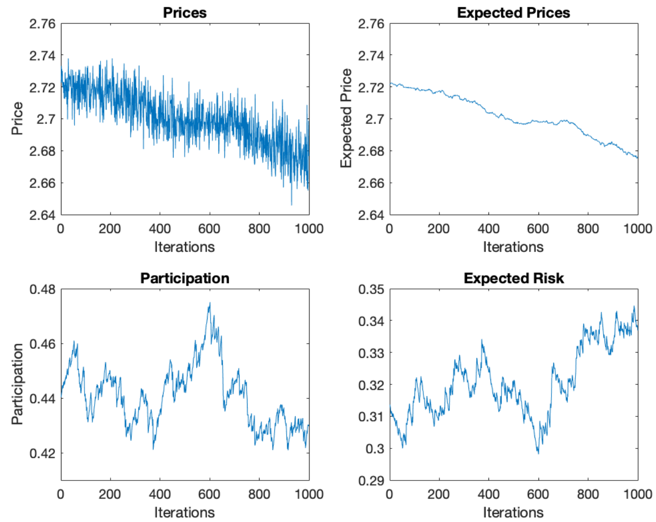

Figure 1.4: Endogenous Risk. For 1000 iterations, $\gamma_1 = 0.05$, $\gamma_2 = 0.0005$, $\sigma_S^2 = 0.435$, $A = 1.15$, $\sigma_D^2 = 0.28$, $\mu = 1$, $R = 1.007$, $\rho = 0.45$

to a decrease in the participation rate e_t which feeds back to a decrease in the asset price. Furthermore, a decrease in price will increase realized returns which leads to a temporary decrease in the subjective risk which further decrease the participation rate. This process continues until risk estimates are adjusted and then the mechanism moves in the opposite direction. The learning about risk feedback mechanism is the key driver of volatility in my framework.

1.5 Quantitative Analysis

In order to keep the model tractable and focus on the interplay between learning, the participation channel, and stock prices I made strong simplifying assumptions. Nonetheless, it is illustrative to calibrate the model to give some measure of quantitative importance to the participation channel.

1.5.1 Parameter Values

The parameters are calibrated according to the values in Table 1.2. The risk aversion ρ is calibrated to a value within the range of studies found in Babcock, Choi, and Feienerman (1993) at 0.45. The historical average real interest rate in the U.S. is 2.7% so I take the gross quarterly rate which is $R = 1.007$. Next the volatility of dividends σ_D^2 is taken from a Hodrick-Prescott (HP) filter of quarterly real historical stock market dividend data from 1927 to 2017 from Robert Shiller's database which is 0.28. For mean dividends μ I choose a value of 1 where the ratio of mean dividends to the standard deviation is sufficiently high such that the probability of negative dividends is unlikely. The volatility of supply σ_S^2 is taken from Baker and Wurgler (2000) who estimate the quarterly volatility of shares in the S&P 500 at 0.435. The gain parameter γ_1 is chosen similarly to past studies at 0.05. Branch and Evans (2006) show that this parameter value is consistent with the data.

Next, I choose the cost parameter $A = 1.15$ to match the mean participation rate consistent with the data. Then, the gain for the risk, γ_2 is 0.0005 which is calibrated such that the ratio of gains $\frac{\gamma_1}{\gamma_2}$ is sufficiently high to insure stability. Branch and Evans (2011) show that it is important that the gain for the risk be smaller than the gain for expected prices to insure stability. In particular, if the gain for the returns moves too much, it may be enough to move the learning path away from the REE. The endogenous share parameter ξ is chosen conservatively to be 0.3 which means that the shares start to become endogenous when prices decline to 30% of the fundamental value. Finally, I am interested in the unconditional moments so I take a long, transient simulation of two million iterations and burn-in the first one million.

Parameters	**Meaning**	**Value**	**Source/Target**
ρ	Risk Aversion	0.45	Babcock, Choi, Feinerman (1993)
σ_D^2	SD of Dividend	0.28	HP filtered dividend volatility
σ_S^2	SD of Supply	0.435	Baker and Wurgler (2000)
R	Real Interest Rate	1.007	Average U.S. Real Interest Rate
γ_1	Price Gain	0.05	Branch and Evans (2006)
γ_2	Risk Gain	0.0005	No prior reference
ξ	Endogenous Supply	0.3	Projection facility

Table 1.2: Parameter Values.

1.5.2 Moments

The moments I am interested in matching are as follows. The quarterly volatility of the HP filtered log prices from 1927 - 2017 is 0.132. The mean participation rate from the SIPP participation data from 1995 - 2013 is 0.373 for both direct and indirect stock holdings, and the volatility of participation rates is 0.008. I do a quarterly interpolation of the stock market participation data and HP filter it. I stress that this number is a noisy indicator of the true parameter and that future studies may want to find a more comprehensive way of measuring the volatility of participation rates. Nevertheless, it proves instructive to see how the model performs. The mean and standard deviation of annualized excess returns are 1.061 and 0.313. Finally, the autocorrelation of quarterly HP filtered log prices is 0.842.

Moment	No Risk	Risk	Data
$Sd(p)$	0.011	**0.031**	0.132
$Mean(e)$	0.368	0.37	0.373
$Sd(e)$	0.004	**0.008**	0.008
$Mean(Re)$	1.032	1.032	1.061
$Sd(Re)$	0.048	0.047	0.313
$\rho(p, p_{-1})$	0.625	0.959	0.842

Table 1.3: Moments Table. $Sd(\cdot)$ is the standard deviation, Re is the excess returns, and $\rho(\cdot)$ is the correlation coefficient.

1.5.3 Results

Table 1.3 documents the calibration results. As we can see, the learning about risk model does well on many dimensions, particularly when taking into account that the model is highly stylized. Even with the model abstracting away from serially correlated shocks we can see that the risk specification can match upwards of 50% of the volatility in stock prices. That is, as the stock market participation rates go down, stock volatility goes up. That is because, as participation rates go down, the market becomes more thin and volatility increases. I also find that learning about risk generates 3 times more volatility than learning about returns. Therefore, we can attribute most of the volatility from learning about risk rather than learning about returns. Next I can also match the volatility of participation rates which is 0.008. In contrast, the model without learning about risk is unable to generate the necessary volatility in the participation rate and generates a standard deviation of 0.004.

I can also match half of the mean excess returns at 1.032 and I do much better at matching the autocorrelation of stock prices at 0.959 while without learning about risk, the autocorrelation is 0.625. I am unable to match the standard deviation of excess returns at

0.047. I argue that the current model with iid shocks is not a good model for returns. With iid dividends and autocorrelation of prices, due to learning about risk, prices are moving in the same direction per period which removes the agents' capital gains.[5]

Learning about returns matters but learning about risk is necessary to generate volatility that matches the magnitudes found in the data. As before, with learning about risk the key mechanism is as follows. An increase in the subjective risk, decreases the participation rate which leads to a decrease in the price leading to a decrease in the subjective risk. This process continues until risk estimates are adjusted and then the mechanism moves in the opposite direction. These cyclical movements depend on the magnitude of the shocks and the magnitude of steady-state deviations. As enough data is realized, the process stabilizes around the steady-state values.

The model mechanism is also externally validated by survey responses provided in Arrondel et al (2014) where 20.7% of the sample stated the reason they do not invest in the stock market is that it is too risky. If one takes risk to be the variance of returns as in the context of our model, then it provides a natural explanation for limited participation rates and excess volatility in stock prices.

1.6 Conclusion

I have demonstrated that a simple asset pricing model with a participation decision can do well at matching moments of the data when allowing for agents to adaptively learn about risk and returns. The model adds a participation channel and endogenizes the risk which allows feedback effects to occur when combined with expectations and learning. The two key mechanisms are due to learning about risk and learning about returns. The learning about returns mechanism works as follows. When expected returns increase, participation increases, which leads to an increase in the price. This leads to a decrease in the expected

[5] Suitably extended versions of the model can explain returns such as in Adam et al (2017).

return and hence decreases participation. Similarly for learning about risk, when expected risk increases, participation decreases which decreases the price. This leads to a decrease in the expected returns which further increases the expected risk and hence further lowers participation. When risk estimates are finally corrected, the feedback mechanism moves in the opposite direction. The combination of these two channels are what leads to the agent's subjective risk being an important driver of stock price volatility, with learning about risk being quantitatively more important.

Future research will take the quantitative implications seriously by introducing serially correlated shocks and heterogenous agents. Also a focus of future empirical research will be to collect and better understand the time-series of participation rates.

Chapter 2

Expectations and Stock Market Participation: Theory and Evidence

2.1 Introduction

One of the main stylized facts in household finance is that stock market participation rates are significantly lower than predicted by standard asset pricing models, the so-called "limited participation puzzle" (Mankiw and Zeldes 1991). Although transaction costs, incomplete markets, and liquidity constraints all help explain limited participation, facets of the data are difficult to reconcile with these explanations (Guiso and Sodini 2013). For example, within the 80th percentile of the U.S. wealth distribution where a typical household has $200,000 in financial assets, 20% do not participate in the stock market (Campbell 2006).

Limited participation among the wealthy is difficult to reconcile solely with transaction costs or liquidity constraints and poses a significant challenge to the theory. More recently, there is strong empirical evidence that subjective returns are a key determinant of stock market participation (Hurd et al 2011). In addition, there is substantial heterogeneity in the subjective expectations of stock market returns within survey data (Dominitz and Manski 2011). These two facts suggest that differences in subjective expected returns may play an important role in explaining limited participation among households who are not liquidity constrained.

While information acquisition is a popular interpretation of differences in subjective expected returns (Van Nieuwerbaugh and Veldkamp 2010), private information sets are notoriously difficult to elicit from the data. In contrast, recent empirical evidence by Malmendier and Nagel (2011, 2016) demonstrates that "learning from experience", where households place greater weight on data that occurs within their lifetime, can generate heterogeneity in subjective expected returns independent of private information. Thus I provide an alternative interpretation where heterogeneity in subjective expected returns are generated by a learning process. Taking participation costs as a primitive, I ask to what extent subjective expected returns explain limited participation. The experimental results show that heterogeneous participation costs with rational expectations (RE) are not enough to explain limited participation. In addition, heterogeneity in subjective expected returns along with deviations from RE are needed.

I first write down a simple asset pricing model with heterogeneous participation costs where, depending on the distribution of costs, limited participation is a steady-state outcome. This allows me to have a framework to understand the mechanisms behind limited participation. While the most common approach in testing expectations-based models is using survey data, stock market participation is difficult to test because dynamic participation data is generally unavailable. In addition, standard surveys like the SCF and PSID do not elicit subjective expected returns. Moreover, researchers have little to no control over the subjects' information sets given that the environment is constantly changing, which is especially critical in asset pricing. In particular, the data-generating process (DGP) is unknown to the researcher and important parameters must be calibrated or estimated.

Instead I test the implications of the theory, using the experimental method, which provides an ideal tool to analyze models of limited participation. In the experimental laboratory the researcher is able to control the fundamentals of the environment and jointly elicit expectations along with individual participation data while having complete knowledge of subjects' history and information sets. It also provides a tool to test the comparative statics of limited participation models while analyzing the effects of induced participation costs, subjective expected returns, and learning on individual participation decisions. Thus, the laboratory presents an environment to diagnose the causes of potential deviations from the theory, not identifiable solely using survey data.

Models with forward-looking expectations are inherently difficult to test in the laboratory. Rational expectations requires that agents not only make optimal forecasts, but optimal decisions conditional on their forecasts. This joint optimization task can prove challenging for subjects in the laboratory. A learing-to-forecast (LtF) experiment separates these tasks such that subjects are only asked to forecast stock market prices while the optimization decision is done by an automated auctioneer. In particular, subjects are incentivized based on the accuracy of their forecasts and an automated auctioneer constructs an optimal portfolio based on the subjects' forecasts. Hence a LtF experiment allows for a clean way to elicit individual expectations while preserving the self-referential nature of beliefs and outcomes.

I extend the standard LtF asset pricing experiment (Hommes et al 2005) to include a stock market participation decision. Subjects make both a participation decision and forecast prices in each period. As a preview, I find mixed results for the homogeneous expectations model with participation costs. While the main predictions of the model do well, where mean prices and participation rates are generally lower in treatments with higher participation costs, subjects with higher induced costs do not necessarily have lower subjective returns. Moreover, the heterogeneity in subjective returns are more disperse and persistent than predicted by the benchmark model. Given these results, I analyze the stock market participation decision directly. I find that while participation is increasing in subjective expected returns, contrary to the model, participation also depends on past realized returns. That is, higher past realized returns from participating increases the likelihood of stock market participation.

I next look at a learning-based explanation of the observed heterogeneity in subjective returns. Since the fundamentals and aggregate variables are common knowledge, heterogeneity in subjective expected returns cannot be explained by private information. The experimental method provides me with a novel tool to look at the subjects' expectations updating rule directly. I find strong evidence that subjects update their expectations conditional on realized returns. Contrary to standard learning models, subjects who experience high returns have higher subjective returns and subjects who experience low returns have lower subjective returns than subjects who do not participate. In particular, subjects overweigh the price trend when experiencing high returns and under-weigh the price trend when experiencing low returns, with low returns having a stronger quantitative impact. Thus I find evidence that limited participation can be perpetuated by subjects with low subjective returns due to low experienced returns.

Finally, my experiment allows for documentation of a novel behavioral phenomena. I find evidence of "discouraged investors" who exit the stock market for the remainder of the experiment after receiving consecutive low returns relative to the risk-free return. This finding provides direct experimental evidence of a behavioral finance phenomena as in Strahilevitz et al (2011) where investors are reluctant to hold stocks once they realize low payoffs.

2.2 Literature Review

This paper contributes primarily to three literatures. First, I contribute to the literature on endogenous participation in asset markets. The phenomenon of limited participation is first documented by Mankiw and Zeldes (1991). The first paper to endogenize participation in an asset pricing setting is Allen and Gale (1994). Following Allen and Gale (1994), there have been numerous extensions such as Orosel (1998), Vissing-Jorgensen (2002), Gomes and Michaelides (2005), and more recently Shin (2018).

I also contribute to the empirical literature on limited participation, in particular, to the strand of the literature dealing with subjective returns. Guiso and Sodini (2013) provides a survey of alternative explanations such as participation costs, trust, and non-standard preferences. Hurd et al (2010) and Arrondel et al (2014) use novel datasets on expectations and asset positions to estimate a causal effect between subjective expected returns and likelihood of stock market participation. Malmendier and Nagel (2011) shows that living through a period of low stock market returns reduces the likelihood of stock market participation. My paper consolidates these previous findings by establishing the connection between subjective expected returns, experienced realized returns, and its subsequent effects on stock market participation.

Second, I contribute to the literature on LtF asset pricing experiments following Hommes et al (2005). Following the pioneering work of Hommes et al (2005), there have been numerous extensions. For a survey of the literature see Hommes et al (2011) and Duffy (2016). A related paper is Hennequin (2018), who analyzes the effects of stock market experiences on bubble formation. My paper differs from theirs in that stock market experiences are endogenous in my experiment and that I focus on limited participation and not bubbles.

Finally, this paper contributes to the literature on adaptive learning and asset pricing. This paper follows a strand of literature put forth by Marcet and Sargent (1989) and Evans and Honkapohja (2001) which relaxes the RE hypothesis and replaces it with an econometric learning rule. The first paper to analyze learning in an asset pricing model is Timmermann (1993) who shows that adaptive learning can generate excess volatility. Branch and Evans (2010, 2011) use a similar asset pricing model to explain bubbles and crashes and regime-switching returns. Finally Adam et al (2016) formally test a consumption asset pricing model with learning. As far as I know, this paper is the first to explore endogenous participation in a LtF environment.

2.3 Model

First I describe an asset pricing model with an endogenous participation decision and heterogeneous participation costs with the goal of implementing it in the laboratory. This model provides a framework for understanding the determinants of limited participation in equilibrium. I then introduce a special case of the model which includes parameterizations and features of the experimental design.

2.3.1 Endogenous Participation

The benchmark model is a CARA asset pricing model with heterogeneous participation costs and an endogenous participation decision. In order to map the model to the simplest laboratory environment, I introduce two agents that make the participation and portfolio decisions separately. This formulation allows me to get a closed-form solution by making the participation decision linear in expectations, while keeping the same comparative statics as the model with one risk-averse agent as in Shin (2018).

Time is discrete and continues forever. There are M households who are each paired with a financial advisor. Households have CARA utility of the form: $u^P(c_t) = -e^{-\rho c_t}$, where ρ is the coefficient of absolute risk aversion and financial advisors are risk-neutral: $u^M(c_t) = c_t$. Both households and advisors are 1-period maximizers. There are two assets: a risk-free 1-period bond which pays a gross return $R = 1 + r > 1$ and a risky asset with price p_t that pays an ex-ante dividend $D_t = \mu + \epsilon_t^D$, where $\epsilon_t^D \sim N(0, \sigma_D^2)$. The risk-free 1-period bond is in infinitely elastic supply. The supply of assets is iid with mean S and follows $S_t = S + \epsilon_t^S$, where $\epsilon_t^S \sim N(0, \sigma_S^2)$.

To participate in the risky asset market, financial advisors have to pay a research cost κ^i which is specific to each advisor i. Here the research cost is a proxy for both financial and psychological participation costs that are documented in the literature (Guiso and Sodini 2013). The financial advisor makes the participation decision and the household makes a subsequent portfolio allocation decision. Advisors make a fraction $\alpha \in [0,1]$ of the total payoff from the portfolio returns and the household keeps fraction $1-\alpha$. I assume that the advisor and households do not have the technology to change this contract.[1]

Let $W_{t+1} = RW_t + (1-\alpha)x_t^i(p_{t+1} + D_{t+1} - Rp_t)$ be the next period's wealth for the household, where x_t^i is the fraction of wealth held in the risky asset by household i. CARA utility insures that optimal asset holdings are independent of wealth such that W_t need not be indexed by i. I assume that the endowment is sufficiently large that agents are not liquidity constrained. Since all the random variables are normally distributed and the household has CARA utility, the maximization problem is:

$$\max_{x_t^i} \; E_t W_{t+1} - \frac{\rho}{2} V_t W_{t+1}$$
$$\text{s.t. } W_{t+1} = RW_t + (1-\alpha)x_t^i(p_{t+1} + D_{t+1} - Rp_t)$$

where i denotes household i. If the financial advisor does not participate, then the household places all their wealth into the riskless bond. Taking the F.O.C.:

$$x_t^i = \frac{E_t p_{t+1} + \mu - Rp_t}{(1-\alpha)\rho\sigma_p^2} \tag{2.1}$$

where $V_t W_{t+1} \equiv \sigma_p^2$. Equation (2.1) is the standard mean-variance asset demand adjusted for α. The participation decision for the advisor is:

$$\max_{n_t^i} \; \alpha x_t^i (E_t p_{t+1} + \mu - RE_t p_t) - \kappa^i$$

[1] α can be microfounded through an optimal contract design.

where n_t^i is 1 if the advisor decides to participate in the risky asset market and 0 otherwise. $E_t p_{t+1} + \mu - RE_t p_t$ is the expected return from participation for the advisor and αx_t^i is the fraction of the profits that the advisor makes per share x_t^i. There is an expectation E_t on p_t because I assume p_t is unknown to the advisor before participating. I rationalize this by treating the participation decision as a market order where realized prices can be different from the quoted price during the time of the order. In contrast, the portfolio decision is a limit book, where agents give the auctioneer an asset position for every price. While this is reasonable for the portfolio decision, it is unreasonable for the participation decision because the decisions are made sequentially.

I solve by backward induction, where the advisor takes x_t^i as given. Then the cutoff decision for the advisor is:

$$n_t^i = \begin{cases} 1 & \text{if } E_t p_{t+1} + \mu - RE_t p_t \geq b\sqrt{\kappa^i} \\ 0 & \text{else} \end{cases} \tag{2.2}$$

where $b \equiv \sqrt{\frac{1-\alpha}{\alpha}\rho\sigma_p^2}$. Notice that the threshold in the participation decision depends positively on the participation cost, risk aversion, and negatively on the surplus share α. I can rewrite $k^i \equiv b\sqrt{\kappa^i}$ and call it the effective participation cost for advisor i. The advisor only cares about the subjective expected returns relative to the participation cost per fraction of earnings. Thus, I can always rewrite the distribution of costs as a function of bk^i, where bk^i is the per asset cost of participation for advisor i. The optimization problem reduces to a mean-variance portfolio problem with a risk-neutral participation decision. Again, the structure of the model maps into the simplest laboratory implementation.

Aggregate participation is $N_t = \frac{1}{M}\sum_i n_t^i$ which is the fraction $N_t \in [0,1]$ of the population participating in the risky market and aggregate asset holdings is $X_t = \sum_M x_t^i$.

Market-clearing $N_t X_t = S_t$ implies that the equilibrium pricing equation is:

$$p_t = R^{-1}\left[\frac{1}{N_t}\sum_{N} E_t p_{t+1} + \mu - \frac{S_t}{N_t}(1-\alpha)\rho\sigma_p^2\right] \tag{2.3}$$

where the summation on $E_t p_{t+1}$ is over all stock market participants N_t. Since market-clearing implies prices must be positive, the participation rate N_t is always positive in equilibrium and hence Equation (2.3) is well-defined. Given that prices are not defined for $N_t = 0$, I utilize an automated mutual fund who always participates during the experiment.

Participation affects prices through two channels. First, only expectations of market participants are priced so participation endogenously reduces the number of agent's expectations that are priced. Second, participation affects prices directly through the supply S_t. As N_t increases, prices increase. In particular, lower participation rates lead to lower prices and higher participation rates lead to higher prices. When $N_t = 1$ the model collapses to the standard model without a participation decision. Thus, N_t acts as a demand multiplier for aggregate asset holdings. Because N_t is a decreasing function of $E_t p_t$, higher $E_t p_t$ leads to lower participation which leads to a lower price. This is because higher $E_t p_t$ lowers expected returns in that the advisor expects to pay a higher price to hold the asset. The advisor always wants to pay the lowest price possible.

Since b is a parameter, for simplicity I set this to 1. I do this so that I can arrive at a form similar to Hommes et al (2005) but also to analyze the interactions between the expectational feedback mechanisms directly.

$$p_t = R^{-1}\left[\frac{1}{N_t}\sum_{N} E_t p_{t+1} + \mu - \frac{S_t}{N_t}\right] \tag{2.4}$$

Although it is standard to analyze the case with $S = 0$, I analyze the case with $S > 0$ because it is more interesting to analyze the direct effects of participation especially if we believe feedback between expectations and prices matter. If $S = 0$, the participation effects drop out from Equation (2.4) and participation only indirectly affects prices. I rationalize my approach by arguing that in models with reasonable learning rules, if the steady-state price is a constant, agents can learn the steady-state values (Bao et al 2017). In particular, any LtF experiment with risk-averse preferences will have to implement a positive supply.[2]

2.3.2 Steady-state Equilibrium

In order to solve for the steady-state, I specify the distribution of participation costs. I make the distribution of costs uniform $U(0, \bar{k})$, where $\bar{k}$ is the upper support of the distribution and represents the agent with the highest participation cost. I assume that each agent can be represented by a point in the distribution, that is, agents are equally spaced along the distribution and no two agents can have the same cost. Implicitly, I also assume that one agent has a participation cost of zero. Then I can represent the distribution with a participation cost function $c(N_t) = N_t\bar{k}$, where $N_t = 1$ is full participation and the aggregate cost of full participation is $\bar{k}$. $c(N_t)$ is the limiting cost function that arises as the number of agents M approaches infinity.

Once I specify the cost function, the equilibrium participation rate is characterized by the marginal agent who is indifferent between participating and not participating. I set expected returns equal to the cost function $c(N_t) = N_t\bar{k}$. In the steady-state, $p_t = p_{t+1} = \bar{p}, S_t = S$. Then the steady-state values are:

$$\bar{p} = \frac{\mu - \frac{S}{\bar{N}}}{R - 1}$$
$$\bar{N} = \min \{\sqrt{\frac{S}{\bar{k}}}, 1\}$$

The steady-state is the fundamental price and participation rates that are equilibrium best-responses with a given cost distribution for all agents. Steady-state participation $\bar{N}$ depends on the supply S and the upper support of the cost distribution $\bar{k}$. The minimum operator

[2]In Hommes et al (2005) setting $S = 0$ implements risk neutrality.

insures that the participation rate is never greater than 1.

The comparative statics are intuitive, higher cost implies that the participation rate is lower and prices are lower. Higher supply when $S \geq 1$ implies that there are more shares and the participation rate is higher and prices are higher. The key feature of the model is that higher costs $\bar{k}$ lead to lower participation rates $\bar{N}$ because it lowers aggregate demand for stocks. The specific mechanism is that participation costs k^i reduce the expected utility of participation and acts as a wedge for agents with high costs.

2.3.3 Parameterization

To implement the model in the laboratory, I make three additional assumptions. First, I make the dividends constant $D_t = \mu$. This has no bearing on the equilibrium but simplifies the instructions. Second, I normalize S to 1. Finally, I create bounds on the advisor's excess returns. That is, advisors can only make $\underline{\pi}$ to $\bar{\pi}$ excess returns from participation. This is to map closely to the experiment where payoffs from stock market participation are bounded and has no bearing on the model equilibrium since these thresholds are chosen such that it will not bind in the model. I rationalize my experimental design choices in the subsequent section. Then the steady-state values become:

$$\bar{p} = \frac{\mu - \bar{N}^{-1}}{R - 1}$$
$$\bar{N} = \min\{\sqrt{\bar{k}^{-1}}, 1\}$$

2.3.4 Model Hypotheses

The treatments vary the cost parameter $\bar{k}$.[3] Based on the theory in the last subsection I present the following hypotheses which are tested in the experiment. The model has a

[3]I do not vary the supply parameter S since this information is usually not given to the subjects in a LtF experiment.

clear implication that subjects with lower participation costs should participate more than subjects with high participation costs. Any deviations from this behavior must be due to factors unrelated to participation costs.

Hypothesis 1: *Treatments with a higher cost distribution parameter* $\bar{k}$ *have lower mean prices* $\bar{p}$ *and participation rates* $\bar{N}$. *In particular, subjects with higher induced participation costs* k^i *have lower mean participation rates* $\bar{n}^i$.

Hypothesis 1 is a direct test of the aggregate properties of the model. Since a higher cost parameter implies there are more subjects with higher induced participation costs , the steady-state participation rate is lower. Similarly, the lower participation rate implies a lower price since there is lower aggregate demand for the asset in the steady-state.

Hypothesis 2: *Participation is increasing in subjective expected returns* $E_t^i p_{t+1} + \mu - RE_t^i p_t - k^i$. *In particular, participation is an increasing function of the 2-period ahead forecast* $E_t^i p_{t+1}$ *and a decreasing function of the 1-period ahead forecast* $E_t^i p_t$, *and induced cost* k^i.

Hypothesis 2 falls from the utility function and looks more closely at the individual participation decision. Here I allow for subjective expected returns to be subscripted by i as in $E_t^i p_t + \mu - RE_t^i p_t - k^i$. In the model, subjective expected returns for agent i are increasing in the 2-period ahead forecast and decreasing in the 1-period ahead forecast and induced participation costs. Thus the model predicts that at the individual level, subjects with higher subjective returns should be more likely to participate than those with lower subjective returns.

Hypothesis 3: *Under adaptive learning, all subjects update their subjective returns towards the forecast error and subjective returns converge over time. Heterogeneity in subjective re-*

turns $Var_t(E_t p_{t+1} + \mu - RE_t p_t - k^i)$ *is due to heterogeneity in costs* $Var(k^i)$. *Moreover, all subjects have expected returns* $E_t p_{t+1} + \mu - RE_t p_t$ *net participation costs* k^i.

Hypothesis 3 investigates the model implications under standard learning mechanisms i.e. learning processes that update expectations towards the forecasting error. In the benchmark model, under homogeneous expectations, since all subjects have access to the same history of aggregate variables, their forecasts can only differ due to their induced participation cost. With rational expectations (RE), objective and subjective probabilities must be equal so differences in expectations must be due to differences in their induced costs. With adaptive expectations, if the underlying rational expectations equilibrium (REE) is a constant plus a noise, then both the 1-period and 2-period ahead forecasts must be equal and subjective returns should converge to the objective returns.[4] In particular, the dispersion of subjective returns converges to the variance of the induced cost distribution. The strong empirical evidence on "learning from experience" and the stylized facts on heterogeneous subjective returns in survey data suggest an alternative hypothesis:

Hypothesis 3b: *Subjects' with higher past returns have higher subjective expected returns and subjects' with lower past returns have lower subjective expected returns i.e. subjects update their forecasts differently depending on past payoffs. Heterogeneity in subjective returns persist over time.*

Empirical evidence from survey data along with psychological evidence demonstrates that agents who experience lower returns expect lower returns and higher returns expect higher returns. The experiment provides an ideal environment to test this hypothesis since the economic fundamentals are common knowledge and subject' information sets are known to the researcher. Moreover, the introduction of a 1-period ahead forecast elicitation along with

[4] Shin (2018) demonstrates that in this model, there exists a noisy steady-state REE which is a constant plus a noise that is E-stable.

the 2-period ahead forecast, which will be described in detail in the next section, provides a novel tool for understanding subjects' expectations updating process.

2.4 Experimental Design

The experiment was designed to test Hypotheses 1-3. I utilize a learning-to-forecast (LtF) design. In a LtF experiment, the experimenter elicits expectations of the subjects while the optimization decision is done through an automated auctioneer. This provides a clean way of eliciting expectations without having to deal with the potential complications of the joint forecasting-optimization task. My experiment differs from the standard design in that I introduce a simple, binary optimization task in the form of a participation decision.

I implement the model in Section 2.3 in the laboratory. I vary the induced cost distribution $\bar{k}$ and test the model against a baseline version where I shut down the participation decision. This allows me to test the comparative statics of the model as well as determinants of the participation decision. In particular, I implement 4 treatments. The first treatment tests the benchmark case which is Hommes et al (2005) with the addition of a 1-period ahead forecast. The rest of the treatments vary the cost parameter $\bar{k}$ from 0 to 1.5 to 4 which is a shift in the distribution of participation costs. I conduct 16 sessions with 4 sessions per treatment. Each session has 8 subjects for a total of 128 subjects.

2.4.1 Experimental Instructions

The experiment was programmed using oTree. The experimental design is a standard LtF asset pricing experiment with the addition of a stock market participation decision. 8 subjects are told they are advisors to a household.[5] Households need advice on whether they

[5]Originally subjects were advisors to fund managers. The current formulation just changes "fund manager" to "households" within the instructions and changes nothing else in the experimental setting.

should participate in the stock market along with a forecast of stock prices in each period (in particular for this period t and the next $t+1$). Subjects are told that the household will make an optimal portfolio allocation split between stocks and risk-free bonds conditional on these forecasts and completely follow any participation advice. At the beginning of the experiment, each subject is randomly given an induced participation cost drawn from a uniform distribution without replacement. Subjects are told the distribution $U(0, \bar{k})$ and the drawn induced cost k^i is fixed throughout the experiment. The units are in francs which is a common experimental currency.

As is standard in the literature, subjects are only told qualitative information about the data-generating process. In particular, they are told that higher price forecasts lead to higher asset purchases and that stock market prices are determined by supply and aggregate demand. They are told that aggregate demand depends on the decisions of the other households who are also advised by other subjects in the experiment. The exact number of subjects are not revealed. Moreover, they are also provided with information on the dividend μ and interest rates R which are also fixed. Subjects are told that there is one pension fund who always participates in the market. Finally, they are told there is a small, exogenous demand for stocks by private investors. This is a proxy for the stochastic supply which is formally equivalent to noise traders in the model.

2.4.2 Pricing Mechanism

The pricing mechanism is generated by Equation (2.4) which is repeated here:

$$p_t = R^{-1}\left[\frac{1}{N}\sum_i E_t^i p_{t+1} + (1-h(N_t))\bar{p} + \mu - \frac{S_t}{N_t}\right]$$

where $1-h(N_t)$ is the weight placed on the automated fund which depends on the number of participants in the market and $\bar{p}$ is the steady-state price implied by the model parameters.

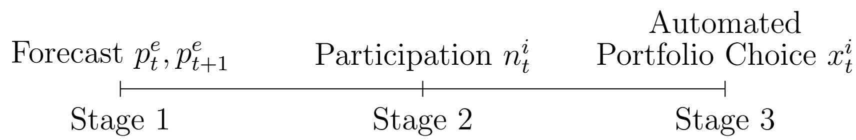

Figure 2.1: Timeline

The automated pension fund participates every period and forecasts the model implied fundamental price $\bar{p}$ to insure prices exist. The automated pension fund plays two roles: first they insure that prices exist every period. Second they play a stabilizing role given that they forecast the fundamental price every period.[6]

2.4.3 Timing and Decisions

There are 50 periods and each period is divided into 3 stages. In the first stage, subjects are told to make a 1 and 2-period ahead forecast of prices p_t^e, p_{t+1}^e. p_{t+1}^e is called a 2-period ahead forecast because p_t is not revealed until the end of the period thus using an information set up to period $t-1$. I add an upper bound $\bar{p}_{t+1}^e = 100$ similar to past studies to rule out potential bubbles. In the second stage, subjects are asked to give participation advice n_t^i to the household to either participate in the stock market or not. At the third stage, all decisions are given to the automated auctioneer who clears the market. The price p_t and participation rate N_t are then revealed. Throughout the experiment, subjects are provided with the history of past prices, past participation rates, subject-specific expectations, and past payoffs. The following timeline shows the sequence of decisions:

It is important to note that my experiment makes two deviations from the standard LtF design. First, subjects are asked to make a participation decision in the form of advice to households. Second, subjects are asked to make a 1-period ahead forecast along with their 2-period ahead forecast. This is because the profits from participating in the stock market

[6]Hommes et al (2005) also have a stabilizing fund and show that removing it has no qualitative impact on their results.

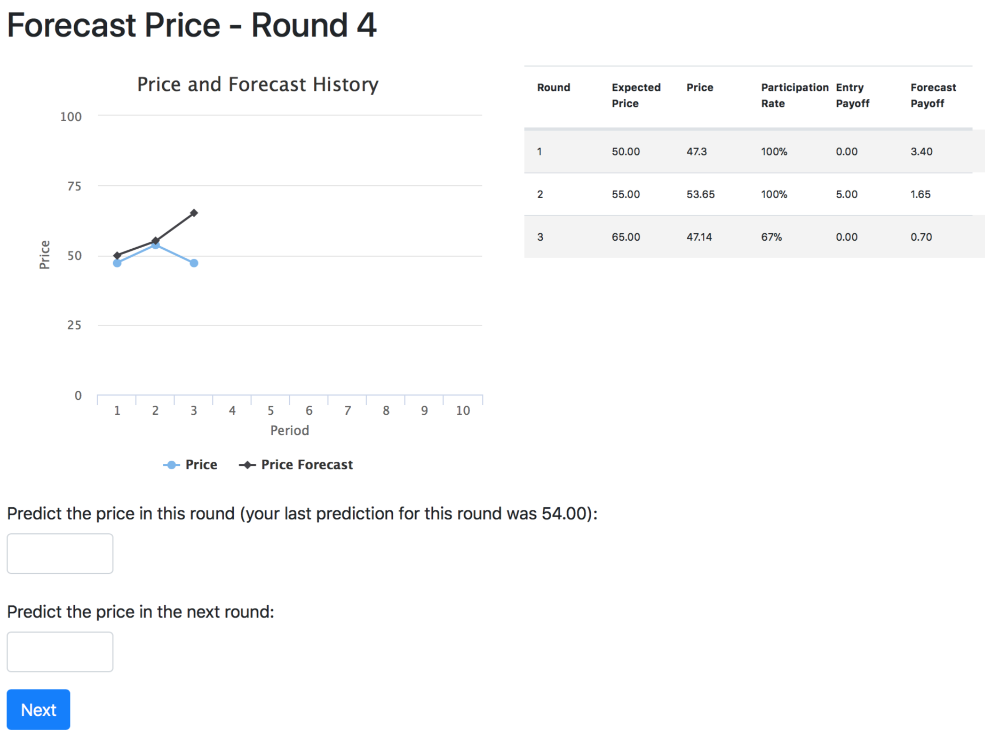

Figure 2.2: Experimental Screen: Forecasting Prices

depend on their realized returns which depends on p_t which is unknown to the subjects when making the participation decision. To make an informed decision, subjects must also make a forecast of the 1-period ahead price. I also provide subjects with their subjective expected returns $p^e_{t+1} + \mu - Rp^e_t$ conditional on their forecasts along with their previous forecast of p^e_t. Figures 2.2 and 2.3 provide screens of the experiment for the forecasting and participation decisions.

2.4.4 Parameterization

I follow Hommes et al (2005) and choose the experimental parameters as follows: the mean dividend $\mu = 3$, the gross interest rate $R = 1.05$, and the standard deviation of the supply shock $\sigma^2_S = 0.25$. For simplicity, I set the supply $S = 1$ which acts as a natural benchmark normalization but future studies can explore what happens when you vary S. In Treatment

Participation - Round 4

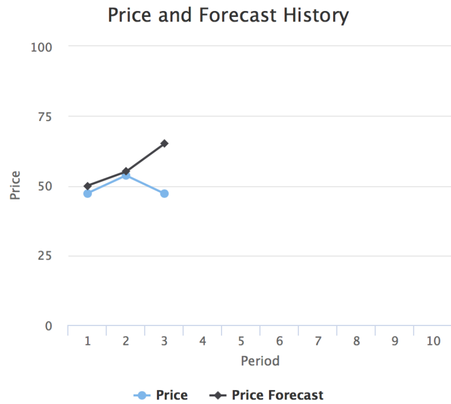

Round	Expected Price	Price	Participation Rate	Entry Payoff	Forecast Payoff
1	50.00	47.3	100%	0.00	3.40
2	55.00	53.65	100%	5.00	1.65
3	65.00	47.14	67%	0.00	0.70

Do you want to participate this period? (Your cost of participating is 2.0. Your prediction for this round is 50.00 and next round is 55.00. You expect the manager's profits to be 4.5). :

- No
- Yes

Next

Figure 2.3: Experimental Screen: Participation Decision

2, all subjects have an induced participation cost k^i of zero which corresponds to $\bar{k} = 0$. The steady-state price in this case is $\bar{p} = 40$. Treatments 3 and 4 have cost parameters at $\bar{k} = 1.5$ and 4 respectively which lead to $\bar{N} = 0.82$ and 0.5 and $\bar{p} = 35.5$ and 20.

With 8 subjects and 1 automated fund, this implies steady-state participation is 6-7 subjects participating in Treatment 3 and 4-5 subjects participating in Treatment 4. The predictions are in between because the discrete nature of the experiment leads to steady-state predictions that are fractions. The cost parameters were chosen such that there is sufficient variation within steady-state participation rates while the steady-state price is sufficiently away from 50 since the average initial guess tends to be around 50. Table 2.1 shows the parameterization.

Parameter	Meaning	Value
μ	mean dividend	3
R	gross interest rate	1.05
σ_S^2	supply shock	0.25
S	supply	1
$\bar{k}$	cost distribution	0 to 4

Table 2.1: Experimental Parameters

2.4.5 Payoffs

The conversion rate in the experiment is 15 points = \$1. The rate was determined through pilot studies to insure sufficient compensation for subjects' time. Subjects are paid according to two criteria: forecasting accuracy and the 1-period return from participating in the stock market. Forecasting payoffs are given by the following equation:

$$\pi_t^f = \frac{16}{2 + |p_t - p_{t,t-1}^e| + |p_t - p_{t,t}^e|} \tag{2.5}$$

where $p_{t,t-1}^e$ is the 2-period ahead forecast of price p_t made in period $t-1$ and $p_{t,t}^e$ is the 1-period ahead forecast of price p_t made at the beginning of period t. Thus subjects are paid based on the accuracy of both their forecasts to insure truthful revelation. Equation 2.5 is an adjusted Brier score with sharp declining payoffs to avoid the flat maximization problem (Camerer 2003). In particular, if the payoff function is sufficiently flat such that differences from the optimum only lead to small changes in the payoffs, then the payoff function may not be sufficient to induce truthful revelation. The maximum payoff subjects can make per period from forecasting is 8 points and sharply declines with the forecast error.

Participation payoffs are given by:

$$\pi_{i,t}^p = \begin{cases} \min\{5, 3 + MP\} & \text{if } n_{t-1}^i = 1 \text{ and } p_t + \mu - Rp_{t-1} - k^i \geq 0 \\ \max\{1, 3 + MP\} & \text{if } n_{t-1}^i = 1 \text{ and } p_t + \mu - Rp_{t-1} - k^i \leq 0 \\ 3 & \text{if no participation } n_{t=1}^i = 0 \end{cases}$$

where $MP = p_t + \mu - Rp_{t-1} - k^i$ is the net 1-period return and n^i_{t-1} is the participation decision where $n^i_{t-1} = 1$ is to participate in the stock market. Hence, if there is positive profit from participation, subjects make a high payoff up to 5, else they make a low payoff down to 1, which is linear in realized returns. Subjects can also take the risk-free option which is chosen in between the high and low payoffs at 3. Notice that participation payoffs depend on subjects' participation decision in period $t - 1$. This is because the 1-period return depends on the realization of next period's prices. The payoffs are incentive compatible because since subjects' subjective expected returns are their best forecasts, risk-neutral subjects can do no better than participating if it is positive and not participating otherwise. In equilibrium, subjects with high participation costs will lose payoff in expectation.

I now rationalize my experimental design choices. With unbounded payoffs, the maximum payoff would be 103 and the minimum would be -98 because $MP = p_t + \mu - Rp_{t-1} - k^i$. Thus, to make participation payoffs comparable to forecasting payoffs, I chose to bound them at a level similar to the forecasting payoffs. In particular, I ran a pilot study to see what average forecasting payoffs were. If one task gives more payoff on average than the other, then subjects have an incentive to pay more attention to one task. I choose the loss and gain thresholds to be symmetric such that the expected payoff from participation are symmetric. If there is more downside risk then there may be over-participation independent of the model since subjects can potentially gain much more from participating than not. Future studies can explore the effects of changing the threshold payments.

Next, I choose to place cutoffs on the payments rather than give subjects an endowment due to potential wealth effects. Of course, since my model removes wealth effects from consideration, I take the interpretation from Crockett et al (2018) that subjects may bring intrinsic utility to the laboratory that differs from induced utility. Thus to avoid potential complications, I remove wealth effects from consideration. Future studies can explore the implications of adding endowments with unbounded payoffs.

2.4.6 Subjects and Payments

Subjects are undergraduate students from the University of California, Irvine. As part of the instructions, each subject is required to complete a quiz to test comprehension. Instructions are reproduced in the appendix. Subjects are recruited for two and a half hour sessions but a typical session lasts two hours, including instructions and the quiz. Treatment 1 lasts around 1 hour and 45 minutes on average because it includes no participation stage. At the end of the last round and before the realization of payoffs, subjects are given a risk elicitation task which takes around 10 minutes. This is to insure that priming does not play a role in the risk elicitation.

Payoffs are earned from every period for each task. One task is selected at random at the end of the experiment and subjects are given points for that task. Random selection helps insure that subjects pay equal attention to both tasks. The mean payment was $18.49 including the show-up payment of $7. The average payment for the risk elicitation was $2.21.

2.4.7 Risk Elicitation

I elicit risk aversion using a multiple-price list (MPL) as in Drichoutis and Lusk (2016). A MPL provides a list of safe and risky lotteries to subjects and asks them to choose between them. After the survey, the experimenter utilizes a randomization device and one lottery on the list is played. Subjects receive a payment based on their choice for that lottery. The number of safe choices provides an estimate of their risk aversion parameter. Csermely and Rabas (2016) shows that the most reliable risk elicitation surveys are in the form proposed by Drichoutis and Lusk (2016).[7]

[7] Their criterion was based on predictability and consistency. For more details please see their paper.

Treatment	Cost Parameter	S.S. Price	S.S. Participation
1	N/A	40	N/A
2	0	40	100%
3	1.5	35.51	81.67%
4	4	20	50%

Table 2.2: Treatment Summaries

Treatment	Session	Mean Price	S.S. Price	Mean Participation	S.S. Participation
1	1	46.23	40	N/A	N/A
1	2	41.62	40	N/A	N/A
1	3	42.84	40	N/A	N/A
1	4	43.02	40	N/A	N/A
2	5	41.44	40	72.6%	100%
2	6	43.89	40	61.9%	100%
2	7	47.04	40	78%	100%
2	8	40.38	40	82.14%	100%
3	9	45.13	35.51	48.15%	81.67%
3	10	48.37	35.51	63.62%	81.67%
3	11	50.77	35.51	65.58%	81.67%
3	12	40.42	35.51	45.32%	81.67%
4	13	28.70	20	63.83%	50%
4	14	25.82	20	60.13%	50%
4	15	25.76	20	55.56%	50%
4	16	29.51	20	66.45%	50%

Table 2.3: Mean Price and Participation Rates per Session

2.5 Experimental Findings

Table 2.2 documents the four treatments along with the model predictions for the steady-state price and participation rates. Table 2.3 provides aggregate prices and participation rates per treatment and session. While the benchmark model generally does well at the aggregate level, some of the comparative statics are counter to the theory. I provide some justification for why this is the case and develop an extension of the model which can explain the deviations in Treatment 3. I first start explaining the data at the aggregate level and then at the individual level. For the individual level data, I run a probit regression to explain the determinants of the participation decision. I then, demonstrate that subjects have different subjective returns depending on past experienced returns. I finally show that differences in subjective returns are due to differences in how subjects update their expectations.

I report summary statistics at the treatment level. Table 2.4 provides summary statistics for prices, participation rates, subjective expected returns, and realized returns per treatment. Consistent with survey evidence, I first find that there is large heterogeneity among subjective returns. Next I find that subjective returns for participants are systematically higher than non-participants across periods. The fact that Treatment 2 has identical induced costs among subjects suggests that heterogeneous subjective returns are due to factors other than induced costs. Finally, I find that the heterogeneity in subjective returns are persistent throughout the experiment, that there is systematic disagreement on fundamental values over time.

Figures 2.4 and 2.5 provide graphs of the aggregate price and participation rates for each treatment. Each graph has 4 series which are represented by the different sessions along with a dotted series which represents the model steady-state predictions. Consistent with Hommes et al (2005) I find heterogeneity across sessions but each treatment follows a general pattern.

Treatment 2						
	25th pct	Median	75th pct	Mean	Std. Dev.	N
Subj. Returns	0.28	1.03	1.95	1.13	3.44	1632
Realized Returns	0.03	1.13	2.08	0.99	2.14	1126
Price	38.99	43.64	46.19	43.19	5.75	204
Participation	0.67	0.78	0.89	0.74	0.19	204
Subj. Returns (Part)	1.12	1.49	1.91	1.60	0.92	51
Subj. Returns (Non)	-0.32	0.16	0.66	0.13	0.69	51
Average Difference	0.75	1.15	1.89	1.47	1.30	51
Treatment 3						
	25th pct	Median	75th pct	Mean	Std. Dev.	N
Subj. Returns	-0.71	0.32	1.34	0.44	3.31	1479
Realized Returns	-1.91	0.09	1.92	-0.02	3.54	704
Price	40.43	44.8	52.31	46.17	6.76	204
Participation	0.44	0.56	0.67	0.56	0.18	204
Subj. Returns (Part)	0.97	1.33	1.72	1.42	1.20	51
Subj. Returns (Non)	-0.69	-0.40	-0.06	-0.39	0.79	51
Average Difference	1.11	1.77	2.29	1.81	1.37	51
Treatment 4						
	25th pct	Median	75th pct	Mean	Std. Dev.	N
Subj. Returns	-1.62	-0.1	1.75	0.21	4.59	1632
Realized Returns	-2.59	0.25	2.63	-0.45	4.47	905
Price	21.56	26.21	31.96	27.44	8.42	204
Participation	0.44	0.67	0.78	0.62	0.19	204
Subj. Returns (Part)	0.53	1.14	1.47	1.06	1.12	51
Subj. Returns (Non)	-1.52	-0.99	-0.51	-0.89	1.42	51
Average Difference	1.31	1.99	2.78	1.96	1.95	51

Table 2.4: Summary Statistics by Treatment

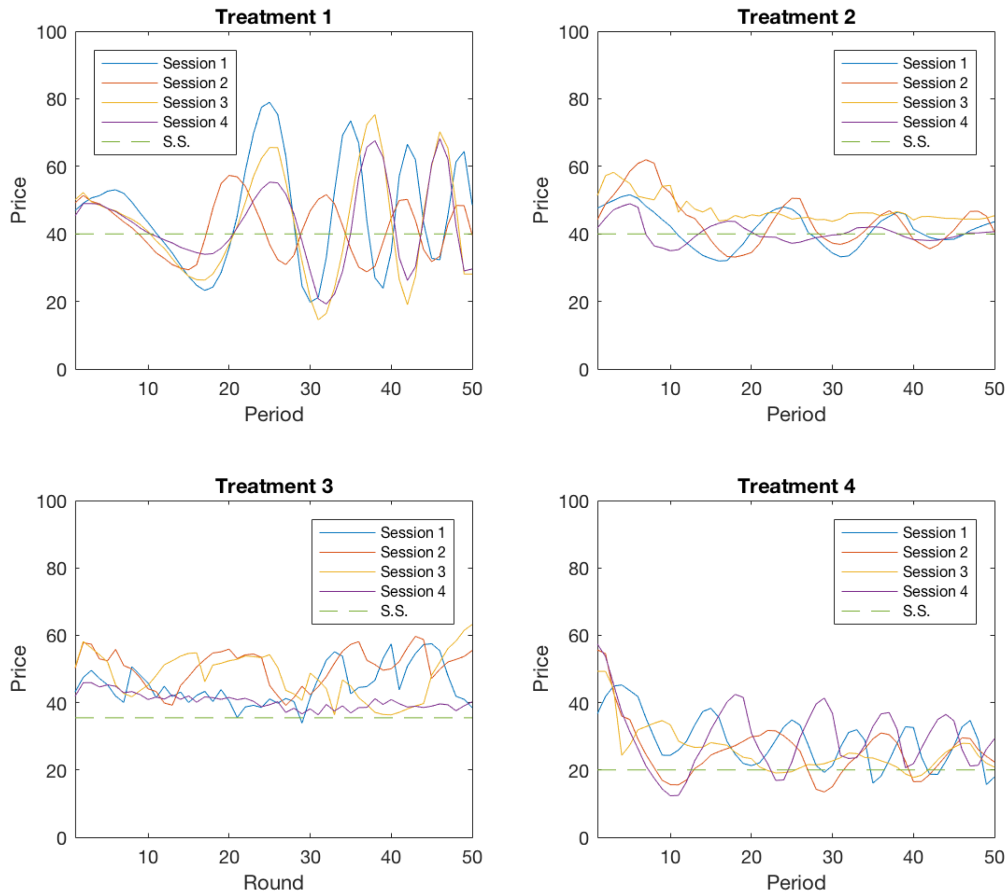

Figure 2.4: Aggregate Prices by Treatment.

2.5.1 Findings for Treatment-level Prices

Consistent with Hypothesis 1:

Finding 1: *For Treatments 2 and 4, higher cost treatments have lower mean prices and participation rates. Treatment 3 has higher mean prices and lower mean participation rates than Treatments 2 and 4. An extension of the model explains this theoretical deviation and is formalized in Finding 1b.*

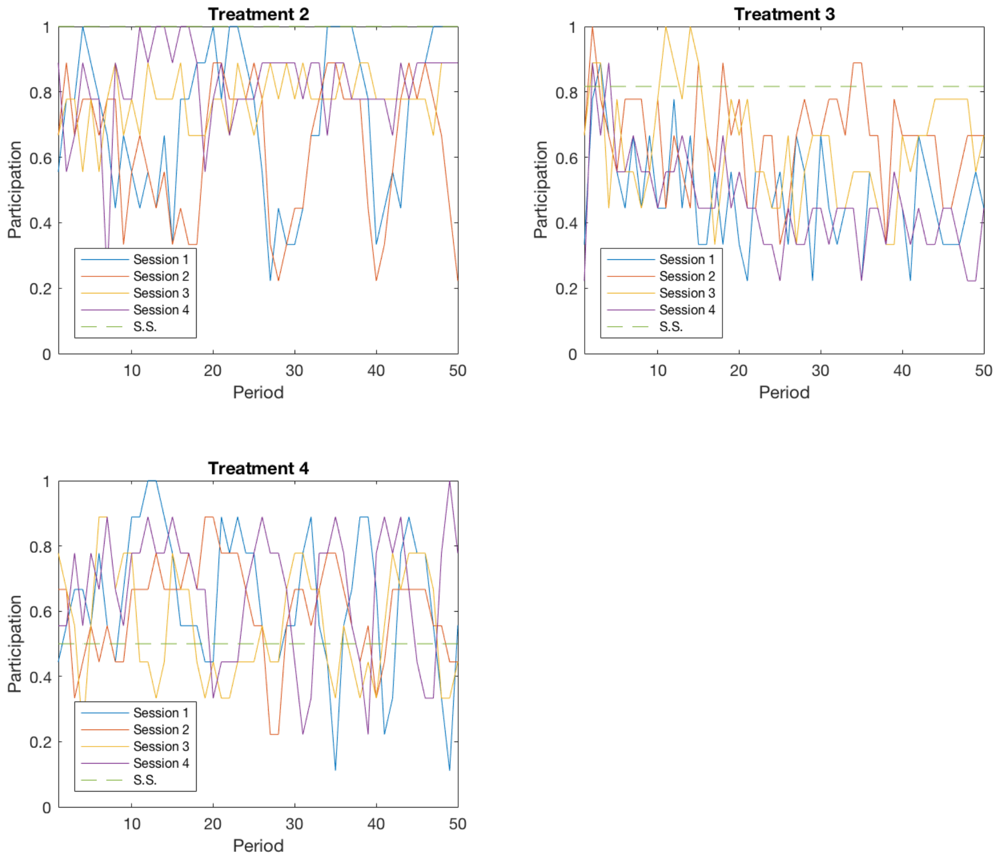

Figure 2.5: Aggregate Participation Rates by Treatment.

To potentially allow for learning over time, it is more accurate to use the mean values of the second half of each session. For robustness, I provide in the appendix the mean and median values for the entire series along with fitting an autoregressive process to each session to find the model implied unconditional means of each series. None of these alternatives have a qualitative effect on my results.

The model predicts Treatment 1 and 2 should have the same mean price and participation rates. Treatment 3 should have lower prices and participation rates than Treatment 2 and Treatment 4 should have lower prices and participation rates than Treatment 3. I summarize these predictions in the following equation:

$$\bar{p}^1 = \bar{p}^2 > \bar{p}^3 > \bar{p}^4 \quad \bar{N}^2 > \bar{N}^3 > \bar{N}^4$$

where $\bar{p}^i$ is the mean price and $\bar{N}^i$ is the mean participation rate for treatment i. In particular, the model predicts that higher cost treatments imply both lower mean prices and participation rates.

In order to formally test my model predictions, I use the Mann-Whitney test. The Mann-Whitney test is a nonparameteric rank sum test commonly used in the experimental literature to compare mean values across different treatments (Moffatt 2015). Table 2.5 summarizes the results from the Mann-Whitney tests.

Price					Participation				
Treatment	1	2	3	4	Treatment	1	2	3	4
2	$=$	-	$<$	$>^{**}$	2	-	-	-	-
3	$>$	$>$	-	$>^{**}$	3	$<^{*}$	-	$<$	-
4	$<^{**}$	$<^{*}$	$<^{**}$	-	4	$<^{*}$	$>$	-	-

Table 2.5: Mann-Whitney Tests for Price and Participation Rates

The asterisks represent the standard significance levels. I also summarize the results in the next equation:

$$\bar{p}^3 > \bar{p}^1 = \bar{p}^2 > \bar{p}^4 \quad \bar{N}^2 > \bar{N}^4 > \bar{N}^3$$

2.5.2 Explanation for Deviations in Treatment 3

All the treatments are in line with the theory except for Treatment 3. I provide an explanation for the deviation from the theory and then extend the benchmark model to explain

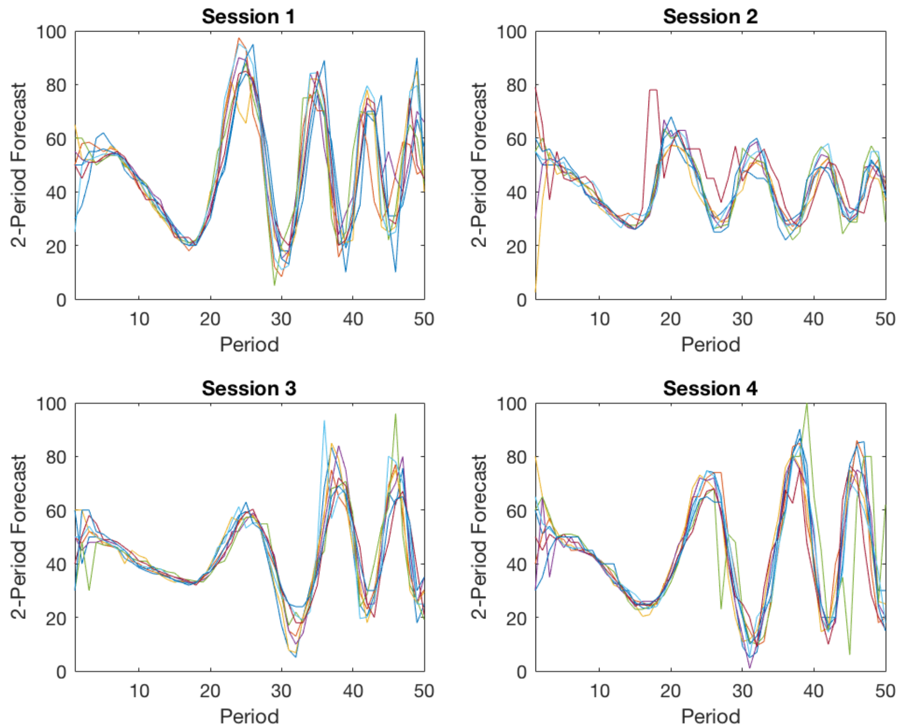

Figure 2.6: 2-period Ahead Forecasts for Treatment 1.

the results. In particular, each session in Treatment 3 has one subject with expectations consistently higher than the average. Since prices depend on average expectations, realized prices deviate from the model implied steady-state price. Figures 2.6 and 2.7 document 2-period ahead forecasts for Treatments 1 and 3.

Each graph provides the 2-period ahead forecasts for all 8 subjects within a session. Figure 2.6 depicts that in a standard session, 2-period ahead forecasts are highly correlated among subjects as in Hommes et al (2005). In contrast, Figure 2.7 shows that there is always one subject with 2-period ahead forecasts that are uncorrelated with the other subjects and consistently higher, in many cases hitting the upper bound. This leads to persistently higher prices, since prices are an average of every subject's 2-period ahead forecast. Moreover if the subject who is consistently providing high 2-period ahead forecasts is also participating in the market with a high induced cost, they can crowd out subjects with lower induced costs, leading to a lower participation rate.

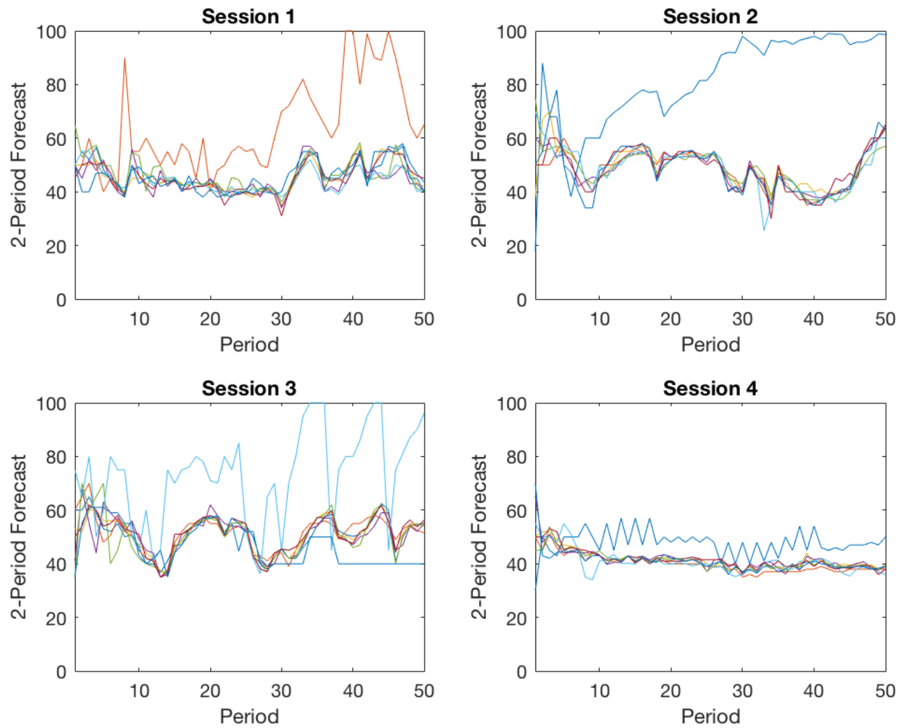

Figure 2.7: 2-period Ahead Forecasts for Treatment 3.

There are multiple potential hypotheses that can explain the behavior of these subjects. Instead of delving into the hypotheses of irrationality and high cognitive loads, I ask the inverse question: given the existence of subjects with high, uncorrelated forecasts, what is the implied behavior of the model? I define subjects with high 2-period ahead forecasts that are uncorrelated with other subjects' forecasts as an "exuberant" subject, loosely borrowing from Shiller (2000).

Finding 1b: *The results in Treatment 3 can be explained by an extension of the benchmark model. I find evidence that "exuberant" subjects crowd-out lower cost subjects from the experiment, leading to higher mean prices and lower mean participation rates than predicted by the benchmark model.*

I extend the model to include an agent with exuberant expectations $\bar{p}^x$ that are taken

Treatment	Session	Mean Price	S.S. Price	Mean Participation	S.S. Participation
3	9	45.13	47.16	48.15%	42.79%
3	10	48.37	51.39	63.62%	28.71%
3	11	50.77	49.56	65.58%	34.8%
3	12	40.42	41.16	45.32%	62.79%

Table 2.6: Calibration for Treatment 3

as exogenous. I derive the extended model in Appendix B. The steady-state price and participation rates are:

$$\bar{p} = \frac{R^{-1}[(1-h)\bar{p}^x + \mu - \frac{S}{N}]}{1 - R^{-1}h}$$

$$\bar{N} = \frac{C + \sqrt{C^2 - 4\bar{k}S(R-1)}}{2\bar{k}}$$

where $1-h$ is the fraction of agents with exuberant expectations, and C is a function of the model parameters. I then calibrate the model by taking the mean expectations of "exuberant" subjects in the experiment and choose the weight $(1-h)$ to be $\frac{1}{9}$ which corresponds to the the number of subjects in the experiment plus an automated fund. Table 2.6 lists the predictions of the extended model with the calibrated values from the experiment.

2.5.3 Interpretation of Treatment-level Findings

Finding 1c: *While higher induced costs k^i lowers the average frequency of positive subjective returns, participation costs alone cannot account for differences in subjective expected returns at the session level.*

As Findings 1a and 1b demonstrate, while the model generally does well in the aggregate, subjects are not participating in the stock market solely based on their induced costs. That is, subjects with a higher induced cost do not necessarily have lower subjective returns than subjects with lower costs. The following graphs show that mean participation rates cannot

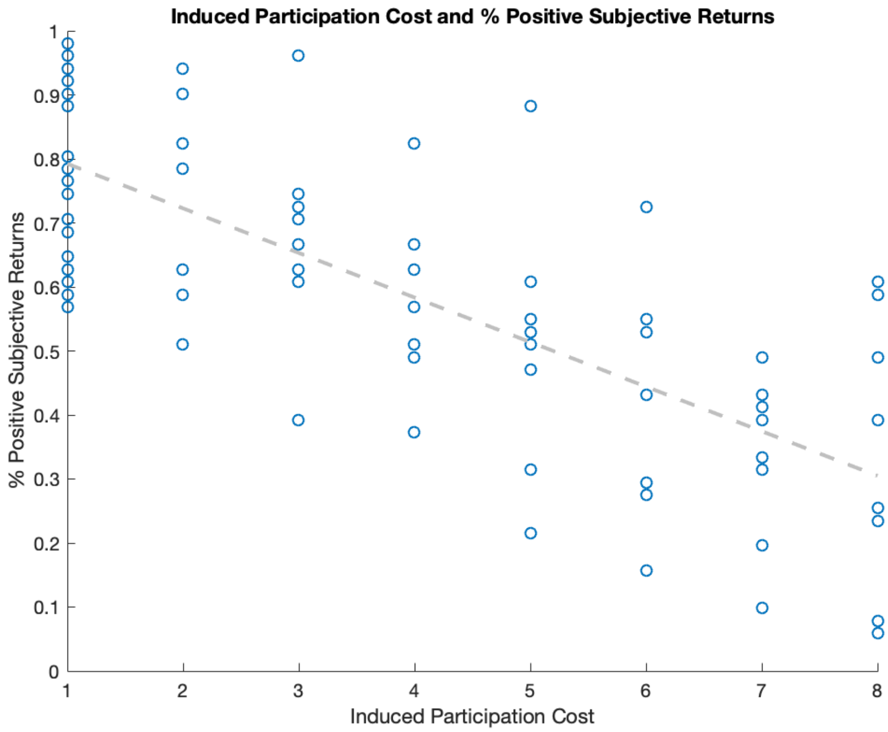

Figure 2.8: Aggregate Induced Cost and Positive Subjective Returns.

be solely explained by differences in induced participation costs. Figure 2.8 shows the percentage of periods that a subject has positive subjective returns and their induced costs. I rank induced costs from 1 to 8 with 1 being the lowest and 8 being the highest to capture session and treatment level heterogeneity: Figure 2.8 shows that on average, induced costs do well at explaining subjective returns. While there is large heterogeneity across sessions, on average, higher induced costs lowers the percentage rate of positive subjective returns. Figure 2.9 decomposes the data into Treatments 3 and 4 and shows that while on average, induced costs explain subjective returns, within a session, participation costs need not align with subjective returns. In particular, subjects with high induced costs can have higher subjective returns and subjects with low induced costs can have lower subjective returns, independent of participation costs.

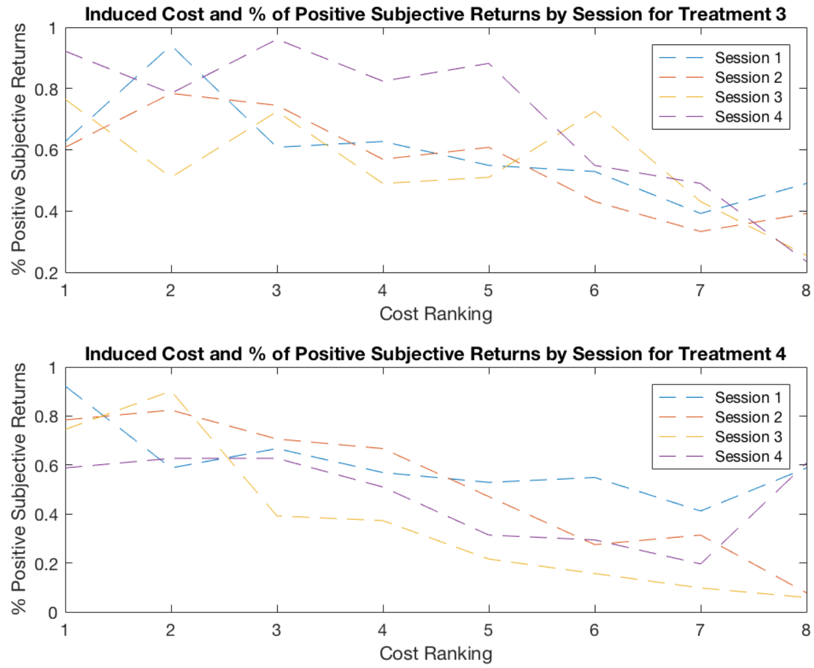

Figure 2.9: Induced Cost and Positive Subjective Returns by Treatment.

2.5.4 Findings for Individual Participation Decisions

Finding 1c suggests that there are determinants for the participation decision independent of participation costs. I now analyze the individual participation decision.

Finding 2: *The probability of participating in the stock market is increasing in subjective expected returns, lagged realized returns, lagged forecasting payoffs, and a price trend.*

In order to understand the participation decision at the individual level, I run a random-effects probit panel data regression with robust standard errors. Random-effects regressions are preferred over fixed-effects under the probit model because coefficient estimates under fixed-effects are biased. In the appendix, I run robustness checks with logit and fixed-effects linear probability models and find that the results are similiar. I run the regression using period, session, and treatment level dummies to capture potential dependence at the period, session, and treatment levels.

The baseline random-effects probit regression is of the form:

$$P(n_{it} = 1|x_{it}, z_i) = \Phi(\alpha + \beta' x_{it} + \gamma' z_i + u_i)$$

where $n_{it} = 1$ is participation, $\Phi(\cdot)$ is the CDF normal, i is the subject, t is the period, x_{it} contain variables that vary both between subjects and periods, z_i contains variables that vary among subjects, β and γ are the regression coefficients, α is the constant term, and u_i is the subject specific term where $V(u_i) = \sigma_u^2$.

I run three specifications. Specification 1 is the simplest regression which just includes subjective expected returns. For Specification 2, I add regressors using the guidance of theory, in particular, a price-trend, risk aversion, and lagged payoffs. The participation and forecasting payoffs are lagged because it is the last payoff that is in the subjects' information set. Then, Specification 3 adds a dummy variable for past experienced payoffs. Table 2.7 shows the results of the regression. I add demographic controls and lags in Appendix B.

2.5.5 Benchmark Regression

Specification 2 is the benchmark regression. I find that subjective returns, lagged realized returns, lagged forecasting payoff, and price trends matter for stock market participation. The economic interpretation is as follows. Conditioning on induced costs, a 1 franc increase in subjective returns leads to a 2.8% increase in the likelihood of participating. Next, a 1 franc increase in lagged realized returns increases the likelihood of participating by 1.8%. A 1 franc increase in past forecasting payoffs increases the likelihood of participating by 1.9%. Finally, subjects place weight on price trends, that is a 1 franc increase in the price trend increases the likelihood of participating by 1.3%.

Table 2.7: Dependent Variable: Individual Participation

Variable	Model 1	Model 2	Model 3
Subjective Returns$_t$	0.03***	0.028***	0.027***
		(0.009)	(0.009)
Actual Returns$_{t-1}$		0.018***	0.004
		(0.005)	(0.006)
Forecast Payoff$_{t-1}$		0.019***	0.01*
		(0.006)	(0.006)
Risk Aversion		0.039	0.035
		(0.026)	(0.023)
Price Trend$_{t-1}$		0.013***	0.01**
		(0.005)	(0.004)
Past Positive Payoff$_t$			0.181***
			(0.023)
Past Negative Payoff$_t$			0.003
			(0.03)
N	4896	4704	4704
Pseudo R^2	0.048	0.157	0.172

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

2.5.6 Past Experienced Payoffs

The fact that participation depends on lagged actual returns suggests that participation may depend on past experiences. To test this hypothesis, I create a dummy variables which splits the dataset into 3 parts. In particular, subjects can receive 3 categories of payments depending on their participation decision. If the subject does not participate, then they receive the risk-free payoff. If the subject participates, then they receive either a high payoff or a low payoff depending on realized returns. Surprisingly, I find that if a subject receives a high payoff from participating, they are 18.1% more likely to participate in the stock market. In contrast, when subjects receive a low payoff, it has no significant effect on the likelihood of participation. The price-trend controls for the fact that past high payoffs may signal future high payoffs i.e. if subjects are in a rising price-trend environment. I find that even controlling for the trending price, subjects are more likely to participate given higher past returns. The fact that the effect is not symmetric suggests that there is some bias in subjects' updating rule.

2.5.7 Interpretation

In line with past empirical evidence, the regression results provide strong econometric evidence that higher subjective returns lead to higher probability of participating in the stock market. More importantly in contrast to the benchmark model, the results show that high past realized returns increases the likelihood of participation. The fact that the coefficient on past positive payoffs are highly positive implies that there is some inertia to participation when subjects receive a high payoff. Nevertheless, since subjective returns are also a function of past realized returns, in order to distinguish the impact of realized returns, that is, whether past positive returns lead to higher subjective returns or if past negative returns lead to lower subjective returns, I look directly at subjects' expectations updating behavior.

2.5.8 Learning from Experience

Finding 3: *Subjects who participate in the stock market in the prior period and receive a low payoff, have lower subjective expected returns. In particular, they place less weight on the price trend.*

Finding 2 demonstrates that subjective expected returns are the main driver of individual participation, while higher lagged realized returns increases the probability of participation. Moreover, Table 2.4 demonstrates that the subjective returns of participants are systematically higher than non-participants. My experimental design allows a novel look at how subjects update their expectations. In particular I find that subjects who experience a low payoff place a lower weight on the forecast trend and thus have a lower subjective return.

To formally test my hypothesis, I run 3 different regressions. I run a regression on subjects' 1-period ahead forecast, 2-period ahead forecast, and forecast trend i.e. $p^e_{t+1} - p^e_t$ on lagged prices and participation along with a dummy for past realized payoffs. The first result is that subjects update their 1-period ahead forecast towards past prices independent of past realized payoffs. That is, subjects update towards the signal (Chambers and Healy 2012). Similarly subjects weigh the 2-period ahead forecast toward past prices but using different weights. A 1 franc increase in lagged prices leads to a 1.3 franc increase in the 2-period ahead forecast while a 1 franc increase in the 2-period lagged price decrease the 2-period ahead forecast by 0.37 francs. Interestingly, subjects' past experiences only have a statistically significant impact on the 2-period ahead forecast, that is, past positive payoffs leads to higher 2-period ahead forecasts and past negative payoffs lead to lower 2-period ahead forecasts. If subjects' experience a low payoff last period, then they lower their 2-period ahead forecast by -0.53 while if they experience a high payoff last period, they increase their 2-period ahead forecast by 0.43.

Table 2.8: Dependent Variable: 1-Period, 2-Period, Forecast Trend

Variable	**1-Period Forecast**	**2-Period Forecast**	**Forecast Trend**
Price_{t-1}	1.24***	1.323***	
	(0.045)	(0.063)	
Price_{t-2}	-0.259***	-0.374***	
	(0.042)	(0.0618)	
$\text{Participation}_{t-1}$	0.502***	0.962***	
	(0.165)	(0.248)	
Price Trend_{t-1}			0.0264
			(0.219)
$\text{Past Positive Payoff}_t$	0.224	0.428*	-0.142
	(0.224)	(0.217)	(0.147)
$\text{Past Negative Payoff}_t$	-0.183	-0.525*	-0.772***
	(0.183)	(0.295)	(0.2)
N	4508	4508	4508
R^2	0.945	0.883	0.025

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

2.5.9 Interpretation

The underlying story is that differences in subjective returns due to differences in experience, can lead to implicit costs to stock market participation. Since the standard framework assumes that participation costs are utility costs, high participation costs are equivalent to having a lower expected utility which depends on subjective expected returns. My results demonstrate that differences in experiences can create "pseudo-costs" to stock market participation that differ from standard participation costs. Moreover, since differences in subjective returns are directly measurable, I provide a testable story for explaining limited participation among non-liquidity constrained households. The underlying mechanism is behavioral in that losses hurt more than gains (Kahneman and Tversky 1979). In essence, confronting an agent that "learns from experience", they will tell you that they expect lower returns since they have directly experienced them.

2.5.10 Discouraged Investors

Finding 4: *A fraction of subjects exit the market for the majority of the experiment after consecutive low payoffs.*

Additionally, the experiment provides a unique environment to identify novel behavioral phenomena that are difficult to elicit from survey data. I find evidence of subjects who exit the market after consecutive low payoffs from participating in the stock market. I call these subjects "discouraged investors". To formally identify discouraged investors, I define a variable called failure rates, where if the subject participated in the stock market in period $t-1$ and received a payoff lower than the risk-free payoff, then it is considered a failure. In particular, there are 10 subjects out of 96 that fit this criterion. I find that discouraged investors end the experiment with failure rates close to 100% since they no longer participate for the rest of the experiment.

This can potentially be rationalized by an ambiguity aversion or robust control model where subjects make robust choices against the worst-case scenario model. This finding is in line with Strahilevitz et al (2011) where investors are reluctant to return to the stock market once they realized low payoffs. Thus I provide experimental evidence for a novel behavioral financial phenomena.

2.6 Conclusion

My experimental results demonstrate that a model of heterogeneous participation costs with rational expectations (RE) alone cannot explain limited participation of the non-liquidity constrained. Moreover, dispersion within subjective expected returns can be due to "learning from experience" where subjects over-weigh public signals with respect to realized outcomes i.e. higher and lower returns.

My experiment provides three answers that are novel to the literature on heterogenous expectations and asset pricing. First I show that heterogeneity in subjective expected returns along with non-rational expectations are needed to explain limited participation among the non-liquidity constrained in an experimental context. Second I provide direct evidence for the conjecture in Malmendier and Nagel (2011) that "true experiences" are the determinants of "learning from experience", that is, it is the subjects' experiences i.e. low versus high realized returns that determine whether they bias their updates or not, not just the history. Finally, I provide strong evidence for the claim in Dominitz and Manski (2011), that heterogeneity in subjective expectations are due to weighing the public signal differently.

While my paper answers some key questions, others remain. For instance, what determines the initial beliefs of subjects? Next, while I showed that heterogeneity can be generated by the same learning rule which weighs outcomes differently, there is a large literature following Hommes et al (2005) that explains dispersion in beliefs through intrinsic heterogeneity, that is, where agents use different learning rules. Future research should look at differentiating between these two modeling approaches. For simplicity, my experiment provides a stationary environment with no private information and common knowledge. A natural extension is to relax each component in turn to see the impact of richer environments on both heterogenous subjective returns and outcomes. Moreover, since investors tend to have different planning horizons, it would be interesting to see how multiple horizons will affect the results.

Finally, my experiment provides a novel extension to the standard LtF design which allows for a participation decision. Many expectations-based models can be augmented to include an extensive margin decision and tested in the laboratory such as the cobweb model with firm entry.

Chapter 3

Learning, Hypothesis Testing, and Restricted-Perceptions Equilibria

3.1 Introduction

While the rational expectations hypothesis is both parsimonious and elegant, it requires demanding assumptions on part of the agents. In response, the adaptive learning literature has moved towards relaxing the hypothesis using econometric learning (Evans and Honkapohja 2001). With econometric learning, agents act as econometricians and must estimate the model parameters over time. A criticism of this approach is the passive nature of the learning mechanism and that the rational expectations equilibria (REE) of a model are the only possible points of convergence. Similarly, deviations from this approach such as restricted perceptions, where agents underparameterize their models, have faced criticism because they do not nest the REE and hence do not give the rational expectations solution a chance. Contrary to agents in learning models, modern econometricians also test the specification of their models and should be able to detect misspecification over time.

More recently, Cho and Kasa (2014) extend the econometric learning model to a setting where a policy maker suspects that their model is misspecified and also tests the specification of their model over time. I extend upon their framework and allow multiple agents to test the specification of their model in a simple asset pricing framework. In my setting, groups of agents are endowed with either the fully-specified model or an underparameterized one i.e. have restricted perceptions. Agents suspect that their model is misspecified and test the specification with new data. If it passes, the model is updated using least-squares learning. If it fails, then the agent draws a new model from a set of models. Following the theme of econometric terminology, I call this procedure hypothesis testing learning.[1] The main question is whether other equilibria naturally arise besides the REE, under hypothesis testing learning.

I find that in a simple asset pricing model, where agents engage in hypothesis testing learning, restricted perceptions equilibria (RPE), that is, the equilibrium points of underparameterized models naturally arise. My environment calls for a new stability definition which I call HT-stability or Hypothesis Testing stability. For an equilibrium to be HT-stable, it needs to be robust to shocks, which I define here to be a model shock. Similar to the idea behind impulse-response functions, I allow the model to converge to a steady-state and force one agent type to switch their model i.e. a model shock. HT-stability then requires that the system return to its original steady-state equilibrium values. Under this definition and reasonable parameters, the REE is not HT-stable and in particular only RPE corresponding to combinations of underparameterized models are HT-stable. In particular, no agents hold the fully parameterized model in HT-stable equilibria.

This finding is similar to those in Sargent (1999) and Cho and Kasa (2014) which relate the strength of the self-referential feedback to the resilience of certain equilibria compared to others. In our case, there are two reasons why the REE fails to be HT-stable. First like in Cho and Kasa (2014), the self-referential feedback of the REE is not as strong as the other

[1] Cho and Kasa (2014) call this "model validation" but the original idea from Foster and Young (2003) uses the term hypothesis testing.

equilibria when the other agents hold underparameterized models. Next, the REE in our base case has three different parameters which are tested and hence is much easier to reject in environments where the other agents have misspecified models. In particular, with three parameters, the agent is easily able to detect shifts in the model parameters and likely to discard their model.

Finally, given that the model has multiplicity of RPE, I implement constant-gain learning to characterize the time agents spend within each RPE. I find that the model spends most of its time within a subset of the RPE and in particular mainly between a subset of them. Surprisingly, the model spends most of its time in the hybrid equilibria where all models are used. Using the language of Cho and Kasa (2014), the "dominant model" i.e., where all agents hold one model most frequently, is not the fully-specified model but rather the dividend only model. The key implication of my analysis is that agents can have misspecified models arising naturally from a more realistic learning process over time.

3.2 Literature Review

This paper contributes to the literature on adaptive learning and more specifically to the budding literature which examines learning mechanisms where agents are endowed with multiple models. I also contribute to the literature dealing with alternative equilibrium concepts to rational expectations such as the restricted-perceptions equilibrium (RPE).

The first paper to propose hypothesis testing as a learning mechanism is Foster and Young (2003) who have agents test the specification of their models in a stationary game-theoretic environment. They find that this type of learning converges to a solution that is approximately a Nash Equilibrium. Next, Branch and Evans (2007) deals with model uncertainty and its effects on volatility. This paper extends Cho and Kasa (2014) who examines a policy maker's decision-making and dynamics under model validation. My paper differs from theirs in two key aspects. First, my approach focuses on the stability of equilibria in a class of models using the techniques introduced in Cho and Kasa (2014). Next, I allow multiple agents to engage in model validation rather than just one and specifically test which models are "dominant" in the long-run via simulation. More recently, Norman (2015) introduces hypothesis testing learning to a macroeconomic setting without adaptive learning and Cho and Kasa (2017) introduce a sequential Lagrange multiplier (LM) test into a model with Bayesian averaging.

Next, in the literature on RPE, Branch and Evans (2010, 2011) also deal with RPE in an asset pricing framework. The closest paper in this strand of the literature is Branch and Evans (2010) who deal with an RPE switching framework in an asset pricing model. This paper is different from theres in that I allow fully-specified models to be chosen by agents through hypothesis testing. As far I know, this paper is the first to extend specification tests to multiple agents in an asset pricing model and test the model properties through numerical simulation.

3.3 Model

Time is measured in discrete periods. There is a finite set of models M indexed by $m = 1, 2, ..., M$. For tractability, agents assume that all models are linear in the exogenous variables. There is a unit measure of agents divided into m partitions who are each endowed with an initial model indexed by $m \in M$ at time 0. Hence an agent i with model m will be indexed i_m. The fraction of agents with initial model m is n_m where $\sum_m n_m = 1$. All agents are infinitely lived and have CARA utility functions of the form: $u(c) = -e^{-ac}$, where $a > 0$ is the coefficient of absolute risk aversion. Agents also maximize their one-period portfolio.[2] There are two assets traded in perfectly competitive markets: a risky Lucas tree

[2]Alternatively, one can employ an OLG structure which is equivalent to an infinitely lived agent maximizing their one-period portfolio.

and a riskless one-period bond. The risky asset yields a dividend stream $\{D_t\}_{t=0}^{\infty}$ and sells at price p_t. The riskless one-period bond gives an exogenous gross return $R > 1$ and the supply is infinitely elastic.

Dividends follow a stationary AR(1) process:

$$\tilde{D}_t = (1-\rho)\mu + \rho\tilde{D}_{t-1} + \epsilon_t^D$$

where $\rho \in [0,1)$ is the AR(1) coefficient, $\mu > 0$, and ϵ_t^D is white noise with distribution $N(0,\sigma_D^2)$. Similarly, supply follows a stationary AR(1) process:

$$\tilde{S}_t = (1-\phi)S + \phi\tilde{S}_{t-1} + \epsilon_t^S$$

where $\phi \in [0,1)$ is the AR(1) coefficient, $S > 0$, and ϵ_t^S is white noise with distribution $N(0,\sigma_S^2)$. The supply shocks are also possibly correlated with the dividend process such that the covariance $\sigma_{DS} \neq 0$. The stochastic supply is a proxy for volatility in asset float where firms create new issues, provide options and warrants that are periodically exercised and change the available supply at a given time. The importance of asset float is well documented in the literature (Baker and Wurgler 2000). I also denote the random variables with tildes to differentiate them from their mean deviation forms: $D_t = \tilde{D}_t - \mu$ and $S_t = \tilde{S}_t - S$ which will prove useful when deriving the restricted perceptions equilibria (RPE).

Then the pricing equation is:

$$p_t = R^{-1}[E_t^*(p_{t+1} + \tilde{D}_{t+1}) - \tilde{S}_t a\sigma_p^2]$$

where $\sigma_p^2 \equiv Var_t(p_{t+1} + D_{t+1})$ and E_t^* is the (potentially) non-rational conditional expectation at time t. With rational expectations, $E_t^* = E_t$. If agents instead deviate from rational expectations or use misspecified models, then $E_t^* = \frac{1}{M}\sum_m E_t^m$, which is the weighted average of all expectations in the economy conditional on their model, where E_t^m is the conditional expectation at time t with respect to model m.

3.3.1 Base Case

A key assumption here is that some agents have underparameterized models, which is called restricted perceptions. I motivate restricted perceptions by stating that agents face a degree of freedoms problem. If there are multiple factors that may influence the dividend or supply process, it will be difficult for the agent to estimate that process, which may cause agents to omit some variables. This specification is similar to Branch and Evans (2010) with the difference that I allow one of the models to be correctly specified.

I implement a simple case of the model with a minimal number of elements in the model class. A natural benchmark case includes only two underparameterized models where agents omit one of the two exogenous variables S and D. Hence there are 3 models in the model set M, in particular, a model where dividends D are omitted and a model where supply S are omitted, along with the correctly specified model. The following models are:

$$p_t = A_1 + B_1 D_t + C_1 S_t + \nu_{1t}$$

$$p_t = A_2 + B_2 D_t + \nu_{2t}$$

$$p_t = A_3 + C_3 S_t + \nu_{3t}$$

where:

$$E_t^1 p_{t+1} = A_1 + B_1 \rho D_t + C_1 \phi S_t$$

$$E_t^2 p_{t+1} = A_2 + B_2 \rho D_t$$

$$E_t^3 p_{t+1} = A_3 + C_3 \phi S_t$$

Since there are three potential models, each agent can hold one of the three models. Let n_1 denote the fraction of agents with model 1, n_2 denote the fraction of agents with model 2, and $1 - n_1 - n_2$ denote the fraction of agents with model 3. To simplify the analysis, I have this fraction be $\frac{1}{3}$ for all models. In discussions, model 1 will be called the fully-specified model, model 2 will be the dividend model, and model 3 will be the supply model. In this setup there are potentially two groups of agents in the economy that underparameterize their forecasting model by either omitting dividends or supply and one group that has the fully-specified model

In an econometric learning model where switching is determined by forecasting performance, the REE is the asymptotic solution. I do not have agents determine which models to use based on forecasting performance here for two reasons. First when there are structural changes in the economy like in our setting, choosing models based on past forecasting performance is not necessarily robust. Second in an environment where agents may not know the entire set of models, agents will be unable to rank forecasting performance, that is, I assume that agents do not know the entire set of models and hence cannot judge if their model is better than another one. This interpretation is borrowed from Foster and Young (2003) who view model selection as formalizing the notion of agents developing "hunches" of how the economy works over time. Hence, in our setting with hypothesis testing learning, whether the REE will be the long-run solution is not entirely obvious.

3.3.2 Equilibrium

The pricing equation in the base case then becomes:

$$p_t = R^{-1}\left[\frac{1}{3}(E_t^1 p_{t+1} + E_t^2 p_{t+1} + E_t^3 p_{t+1}) + E_t D_{t+1} + \mu - (S_t + S)a\sigma_p^2\right]$$

where E_t^m is the subjective expectation of the agent for model m and at this point I replace the exogenous processes $\tilde{D}, \tilde{S}$ with their mean deviation forms D, S. Plugging in for the expectations I get:

$$p_t = \xi_0 + \xi_1 D_t + \xi_2 S_t$$
$$\xi_0 = R^{-1}[\frac{1}{3}(A_1 + A_2 + A_3) + \mu - Sa\sigma_p^2]$$
$$\xi_1 = R^{-1}[\frac{1}{3}\rho(B_1 + B_2) + \rho]$$
$$\xi_2 = R^{-1}[\frac{1}{3}\phi(C_1 + C_3) - a\sigma_p^2]$$

To reiterate, agents with underparameterized models are said to have restricted perceptions. Two key objects in the learning literature are the perceived law of motion (PLM) and the actual law of motion (ALM). A PLM is the model that the agents believe are the true data-generating process while an ALM is the actual model that determines the economy. One facet of restricted perceptions is that the perceived laws of motion (PLM) of the restricted perceptions models do not nest an REE and thus it is impossible for them to converge to it. Nevertheless, models with restricted perceptions can converge to an RPE which is optimal within a limited class of PLMs.

Although the restricted perceptions models are underparameterized, I will require them to forecast in a statistically optimal manner and that the model parameters be optimal linear projections. Intuitively, when forecasts satisfy orthogonality conditions, the agent will be forecasting optimally within their restricted perceptions and will be acting optimally within their model, that is, in an RPE, agents will not be able to detect that their models are misspecified. The orthogonality conditions for the 3 models are thus,

$$E(1, D_t, S_t)'(\xi_0 + \xi_1 D_t + \xi_2 S_t - A_1 - B_1 D_t - C_1 S_t) = 0$$
$$E(1, D_t)'(\xi_0 + \xi_1 D_t + \xi_2 S_t - A_2 - B_2 D_t) = 0$$
$$E(1, S_t)'(\xi_0 + \xi_1 D_t + \xi_2 S_t - A_3 - C_3 S_t) = 0$$

where:

$$A_j = \xi_0, j = 1, 2, 3$$

$$B_1 = \xi_1$$

$$B_2 = \xi_1 + \xi_2 r$$

$$C_1 = \xi_2$$

$$C_3 = \xi_2 + \xi_1 \tilde{r}$$

where $r = \frac{ED_t S_t}{ED_t^2}$ and $\tilde{r} = \frac{ED_t S_t}{ES_t^2}$. Then an RPE is a set of coefficients $\{\xi_0, \xi_1, \xi_2, A_1, A_2, A_3, B_1, B_2, C_1, C_3\}$ that solves these systems of equations. The coefficients become rather involved so we leave them for Appendix C.

3.3.3 Adaptive Learning

Agents do not know the model parameters and must estimate them over time. Here, I deviate from rational expectations and implement an adaptive learning rule for two reasons.[3] Rational expectations requires a full understanding of the model as well as beliefs of other agents. In this sense it is a Nash equilibrium, hence coordination between agents requires strong cognitive and informational assumptions. Instead, many applied econometricians estimate econometric forecasting models and adjust the coefficients in light of new data. Here I adhere to the Cognitive Consistency Principle (Sargent 1993) which requires agents and econometricians to be on equal footing.

Next in an environment such as this where agents possibly entertain multiple models, it may take time for a new model to converge to its RPE. I allow agents to update their parameters over time to allow each model to adapt to the environment. With adaptive learning, parameters are updated with respect to observable data over time. Hence, a misspecified model can potentially still converge to some equilibrium value because the model parameters can adjust such that it fulfills the orthogonality conditions as above.

[3] The current formulation can be seen as a midpoint between RE and restricted-perceptions.

Agents in the model update their expectations conditional on data. Let $z_t^{i'}$ be the data vector depending on the agent's model which can include some combination of 1, D_t, and S_t. Then let Φ_t^i be the vector of coefficients which include some combination of A_t^i, B_t^i, and C_t^i depending on the model. Agents update their pricing coefficients Φ_t^i as follows:

$$\Phi_t^i = \Phi_{t-1}^i + \gamma_1 (R_{t-1}^i)^{-1} z_t^i [p_t - \Phi_{t-1}^{i'} z_t^i]$$
$$R_t^i = R_{t-1}^i + \gamma_1 [z_t^i z_t^{i'} - R_{t-1}^i]$$

where γ_1 is the gain parameter and R_t^i is the moment matrix for agent i with some model.

3.4 Hypothesis Testing

In addition to updating their models over time, agents also test the specification of their model. Modern econometricians use statistical tools to test for model misspecification and update their model specification accordingly. Following the analogy of Cho and Kasa (2014), one inconsistency of the current learning approach is that we assume agents have taken the first semester of econometric theory but not the second. That is, the agents understand estimation but not inference.

We would expect such a procedure to approach the true model, but endogenous data and misspecification on part of the other agents creates issues. As stated above, I call this procedure hypothesis testing learning as in Foster and Young (2003). There are potentially many ways to test the specification of a model. Following Cho and Kasa (2014), I endow agents with a Lagrange multiplier (LM) test. The Lagrange multiplier (LM) or score test, uses the score and the Fisher information to test the sensitivity of the likelihood function. The LM test is a natural choice as a hypothesis testing mechanism because it does not require that one explicitly specify an alternative hypothesis. Hence it is usually known as a misspecification test. Moreover, the LM test has a recursive formulation which allows it to be naturally added to the tools of stochastic recursive algorithms of the adaptive learning literature.

The LM test is defined as follows:

$$\Lambda_t^i = (R_{t-1}^i)^{-1} z_t^i [p_t - \Phi_{t-1}^{i'} z_t^i]$$

where the null hypothesis is $H_0 : \Lambda_t^{i'} (\Omega_t^i)^{-1} \Lambda_t^i \leq \tau_i$ and where Ω_t^i is the variance of the score Λ_t^i and defined as:

$$\Omega_t^i = \Omega_{t-1}^i + \gamma_2 (\Lambda_t^i \Lambda_t^{i'} - \Omega_{t-1}^i)$$

Agents update the LM test-statistic as follows:

$$\theta_t^i = \theta_{t-1}^i + \gamma_2 (\Lambda_t^i (\Omega_t^i)^{-1} \Lambda_t^i - \theta_{t-1}^i)$$

Here τ_i is a threshold value that determines the test-statistic where agents would reject their model. If the score statistic θ_t^i is less than τ_i then the agent keeps their model, else they reject it and draw another model. I follow Foster and Young (2003) where agents draw a new model at random from the set of models. Following their definition, I call this procedure experimentation, where agents are unaware of the different elements in the model class and that drawing a model at random is equivalent to having a "hunch". In the case of 3 models, the probability of drawing a model m is $\frac{1}{3}$ after rejecting their existing model. For now I allow the agent to redraw their model mainly for computational tractability. After a new model is drawn, there is also a grace period ξ between drawing a new model and testing it to give each model's parameters a chance to adjust to the environment and thus not have it potentially be immediately rejected upon being drawn.

Once I implement hypothesis testing, the potential space of equilibria increases. With this environment there are ten equilibria that correspond to the combination of the three different model types. One of them is the REE while the other nine are RPE. Unfortunately with multiple switching, learning, and hypothesis testing I am unable to derive analytical results due to the complex nonlinearities that arise from the learning algorithms. Instead, I use numerical analysis and simulations to characterize the local stability and long-run properties of the different equilibria.

3.5 Numerical Analysis

It is illustrative to describe the numerical algorithm used to simulate and analyze the model. I first write down the dynamical system by gathering the equations from the previous sections. Then I describe the algorithm and timing to provide intuition.

Let $s_t = [n_1,\ n_2,\ n_3]$ where n_i for $i = 1, 2, 3$ denotes the fraction of the agents in the economy that can each be endowed with some model $m = 1, 2, 3$. While this leads to discrete model switching, given that all agents of type n_i switch once their model is rejected, it leads to computationally tractable simulations.[4] Then, s_t denotes the set of models agents have at time t.

I first write down the set of model updating equations:

$$\Phi_t^i = \Phi_{t-1}^i + \gamma_1 \Lambda_t^i$$

$$\Lambda_t^i = (R_{t-1}^i)^{-1} z_t^i [p_t - \Phi_{t-1}^{i'} z_t^i]$$

$$R_t^i = R_{t-1}^i + \gamma_1 [z_t^i z_t^{i'} - R_{t-1}^i]$$

where Φ_t^i is the vector of belief parameters, Λ_t^i is the score for the model, and R_t^i is the moment matrix all for agent i. Through the feedback, these determine the pricing equation or ALM:

$$p_t = R^{-1}[h(s_{t-1}, \Phi_{t-1}^i) + E_t D_{t+1} + \mu - (S_t + S) a \sigma_p^2]$$

[4]Currently this research project is still in an early stage where even running the algorithm is computationally involved. Future work will deal with less discrete updating by using distributions of test thresholds.

where $h(s_{t-1}, \Phi^i_{t-1})$ is the aggregate expectations of the economy subject to beliefs Φ^i_{t-1} and model set s_{t-1}. Notice here, that the pricing equation is also dependent on s_{t-1} which is an index for the number of model types in the economy. The important aspect of hypothesis testing is that s_{t-1} is also a function of the realizations p_{t-1} and test parameters. Next, models are tested by forming the recursive LM test-statistics:

$$\theta^i_t = \theta^i_{t-1} + \gamma_2(\Lambda^i_t(\Omega^i_t)^{-1}\Lambda^i_t - \theta^i_{t-1})$$
$$\Omega^i_t = \Omega^i_{t-1} + \gamma_2(\Lambda^i_t\Lambda^{i'}_t - \Omega^i_{t-1})$$

where the null hypothesis is $H_0 : \Lambda^{i'}_t(\Omega^i_t)^{-1}\Lambda^i_t \leq \tau_i$. And finally, the set of models is determined by:

$$s_t = f(s_{t-1}, p_t, \tau_i, \xi)$$

where s_t depends on the ALM p_t, the test thresholds τ_i, and the grace period ξ. Because the law of motion for s_t is a complex nonlinear equation that depends on the rest of the dynamic system, I rely on numerical simulations. Intuitively, agents update their beliefs via least- squares and their beliefs lead to some realization of the ALM. Using the new data, agents test their models and decide to either accept or reject their model. Finally, the set of models are either changed or left alone. The algorithm is as follows:

Step 1. Agents are endowed with a model $m \in M$ and initial beliefs Φ_0.

Step 2. Endogenous value p_t is realized through market clearing.

Step 3. Each agent updates their beliefs $\Phi_{t,m}, \theta_{t,m}$ with respect to new data.

Step 4. After the grace period ξ, each agent also tests the specification of their model.

Step 5. If the null hypothesis H_0 is rejected, then agents draw another model at random.

Step 6. Repeat.

To further clarify the hypothesis testing mechanism, Figure 3.1 demonstrates a hypothetical model switching tree diagram. Figure 3.1 takes the switching probabilities of the other agents as given and demonstrates the probability of switching for the representative agent with model 1 (m1) after model rejection. Here the triple (n_1, n_2, n_3) denotes the fraction of the population with the respective models, where the first element is model 1, the second element is model 2, and the third element is model 3. For example, the triple $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ means that each representative agent holds one of the three models.

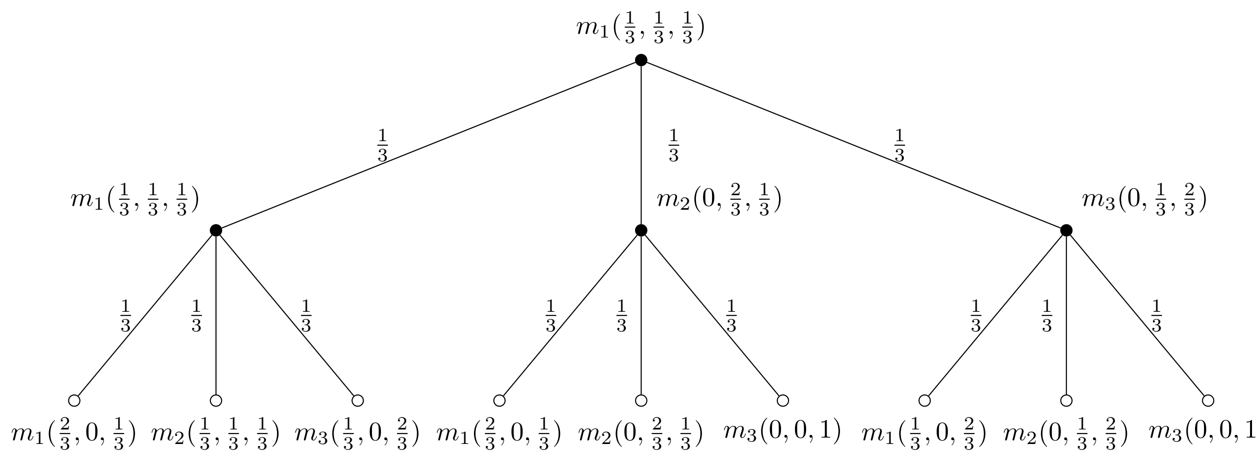

Figure 3.1: Model Switching Tree Diagram. Representative agent 1 given model 1 is rejected.

3.5.1 Parameterization

Since I am not doing a calibration exercise, the choice of parameterization is mainly for expositional purposes. I parameterize the model as follows:

Parameters	**Meaning**	**Calibration**
a	Risk Aversion	0.15
R	Risk-free Rate	1.02
μ	Mean Dividend	1
S	Mean Supply	1
σ_D^2	Std Dev. of Dividend	0.45
σ_S^2	Std Dev. of Supply	0.45
ρ	AR(1) Dividend Coefficient	0.45
η	AR(1) Supply Coefficient	0.45
σ_{DS}	Cov. of Dividend and Supply	0.45
ξ	Grace Periods	5
γ_1	Parameter Gain	t^{-1}
γ_2	LM Test Gain	0.00025
N	# of Iterations	100,000

Table 3.1: Parameter Values

I discuss the parameter choices for the hypothesis testing algorithm. I choose a grace period ξ of 5 to allow models to adapt to their environment when chosen. That is, agents wait 5 periods before testing their models. Next, the parameter gain γ_1 is set to a decreasing-gain initially to check for local stability under learning via a new definition I explain later. In subsequent sections, I also have agents implement a constant-gain as well. The LM test gain γ_2 is not chosen to be a decreasing-gain for the same reason as Cho and Kasa (2014) where I require the gain on the score statistic to be larger than the parameter updates or else the score statistics will depend on the history of the estimates rather than the current magnitudes. Intuitively, I want the updates to be slow enough such that the model is not rejected by a large shock but fast enough such that the test can detect parametric changes.[5]

3.5.2 HT-Stability

In order to test for local stability, I need to specify a definition that fits hypothesis testing learning. One issue with local stability is that, with a sufficiently high test-statistic, all the RPE are trivially locally stable because there is no switching in these scenarios.[6] Instead, I want a robust form of local stability when allowing for model switching. In particular, I want to understand which equilibria are robust in some sense to model changes i.e. a model shock. I call this HT-stability with respect to τ (or Hypothesis Testing stable), which can be thought of as a hypothesis test with a test-statistic threshold of τ.

I call an equilibrium HT-stable(τ) (HT-stable with respect to τ) if it satisfies the following criteria. First the equilibrium must be either an RPE or REE of the system. Second if the system is in an equilibrium, the equilibrium must be locally stable to all one-model shocks to the system. Similar to the idea behind impulse-responses, I define a one-model shock as follows. Suppose an environment converges to its equilibrium values after N periods. At period $N+1$, I force one of the representative agents to switch their model. An equilibrium is then HT-stable(τ) if after the model shock, it returns to the RPE over time.

With the definition above, many if not all the equilibria will possibly by HT-stable for some value τ, that is, if τ is sufficiently large, no amount of misspecification will reject the model. Hence when checking for HT-stability, I use a more stringent method. I say an equilibrium is HT-stable($\bar{\tau}$) if it converges to its equilibrium values after a one-model shock with respect to its equilibrium LM test-statistic value $\bar{\tau}$. I define the equilibrium threshold value $\bar{\tau}$ as the LM test-statistic that the model converges to without model switching. That is, if all agents have some model m, the equilibrium $\bar{\tau}$ is the average τ that results in this environment. One caveat of this definition is that some models will be more misspecified than

[5]Cho and Kasa (2014) explain in detail the technical reasons for this choice.

[6]We can think of the standard econometric learning case as one where the test threshold τ goes to ∞.

others and thus will be harder to reject in equilibrium.[7] Nevertheless, the above definition provides a useful benchmark to understand the interactions between adaptive learning and hypothesis testing with multiple agents.[8]

I find that not all of the RPE are HT-stable($\bar{\tau}$). In particular, the REE is not HT-stable($\bar{\tau}$). The fully-specified model has three parameters that need to be matched which makes the LM test very sensitive to changes in the environment. Hence a model shock causes other agents to eventually discard their model as well and over time it diverges away from the REE values.

I provide some intuition on why the REE is not HT-stable($\bar{\tau}$). The REE is not HT-stable($\bar{\tau}$) because when the REE is subject to a model shock, it raises the LM test-statistic causing the fully-specified model to be rejected. Over time, all models shift because the equilibrium threshold for the REE is close to zero. With a finite grace period ξ the LM test-statistic is unable to adjust below the threshold value and hence the model does not converge but switches between different equilibria.

In Table 3.2, I write down a list of all the potential equilibria (REE, RPE) and list if they are HT-stable($\bar{\tau}$) or not. I number them and call them "regimes" which will prove useful in the following section with constant-gain learning, where parameter values may be away from the equilibrium estimates. To reiterate, model 1 is the fully-specified model, model 2 is the dividend model, and model 3 is the supply model.

[7] While the one-model shock is not a local shock, it does reveal some properties of the strength of a particular RPE.

[8] Two interesting extensions would be to see what happens if we use a common τ across all models and to have a model where all models are equally misspecified and thus have equivalent $\bar{\tau}$.

Regime	**Equilibrium**	**HT-Stable($\bar{\tau}$)**
1	$(1, 0, 0)$	No
2	$(0, 1, 0)$	**Yes**
3	$(0, 0, 1)$	**Yes**
4	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$	No
5	$(\frac{2}{3}, \frac{1}{3}, 0)$	No
6	$(\frac{1}{3}, \frac{2}{3}, 0)$	No
7	$(0, \frac{1}{3}, \frac{2}{3})$	**Yes**
8	$(0, \frac{2}{3}, \frac{1}{3})$	**Yes**
9	$(\frac{1}{3}, 0, \frac{2}{3})$	No
10	$(\frac{2}{3}, 0, \frac{1}{3})$	No

Table 3.2: HT-Stability of RPE and Regime Numbers

I find that the only HT-stable($\bar{\tau}$) equilibria are combinations of models 2 and 3. In particular, since both of these models are already misspecified, their equilibrium threshold $\bar{\tau}$ are higher than the REE model. Next the higher $\bar{\tau}$ values allow certain models to be more robust towards changing beliefs. Intuitively, agents in these models are used to higher test-statistics and are less able to detect that the environment has changed. In contrast, agents with the fully-specified model more easily detect changes in the environment which is signaled through a higher LM test-statistic.[9]

I now show a case of HT-stability for the RPE corresponding to the all dividends model (0-1-0). For HT-stability I simulate the model for N periods and then shock the model at time N + 1 where I force one of the models to switch. The equilibrium LM test-statistic $\bar{\tau}$ for the all dividend RPE is $\bar{\tau} = 2$ and I set that as the threshold. I then run the model for another N periods. Figure 3.2 shows the process of HT-stability. We see that after the shock,

[9] Another potential extension would be to have a distribution of agents with different test-statistic thresholds.

the model stays in a different environment for some time and then as the LM test-statistic starts to increase in the other models, it shifts back to the (0-1-0) RPE.

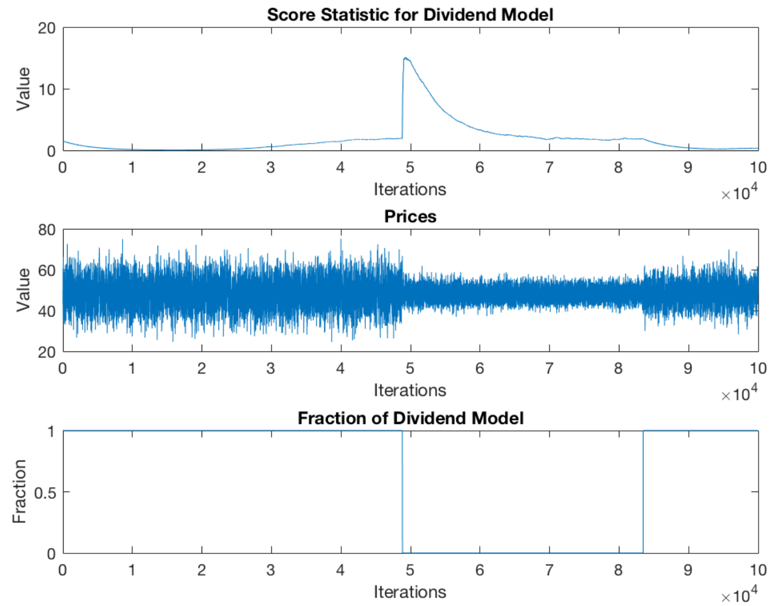

Figure 3.2: HT-stability Test for All Dividends RPE

One concern about HT-stability may be robustness. Both the choice of τ and grace period length ξ are important for convergence after the shock. In order to demonstrate some robustness, I also do the exercise for $\xi = 5$ and 10, $\gamma_2 = 0.0001$ and 0.00025, and $\bar{\tau} \pm 0.1$. I find that the results do not change for the different robustness checks.

3.5.3 Constant-Gain Learning

I now implement constant-gain learning and find that a subset of the RPE are utilized more than others. This is because some models are less sensitive to the LM test, depending on both the number of parameters that need to be estimated and the feedback properties of the restricted-perceptions equilibria (RPE). The constant-gain learning simulation is the computational counterpart to the analytical results of Cho and Kasa (2014), where they demonstrate that some models are "dominant" in the long-run i.e. models that are used "almost always".[10]

I run two different simulations and show the time spent within each model. For the first simulation I use $\gamma_1 = 0.001$ and $\gamma_2 = 0.03$ at 500,000 iterations, where γ_1 is the gain parameter for the estimates and γ_2 is the gain parameter for the score statistic. For the second simulation I use $\gamma_1 = 0.005$ and $\gamma_2 = 0.07$.[11] I set the test threshold τ_i to 2 and initialize the model at the RPE where all models are used $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, where again the triple (n_1, n_2, n_3) denotes the fraction of the population that holds each model. To insure that initial conditions do not heavily influence the results, I also burn in the first half of the simulations.

Again, for notational simplicity, I define the ten different possible model combinations as "regimes" rather than equilibria, given that at any time t, the parameter estimates may be away from the RPE (or REE) values. In Table 3.2, I define the regime numbers for the different states. For example, Regime 1 is $(1, 0, 0)$ which is the regime where all agents hold the correct model. Regime 4 is $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, where a fraction $\frac{1}{3}$ each hold one of the three models.

[10] The technical definition uses large deviation theory and analyzes the invariant distribution of the model set as the gain parameter goes to 0. Intuitively, their definition is similar to Foster and Young (2003), where one model is used almost all the time in the limit.

[11] As in Cho and Kasa (2014), it is convenient to raise γ_2^{α} where here $\alpha = \frac{1}{2}$. With some abuse of notation, $\gamma_2^{\frac{1}{2}} = 0.001^{\frac{1}{2}} \approx 0.03$.

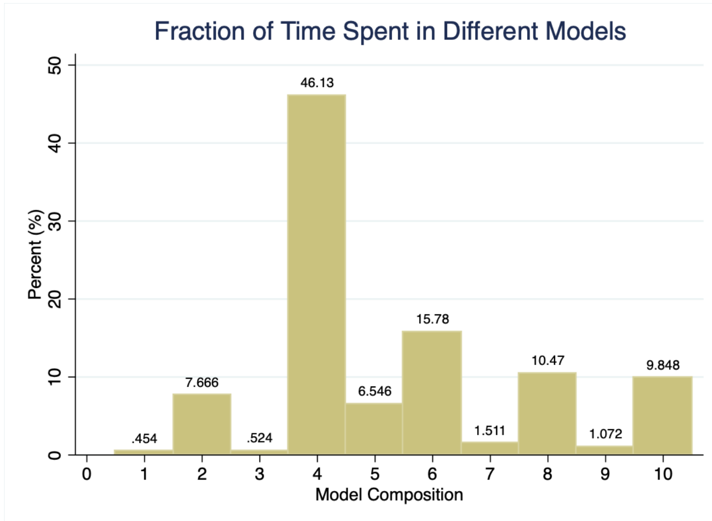

Figure 3.3: Constant Gain Learning Fractions $\gamma_1 = 0.001, \gamma_2 = 0.03$

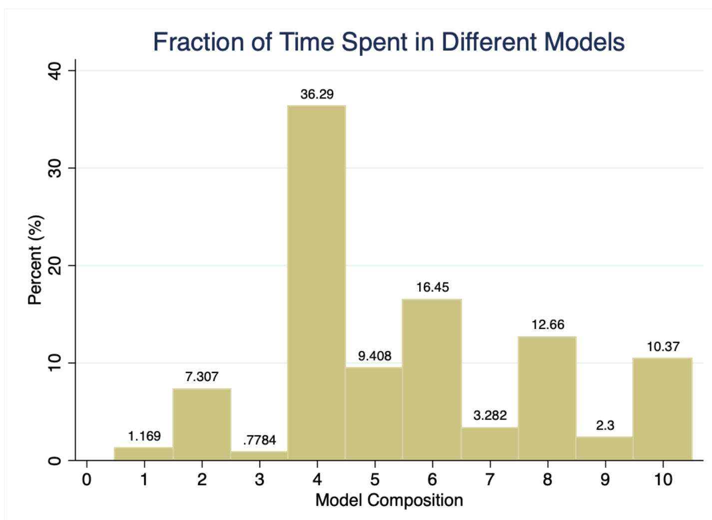

Figure 3.4: Constant Gain Learning Fractions $\gamma_1 = 0.005, \gamma_2 = 0.07$

Figures 3.3 and 3.4 show the fraction of the time spent within each regime. There are three findings from Figures 3.3 and 3.4. First, Regime 1, where all agents use the fully-specified model, is not used very often even though the test threshold is relatively high at $\tau_i = 2$. Next, the regime that the model spends the most time in is the hybrid regime $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, which is surprising given that it is not HT-stable. A key reason for this may be the discrete nature of the model switching mechanism here.[12] Finally, the "dominant model" in the simulations, where all agents hold one model most frequently, is the dividend model $(0, 1, 0)$. While the fact that the dominant model is not the correct model is surprising, it is consistent with Cho and Kasa (2014) who conjecture that the correct model will not necessarily be dominant depending on the feedback parameters of the model.

It is important to provide intuition why the correct model is not used the most by all agents. Similar to the explanation in Cho and Kasa (2014), there are two counterbalancing mechanisms involved. First, while a higher test threshold makes the specified model less likely to be rejected, all else equal, it also makes the misspecified models less likely to be rejected as well. Because the misspecified models are also held by other agents, the fully-specified model is then more likely to be rejected outside an REE. Thus certain equilibria may be more fragile than others depending on the feedback mechanisms within the model.

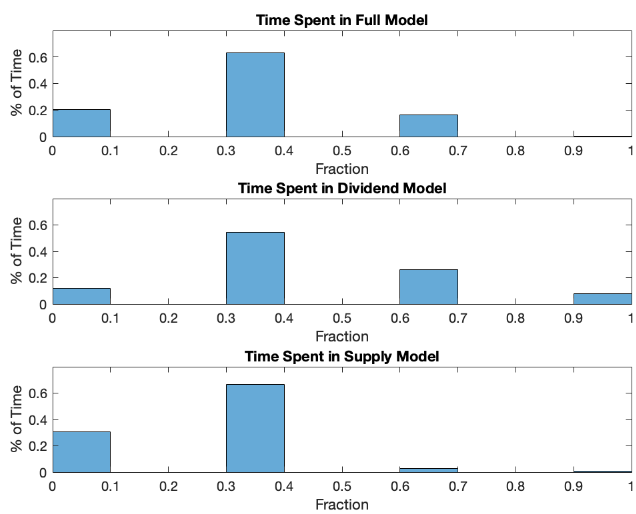

Figure 3.5: Constant Gain Learning Fractions $\gamma_1 = 0.001, \gamma_2 = 0.03$

[12]Future work can see how robust this finding is to different model switching mechanisms.

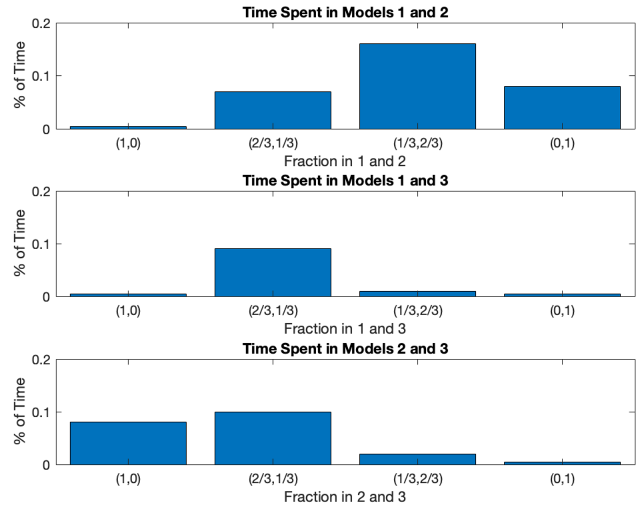

Figure 3.6: Constant Gain Learning Fractions $\gamma_1 = 0.001, \gamma_2 = 0.03$

Another way to visualize the simulations are Figures 3.5 and 3.6 which show how much time agents spend in each model. Figure 3.5 shows the fraction of time spent within each model while Figure 3.6 shows the data for two models holding one of the models fixed. For instance, in the top panel of Figure 3.6, the first column is how much time is spent in the fully-specified model $(1, 0, 0)$ and the second column is the amount of time spent in the regime $(\frac{2}{3}, \frac{1}{3}, 0)$. I remove the hybrid equilibria $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ to make a better comparison between the other regimes. Another interesting finding of the simulations is that the basin of attraction seems to be around regimes with the dividend model and in particular, regimes with model 3, the supply model, tend to not be seen. Finally, adding a higher gain to the simulation does not change the shape of the distribution but adds more weight onto certain hybrid regimes.

3.5.4 Robustness Check

Here I do two robustness checks. First I simulate the model with larger test thresholds $\tau_i = 10$ and a longer grace period $\xi = 20$ and find that it is robust to this specification. Figure 3.7 shows the first robustness check. Again, the finding is consistent with Cho and Kasa (2014) who demonstrate that the dominant model depends on the H-functional and the LD rate function[13] which is ambiguous even in the simple linear case.

For the second robustness check, I change the initial conditions of the new model when agents switch their models following rejection. Currently, when an agent switches their model, the initial conditions of the new parameter estimates are switched to the steady-state values consistent with the RPE (or REE) of the new model.

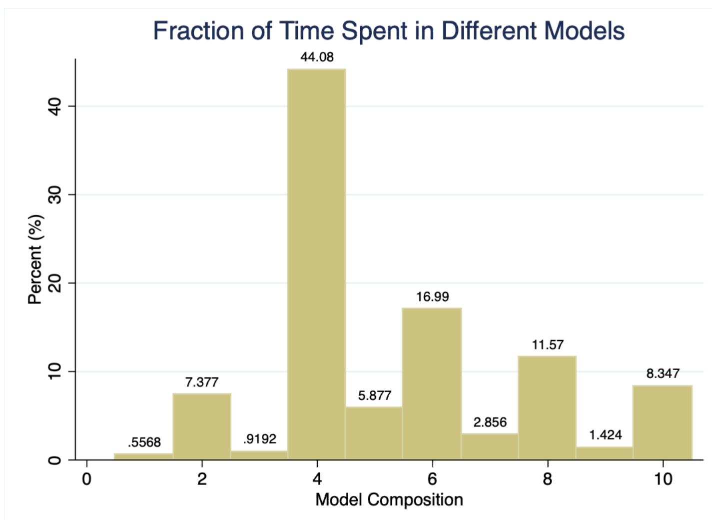

Figure 3.7: Constant Gain Learning Fractions $\gamma_1 = 0.001, \gamma_2 = 0.03, \tau_i = 10$, $\xi = 20$

[13]These functions characterize the probability of escape paths and feedback strengths of the RPE.

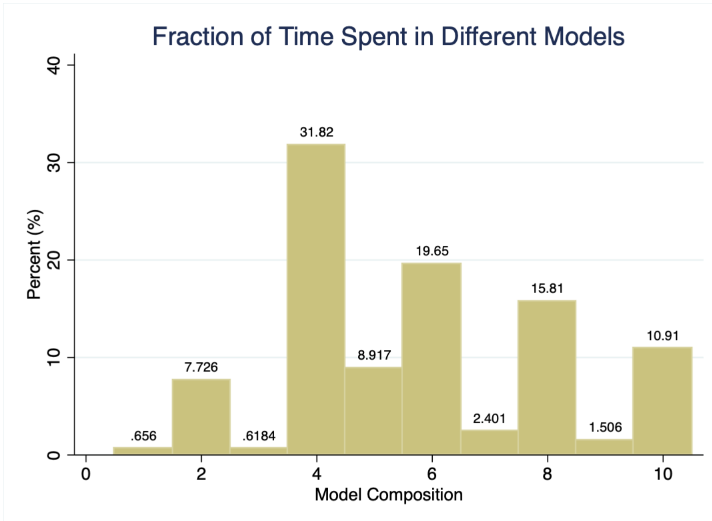

Figure 3.8: Constant Gain Learning Fractions with Different Switching Rule

In the second robustness check, I instead have agents switch to their previous parameter estimates and any estimates on the previously omitted variables (e.g. dividend D or supply S) are given the RPE (or REE) steady-state values as the initial conditions. Figure 3.8 shows that while this spreads the distribution more towards hybrid regimes, the same pattern persists. In particular, notice that instead of 45% of the time spent in Regime 4 ($\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$), 30% of the time is spent there, while more weight is placed on the other regimes.

I suspect that with a large enough test threshold and grace period, the REE may be the dominant model. Nevertheless, with constant-gain learning, the parameter estimates are constantly drifting, such that with a large enough shock, the model enters an escape path which triggers a rejection, making it difficult to return to the REE. Thus, as in Cho and Kasa (2014) the probability of entering the escape path determined by the LD rate function and the feedback properties of the RPE via the H-functional characterize which models are dominant in the long-run.

The constant-gain learning simulations and HT-stability suggest that the connection between them are not as strong as we would suspect, that is, HT-stability equilibria here do not necessarily seem to be regimes that are regularly visited under constant-gain learning.

One possible explanation is that the model shock here is a large shock and that HT-stability tells us more about the strength of the RPE rather than local stability. Future work should try to establish a stronger connection by possibly using a less discrete switching mechanism and refining the definition of stability.

3.6 Conclusion

I have demonstrated that in an asset pricing environment with hypothesis testing learning, the REE is not necessarily the stable equilibrium. Moreover, with constant-gain learning, agents spend most of their time using both the fully-specified and underparameterized models and the model switches between a subset of the RPEs.

The current project can be seen as a numerical examination of a potentially more general claim, which is that RPE can arise via a more natural learning mechanism than standard econometric learning. In particular, I fix the test thresholds, model class, and grace period. Since I have shown that under certain conditions, RPE can arise as the dominant model i.e. models that are used the most in the limit, the next step is to prove this in a more general case.

Cho and Kasa (2014) show that in a linear Gaussian version of their model, there is ambiguity in dominant models even if the model class includes the true model because of the interaction between the model feedback parameters and the functions that determine the escape paths (the H-functional and the Large Deviation (LD) rate function). An important step in the literature would be to see which parameter values in a linear Gaussian case lead to misspecified models being dominant models under hypothesis testing.

Bibliography

[1] K. Adam, A. Marcet, and J. Beutel. Stock Price Booms and Expected Capital Gains. *American Economic Review*, 2017.

[2] K. Adam, A. Marcet, and J. Nicolini. Stock Market Volatility and Learning. *The Journal of Finance*, 2016.

[3] F. Allen and D. Gale. Limited Market Participation and Volatility of Asset Prices. *American Economic Review*, 1994.

[4] L. Arrondel, H. Calvo-Pardo, and D. Tas. Subjective Return Expectations, Information and Stock Market Participation: Evidence from France. *Working Paper*, 2014.

[5] B. Babcock, E. Choi, and E. Feinerman. Risk and Probability Premiums for CARA Utility Functions. *Journal of Agricultural and Resource Economics*, 1993.

[6] M. Baker and J. Wurgler. The Equity Share in New Issues and Aggregate Stock Returns. *The Journal of Finance*, 2000.

[7] T. Bao, C. Hommes, and T. Makerewicz. Bubble Formation and (In)Efficient Markets in Learning-to-Forecast and Optimise Experiments. *The Economic Journal*, 2017.

[8] C. Bertaut and M. Starr-McCluer. Household Portfolios in the United States. *Household Portfolios*, 2002.

[9] W. Branch and G. Evans. A Simple Recursive Forecasting Model. *Economic Letters*, 2006.

[10] W. Branch and G. Evans. Model Uncertainty and Endogenous Volatility. *Review of Economic Dynamics*, 2007.

[11] W. Branch and G. Evans. Asset Return Dynamics and Learning. *The Review of Financial Studiesl*, 2010.

[12] W. Branch and G. Evans. Learning about Risk and Return: A Simple Model of Bubbles and Crashes. *American Economic Journal: Macroeconomics*, 2011.

[13] W. Branch and G. Evans. Bubbles, Crashes, and Risk. *Economic Letters*, 2013.

[14] C. Camerer. Behavioral Game Theory. *Princeton University Pressl*, 2003.

[15] J. Campbell. Household Finance. *The Journal of Finance*, 2006.

[16] C. Chambers and P. Healy. Updating Toward the Signal. *Economic Theoryl*, 2012.

[17] D. Chen, M. Schonger, and C. Wickens. oTree - An Open-Source Platform for Laboratory, Online, and Field Experiments. *Journal of Behavioral and Experimental Finance*, 2016.

[18] I. Cho and K. Kasa. Learning and Model Validation. *Review of Economic Studies*, 2014.

[19] I. Cho and K. Kasa. Gresham's Law of Model Averaging. *American Economic Review*, 2017.

[20] S. Crockett, J. Duffy, and Y. Izhakian. An Experimental Test of the Lucas Asset Pricing Model. *The Review of Economic Studiesl*, 2018.

[21] T. Csermely and A. Rabas. How to Reveal People's Preferences: Comparing Time Consistency and Predictive Power of Multiple Price List Risk Elicitation Methods. *Journal of Risk and Uncertainty*, 2016.

[22] J. Dominitz and C. Manski. Measuring and Interpreting Expectations of Equity Returns. *Journal of Applied Econometrics*, 2017.

[23] A. Drichoutis and J. Lusk. What Can Multiple Price Lists Really Tell us About Risk Preferences? *Journal of Risk and Uncertainty*, 2016.

[24] G. Evans and S. Honkapohja. Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models. *Econometrica*, 1995.

[25] G. Evans and S. Honkapohja. Learning and Expectations in Macroeconomics. *Princeton University Press*, 2001.

[26] A. Fagereng, C. Gottlieb, and L. Guiso. Asset Market Participation and Portfolio Choice Over the Life-Cycle. *The Journal of Finance*, 2017.

[27] D. Foster and P. Young. Learning, Hypothesis Testing, and Nash Equilibrium. *Games and Economic Behavior*, 2003.

[28] F. Gomes and A. Michaelides. Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence. *The Journal of Finance*, 2005.

[29] L. Guiso, M. Haliassos, and T. Jappelli. Household Portfolios. *MIT Press*, 2002.

[30] L. Guiso and T. Jappelli. Household Portfolios in Italy. *Household Portfolios*, 2002.

[31] L. Guiso and P. Sodini. Household Finance: An Emerging Field. *The Handbook of the Economics of Finance*, 2013.

[32] H. Guo. Limited Stock Market Participation and Asset Prices in a Dynamic Economy. *Journal of Financial and Quantitative Analysis*, 2004.

[33] F. Guvenen. A Parsimonious Macroeconomic Model for Asset Pricing. *Econometrica*, 2009.

[34] M. Hennequin. Experiences and Expectations in Asset Markets: An Experimental Study. *Working Paper*, 2018.

[35] C. Hommes. The Heterogenous Expectations Hypothesis: Some Evidence From the Lab. *Journal of Economic Dynamics and Control*, 2017.

[36] C. Hommes, J. Sonnemans, J. Tuinstra, and H. van de Velden. Coordination of Expectations in Asset Pricing Experiments. *The Review of Financial Studies*, 2005.

[37] M. Hurd, M. Van Rooij, and J. Winter. Stock Market Expectations of Dutch Households. *Journal of Applied Econometrics*, 2011.

[38] K. Lansing. Asset Pricing with Concentrated Ownership of Capital and Distribution Shocks. *American Economic Journal: Macroeconomics*, 2015.

[39] L. Ljungqvist and T. Sargent. A Labor Supply Elasticity Accord? *American Economic Review*, 2011.

[40] R. Lucas. Asset Prices in an Exchange Economy. *Econometrica*, 1978.

[41] U. Malmendier and S. Nagel. Depression Babies: Do Macroeconomic Experiences affect Risk-taking? *The Quarterly Journal of Economics*, 2011.

[42] U. Malmendier and S. Nagel. Learning from Inflation Experiences. *The Quarterly Journal of Economics*, 2015.

[43] A. Marcet and T. Sargent. Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory*, 1989.

[44] R. Merton. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, 1971.

[45] P. Moffat. Experimetrics: Econometrics for Experimental Economics. *Palgrave Macmillan*, 2015.

[46] A. Nakov and G. Nuno. Learning from Experience in the Stock Market. *Journal of Economic Dynamics and Control*, 2015.

[47] T. Norman. Learning, Hypothesis Testing, and Rational Expectations Equilbrium. *Games and Economic Behavior*, 2015.

[48] G. Orosel. Participation Costs, Trend Chasing, and Volatility of Stock Prices. *The Review of Financial Studies*, 1998.

[49] R. Rogerson. Indivisible Labor, Lotteries, and Equilibrium. *Journal of Monetary Economics*, 1988.

[50] T. Sargent. Bounded Rationality in Macroeconomics. *Oxford University Press*, 1993.

[51] T. Sargent. The Conquest of American Inflation. *Princeton University Press*, 1999.

[52] M. Shin. Endogenous Participation, Risk, and Learning in the Stock Market. *Working Paper*, 2018.

[53] M. Strahilevitz, T. Odean, and B. Barber. Once Burned, Twice Shy: How Naive Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold. *Journal of Marketing Research*, 2011.

[54] A. Timmermann. How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices. *The Quarterly Journal of Economics*, 1993.

[55] S. Van Nieuwerburgh and L. Veldkamp. Information Acquisition and Under-Diversification. *Review of Economic Studies*, 2010.

[56] A. Vissing-Jorgensen. Limited Asset Market Participation and the Elasticity of Intertemporal Substitution. *Journal of Political Economy*, 2002.

[57] N. Williams. Escape Dynamics in Learning Models. *Review of Economic Studies*, 2019.

Appendix A

Appendix for Chapter 1

Derivation of Equation (1.7): The expected value of the utility of c_{1t} is:

$$E_t u(c_{1t}) = E_t[-e^{-\rho c_{1t}}] = -e^{-\rho E_t c_{1t} + \frac{\rho^2}{2} V_t c_{1t}} \tag{A.1}$$

since c_{1t} is normally distributed. Then:

$$E_t c_{1t} = R + x_t[E_t(p_{t+1} + D_{t+1}) - Rp_t]$$
$$V_t c_{1t} = x_t^2 \sigma_p^2$$

We know:

$$x_t = \frac{E_t(p_{t+1} + D_{t+1}) - Rp_t}{\rho\sigma_p^2}$$

Plugging this into Equation (A.1) I get:

$$\begin{aligned} E_t u(c_{1t}) &= -e^{-\rho R - \rho[E_t(p_{t+1}+D_{t+1}) - Rp_t] + \frac{\rho^2}{2} x_t^2 \sigma_p^2} \\ &= -e^{\rho R - \rho \frac{E_t(p_{t+1}+D_{t+1}) - Rp_t}{\rho\sigma_p^2} + \frac{\rho^2}{2}\left[\frac{E_t(p_{t+1}+D_{t+1}) - Rp_t}{\rho\sigma_p^2}\right]^2 \sigma_p^2} \\ &= -e^{-\rho R - \frac{[E_t(p_{t+1}+D_{t+1}) - Rp_t]^2}{2\sigma_p^2}} \end{aligned}$$

And thus:

$$\Gamma(p_t) = e^{-\rho R} - e^{-\rho R - \frac{[E_t(p_{t+1}+D_{t+1})-Rp_t]^2}{2\sigma_p^2}}$$

Proof of Proposition 1: I focus on interior equilibria where $\frac{S\rho\sigma_p^2}{\mu} < Ae^{-\rho R}[1 - e^{-\frac{\mu^2}{2\sigma_p^2}}]$. Fix a set of parameters. Both $\frac{1}{\bar{x}}$ and $\bar{e}$ are compositions of continuous functions and hence continuous. With market clearing, $\frac{1}{\bar{x}} = \frac{\bar{e}}{S}$ implies $\frac{1}{\bar{x}}$ is below the $\bar{e}$ equation at $\bar{p} = 0$. If the equation $\frac{1}{\bar{x}}$ is above the participation curve, then the equilibrium condition is the intercept of the participation curve. Next I take the limit of $\frac{1}{\bar{x}}$ as $\bar{p}$ goes to $\frac{\mu}{R-1}$. The equation $\frac{1}{\bar{x}}$ approaches ∞. Since $\bar{e}$ is bounded and monotonic in $\bar{p}$, we know that there exists a point on $\frac{1}{\bar{x}}$ where $\frac{1}{\bar{x}} > 1$. Hence by, the intermediate value theorem, there exists a point where they cross. Because both curves are monotonic within the given parameter space and since $\bar{x} > 0$, it is unique. Hence, there exists a unique steady-state $\bar{p}$.

Proof of Proposition 2: I implicitly differentiate $\bar{e}$ with respect to the parameters. μ is not in the equation hence $\frac{\partial \bar{e}}{\partial \mu} = 0$. $\frac{\partial \bar{e}}{\partial R}$, $\frac{\partial \bar{e}}{\partial \sigma_p^2}$, $\frac{\partial \bar{e}}{\partial A}$ all follow from standard differentiation. A appears as a multiplier and hence $\frac{\partial \bar{e}}{\partial A} > 0$. $\frac{\partial \bar{e}}{\partial R} = \frac{A\bar{e}^3\rho(1-e^{\frac{\rho^2\sigma_p^2}{2\bar{e}^2}})}{\bar{e}^3 e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + A\rho^2\sigma_p^2}$. Since $e^{\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} > 1 \implies \frac{\partial \bar{e}}{\partial R} < 0$. $\frac{\partial \bar{e}}{\partial \sigma_p^2} = \frac{\bar{e}A\rho^2}{2\bar{e}^3 e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + 2A\rho^2\sigma_p^2} \implies \frac{\partial \bar{e}}{\partial \sigma_p^2} > 0$. $\frac{\partial \bar{e}}{\partial \rho} = \frac{A\bar{e}^3 R - A\bar{e}^3 e^{\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + A\bar{e}\rho\sigma_p^2}{\bar{e}^3 e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + A\rho^2\sigma_p^2}$. Hence $\frac{\partial \bar{e}}{\partial \rho}$ is positive when $R > e^{\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + \frac{\rho\sigma_p^2}{\bar{e}^2}$, negative if the sign is opposite and 0 at equality.

Proof of Proposition 3: I use the chain rule. Let Ω be the set of model parameters. Then $\bar{p} = f(\bar{e}(\Omega), \Omega)$ which implies $\frac{\partial \bar{p}}{\partial \Omega} = \frac{\partial f}{\partial \Omega} + \frac{\partial f}{\partial \bar{e}}\frac{\partial \bar{e}}{\partial \Omega}$. A only appears in $\bar{e}$. $\frac{\partial \bar{e}}{\partial A} > 0$ hence, $\frac{\partial \bar{p}}{\partial A} > 0$. μ does not appear in $\bar{e}$ and $\frac{\partial \bar{x}}{\partial \mu} > 0$ hence $\frac{\partial \bar{p}}{\partial \mu} > 0$. $\frac{\partial \bar{x}}{\partial R} < 0$ and $\frac{\partial \bar{e}}{\partial R} < 0$, hence $\frac{\partial \bar{p}}{dR} < 0$. $\frac{\partial \bar{x}}{\partial \sigma_p^2} < 0$ and $\frac{\partial \bar{x}}{\partial \rho} < 0$. $\frac{\partial \bar{p}}{\partial \rho} = -\left[\frac{\sigma_p^2}{\bar{e}} - \frac{\rho\sigma_p^2\frac{\partial \bar{e}}{\partial \rho}}{\bar{e}^2}\right]\frac{S}{R-1}$ and $\frac{\partial \bar{p}}{\partial \sigma_p^2} = -\left[\frac{\rho}{\bar{e}} - \frac{\rho\sigma_p^2\frac{\partial \bar{e}}{\partial \sigma_p^2}}{\bar{e}^2}\right]\frac{S}{R-1}$. Then $\frac{\partial \bar{p}}{\partial \sigma_p^2} < 0$ if $\frac{\partial \bar{e}}{\partial \sigma_p^2} < \frac{\bar{e}}{\sigma_p^2} \iff \frac{\sigma_p^2 A\rho^2}{2\bar{e}^e e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + 2A\rho^2\sigma_p^2} < 1 \iff \frac{\sigma_p^2 A\rho^2(-2\bar{e}^3 e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} - 1)}{2\bar{e}^3 e^{R\rho+\frac{\rho^2\sigma_p^2}{2\bar{e}^2}} + 2A\rho^2\sigma_p^2} < 0$ which is always true.

$\frac{\partial \bar{p}}{\partial \rho}$ is positive if $\frac{\partial \bar{e}}{\partial \rho} > \frac{\bar{e}}{\rho}$ and negative if less than, 0 at equality.

Proof of Proposition 4: Proposition 5.2 in Evans and Honkapohja (1995) is the result I use to prove my case. The requirements are that the gain parameter $\gamma > 0$ is a decreasing sequence, the shocks ϵ_t^S are iid with $E(\epsilon_t^S) = 0$, $Var(\epsilon_t^S) > 0$ and either (1) $|\epsilon_t^S| < \alpha$ with probability 1 for all t or (2) $E|\epsilon_t^S|^p$ exists and is bounded in t for each $p > 1$, and the derivatives of G and H are bounded. I claim to satisfy condition (1). First ϵ_t^S is iid by definition. Next for some α sufficiently small, as long as σ_S^2 is sufficiently small, or I bound the distribution of ϵ_t^S, then Proposition 5.2 holds.

Proof of Proposition 5: The proof depends on Evans and Honkapohja (2001) E-stability condition which requires the eigenvalues of the T-map to have negative real parts. First, if the sequence of shocks $\{\epsilon_t^S\}_{t=0}^{\infty}$ are such that $|\epsilon_t^S| < \alpha$ with probability 1 for all t and $\alpha > 0$ is sufficiently small then Proposition 4 holds and there exists a unique noisy steady-state REE. Our PLM is $p_t = a + \nu_t$ which implies that $p_t^e = a + \nu_t = p_{t+1}^e$. Then the ALM is $p_t = R^{-1}\left[a + \mu - \frac{S}{e_t}\rho\sigma_p^2\right]$. Then the T-map is: $\frac{da}{d\tau} = R^{-1}\left[a + \mu - \frac{S_t}{e_t}\rho\sigma_p^2\right] - a$ which can be rewritten as $\frac{da}{d\tau} = a(R^{-1} - 1) + R^{-1}(\mu - \frac{S_t}{e_t}\rho\sigma_p^2)$. Furthermore, since e is a function of a, I need to sign the derivative of e with respect to a which is $\frac{de}{da} < 0$. Given that $\frac{de}{da} < 0$, the T-map satisfies the local stability conditions by definition and hence proves my proposition.

Appendix B

Appendix for Chapter 2

B.0.1 Additional Summary Statistics

Here I list additional summary statistics for prices and participation rates per session. In particular, I list the median, 2nd half mean, and a time-series fit. The time-series fit which is called AR(1) in Table B.1, is a time-series regression on each session's price and participation series to estimate the constant and AR(1) coefficients. Then I take the implied unconditional average of the series using the estimated values. The additional summary statistics show that in general the different measures of average behavior are all very similar.

Treatment	Session	Med. Price	Med. Part.	AR(1) Price	AR(1) Part.
1	1	45.52	N/A	45.85	N/A
1	2	42.45	N/A	39.88	N/A
1	3	41.84	N/A	42.28	N/A
1	4	43.66	N/A	42.05	N/A
2	5	41.32	77.78%	40.82	74.2%
2	6	43.38	66.67%	42.5	58.22%
2	7	45.64	77.78%	46.32	78.59%
2	8	39.84	88.89%	40.15	81.7%
3	9	43.79	44.44%	45.29	49%
3	10	48.77	66.67%	50.74	63.83%
3	11	52.07	66.67%	51.73	65.9%
3	12	40.25	44.44%	40.24	46.4%
4	13	28.46	66.67%	27.52	64.79%
4	14	25.46	66.67%	22.32	59.14%
4	15	23.76	55.56%	23.42	54.92%
4	16	29.26	77.78%	27.22	66.67%

Table B.1: Additional Summary Statistics

B.0.2 Robustness Checks

I perform robustness checks. I first show the regressions for Finding 2 using the linear probability model with fixed effects. Since fixed effects are not biased with linear probability models, they are preferred to random effects. Table B.2 shows the results for the linear probability model.

Table B.2: Dependent Variable: Individual Participation (LPM)

Variable	**Model 1**	**Model 2**	**Model 3**
Subjective Returns$_t$	0.022***	0.021***	0.021***
	(0.005)	(0.005)	(0.005)
Actual Returns$_{t-1}$		0.017***	0.004
		(0.004)	(0.005)
Forecast Payoff$_{t-1}$		0.017***	0.014***
		(0.005)	(0.005)
Price Trend$_{t-1}$		0.011***	0.01**
		(0.004)	(0.004)
Past Positive Payoff$_t$			0.139***
			(0.023)
Past Negative Payoff$_t$			-0.002
			(0.029)
N	4896	4704	4704
R^2	0.059	0.085	0.149

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table B.3: Dependent Variable: Individual Participation (Logit)

Variable	**Model 1**	**Model 2**	**Model 3**
Subjective Returns$_t$	0.054***	0.052***	0.049***
	(0.015)	(0.017)	(0.017)
Actual Returns$_{t-1}$		0.017***	0.004
		(0.005)	(0.006)
Forecast Payoff$_{t-1}$		0.019***	0.015***
		(0.006)	(0.006)
Risk Aversion		0.04	0.036
		(0.027)	(0.026)
Price Trend$_{t-1}$		0.011**	0.008*
		(0.004)	(0.004)
Past Positive Payoff$_t$			0.154***
			(0.027)
Past Negative Payoff$_t$			0.001
			(0.028)
N	4896	4704	4704
Pseudo R^2	0.107	0.181	0.193

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

We can see that the coefficients are very similar to the probit model. Next I run the logit model with random effects. Table B.3 shows the results for the logit model. Again, the results are very similar. Finally, in Table B.4 I show the demographic regressions along with an added lag. The demographic regressions show that they are insignificant to the regression. The lagged regression shows that the only lags that matter are the 2-period lagged price trend and the direction of the 2-period lagged payoff.

Table B.4: Dependent Variable: Individual Participation

Variable	Model 4	Model 5
Gender	0.062	
	(0.059)	
Age	-0.012	
	(0.018)	
Major	-0.004	
	(0.007)	
Subjective Returns_t	0.027***	0.029***
	(0.009)	(0.008)
Subjective Returns_{t-1}		-0.005
		(0.004)
Actual Returns_{t-1}	0.004	0.006
	(0.006)	(0.006)
Actual Returns_{t-2}		-0005
		(0.004)
Past Positive Payoff_t	0.18***	0.188***
	(0.027)	(0.029)
Past Negative Payoff_t	0.003	0.023
	(0.03)	(0.029)
Past Positive Payoff_{t-1}		0.062***
		(0.023)
Past Negative Payoff_{t-1}		0.023
		(0.023)
N	4704	4608
Pseudo R^2	0.172	0.206

B.0.3 Instructions

Overview

Welcome to this experiment in economic decision-making. Please read the instructions carefully as they explain how you earn money from the decisions you make in today's experiment. We ask that you not talk with one another and that you silence your phones. If you have questions at any time please raise your hand and it will be answered in private. There will be a short quiz following the reading of the instructions which you will all need to complete before we can begin the session. **Also, at the end of the last round, we will give you a survey that pays you cash.**

Today's session will involve "rounds". Each round will have 2 "tasks": forecasting and entry advice. For each task you will view some information and make decisions. You will receive points for each task in each round. At the end of the session, we will randomly select 1 task. Your points from this task will be converted into dollars at 15 points = $1. Your earnings from the task, the survey, and your $7 show-up payment will be given privately in cash at the end of the session.

General Information

You are a **financial advisor** to an investment fund manager. The manager has 2 investment options: a risk-free investment and a risky investment. The risk-free investment is putting all the money into a bank account paying a fixed interest rate. The risky investment is holding stocks which requires a *transaction fee* to buy. Your 2 tasks are:

1) Forecast the stock market price as accurately as possible and

2) Provide entry advice to the manager *(hold the stock or not).*

To make the best decision, the manager needs to know what the stock price will be. As the advisor, you have to predict the stock price (in francs) during 51 rounds and tell the manager if he should buy the stock or not in each round.

Each manager has a different transaction fee for holding stocks. The transaction fees are fixed per manager (e.g. each manager's fee does not change in all rounds) and ranges evenly from 0 to 4 francs per round (with no manager having a fee of 0). The manager makes profits each round. If the manager does not buy stocks he makes 3 francs that round. If he buys the stock, he makes uncertain profit: dividends which are 3 francs per round plus a capital gain from stocks (which can be negative). Therefore good entry advice depends on good forecasts. Your points depend on forecasting accuracy and the manager's profits.

Market Information

The stock price is determined by equilibrium between the supply and demand of stocks. The supply of stocks is fixed. The demand for stocks is mainly determined by the total demand of a number of investment funds active in the stock market. Some of these funds are advised by a participant in the experiment, others use a fixed strategy.

The more funds there are in the market, the higher the demand for stocks on average. There is also 1 fund who will always enter the market. There is also some uncertain, small demand for stocks by private investors but their effect on the stock price is small. *Stock prices are determined by equilibrium, that is, the stock price in round t will be the price where total demand equals supply.*

Manager's Investments Information

The exact investment strategy of your manager and the strategies of the other funds are unknown. The risk-free bank account pays a **fixed interest rate of 5% per round**. Stockholders receive a certain **dividend of 3 francs per round**. Stock returns per round are uncertain and depend on dividends and stock price changes.

Based on your stock price forecast and entry decision, your manager will make an optimal investment decision *(e.g. some money into the bank account and some into stocks)*. The higher your price forecast, the larger will be the fraction of money invested by your manager in stocks, so the larger will be their demand for stocks. If you tell the manager not to enter, then he invests everything in the risk-free bank account.

Task 1: Forecasting Prices

Your 1st task is to forecast the stock market price in each round as accurately as possible. The stock price will always be between 0 and 100 francs. The stock price has to be predicted both **one** and **two** rounds ahead. So at the start of each round you will make 2 stock price predictions *(e.g. this round and next)*. If the manager enters, he will use your **two** round ahead forecast to make his optimal investment decision. Your forecasts can be made up to 2 decimal points.

At the start, you have to predict the stock price in the 1st **two** rounds, that is, you have to give forecasts for rounds 1 and 2. After everyone has given their forecasts for the 1st two rounds, along with their entry advice, the stock price in round 1 will be revealed and based on your forecasting error, your points for round 1 will be given. After that, you have to give your forecast for rounds 2 and 3, along with entry advice for round 2. After everyone has given their forecasts and entry advice in round 2, the stock price in round 2 will be revealed and based on your forecasting error and the manager's profits, your earnings for round 2 will be given. This continues for 51 rounds.

It is important to note that you make 2 forecasts for each round's price after round 1. During round 1, you make a forecast of the stock price in round 2 and in round 2, you make *another* forecast of the stock price in round 2. This is because at the start of round t, you do not know the stock price in round t since it is revealed at the end of the round.

Task 2: Entry Advice for the Manager

Your 2nd task is to give entry advice. Your 2 choices are to **enter** the stock market or **not enter**. The manager will follow your advice completely. The manager makes a profit based on your decision. If the manager does not enter, the manager makes a service fee of 3 francs that round. Otherwise, the manager makes a profit:

$$\text{Price}_{t+1} + 3 - 1.05 * \text{Price}_t - \text{Manager's fee}$$

where price$_t$ is price in round t, 3 is the dividend, and 1.05 is the gross interest rate. Hence his profits depend on the price change after entry *(e.g. Price$_t$ and Price$_{t+1}$)*.

Your job will be to make sure the manager makes the decision that maximizes his per round profits. In each round, you will receive 3 points plus the manager's profits if the manager's profits are positive (up to 5 points) and 3 points minus the manager's profits if they are negative (down to 1 point). If he does not enter, you will receive 3 points. The following is a timeline of your tasks in each round:

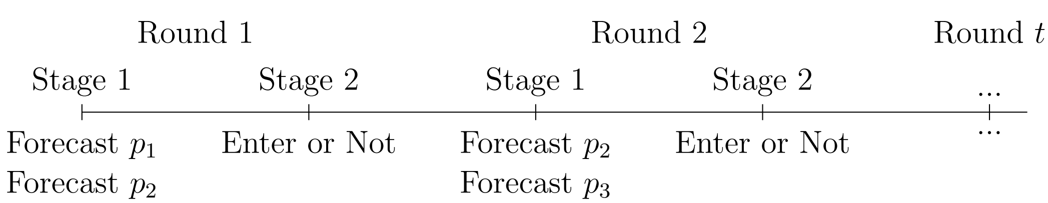

Points

Your points will depend on your forecasting accuracy and the manager's profits. For forecasting, the better you predict the stock market price in each round, the more points you get. Your points for forecasting are:

$$\text{Forecast points in round } t = \frac{16}{2 + |\text{Price in round } t - \text{1st forecast}| + |\text{Price in round } t - \text{2nd forecast}|}$$

where $|\cdot|$ is an absolute value (deviation), e.g. $|10 - 13| = 3, |5 - 4| = 1$. **The accuracy of both your 1st and 2nd forecast will matter.** You can earn up to 8 points in each round if you predict the stock price exactly both times.

Your points for entry are:

$$\text{Entry points in round } t = \begin{cases} \min\{3 + MP, 5\} & \text{if entered in t-1 and } \underbrace{p_t + 3 - 1.05p_{t-1} - k_i}_{\text{manager's profits}} \geq 0 \\ \max\{3 + MP, 1\} & \text{if entered in t-1 and } \underbrace{p_t + 3 - 1.05p_{t-1} - k_i}_{\text{manager's profits}} < 0 \\ 3 & \text{if did not enter in t-1} \end{cases}$$

where MP is the manager's profits, p_t is the price in round t, and k_i is the manager's transaction fee. And where max chooses the maximum value, e.g. $\max\{3, 5\} = 5$ (min chooses the minimum value). Hence the maximum points you can make is 5 and the minimum is 1. We will provide you with your manager's expected profits based on your forecasts. *Note that entry points in round t depend on your entry choice in round $t-1$.* **Remember we will pick 1 task (forecasting or entry) at random to pay you at the end of the experiment.**

Forecasting Example: Suppose the price in round 7 was 70 francs and you guessed 65 in round 6 and 60 in round 7. Then your points:

$$\frac{16}{2 + |70 - 65| + |70 - 60|} = \frac{16}{2 + 5 + 10} = 0.94 \text{ points}$$

Entry Example: Suppose the price in round 8 was 80 francs and the price in round 7 was 70 francs and the manager's transaction fee was 1. Then your points in round 8 if you entered in round 7 would be:

$$80 + 3 - 1.05 * 70 - 1 = 9.5 + 3 = 12.5 \geq 5 \implies \text{ 5 points}$$

If you did not enter, you would get 3 points. If it were negative you get less than 3 points.

The following is a table of possible points for forecasting. It is important to note that the table does not give all values and that points are rounded to 2 decimal points.

Earnings Table (in points)

	Error 2 = 0	0.5	1	1.5	2	2.5	3	3.5
Error 1 = 0	8	6.4	5.33	4.57	4	3.56	3.2	2.91
0.5	6.4	5.33	4.57	4	3.56	3.2	2.91	2.67
1	5.33	4.57	4	3.56	3.2	2.91	2.67	2.46
1.5	4.57	4	3.56	3.2	2.91	2.67	2.46	2.29
2	4	3.56	3.2	2.91	2.67	2.46	2.29	2.13
2.5	3.56	3.2	2.91	2.67	2.46	2.29	2.13	2
3	3.2	2.91	2.67	2.46	2.29	2.13	2	1.88
3.5	2.91	2.67	2.46	2.29	2.13	2	1.88	1.78

Now please complete a short quiz. You can use the instructions sheet. Please raise your hand when you are done and we will come around to check your answers. After everyone has finished, we will let you know when you can begin the experiment.

Question 1

What is the smallest and largest prices the stock can be?

Suppose you are at the beginning of round 10. Which rounds' prices do you need to predict?

Question 2

If you advise your manager to enter, which forecast will your manager use to make his investment decision (1 or 2 rounds ahead)?

How much does the stock give in dividends (in francs) per round?

What is the interest rate that the bank account pays per round?

Question 3

What is the range of all manager's transaction fees?

What is the most amount of points you can earn for entry advice? What is the least amount? How many points do you receive if you do not enter?

If you advise the manager to not enter the stock market, will he buy stocks?

Question 4

Suppose the price for round 20 is 50 francs. You guessed 45 in round 19 (for price in round 20) and 50 in round 20 (for price in round 20). How many points would you receive for forecasting the price in round 20?

Suppose your manager's profit in round 6 is 3. How many points would you earn? (Hint: $\min\{3 + MP, 5\}$). Suppose your manager's profit in round 6 is -1. How many points would you earn? (Hint: $\max\{3 + MP, 1\}$)

Suppose the price for round 21 is 55 francs and price for round 20 is 50 francs. Your manager's transaction fee is 1. If you entered in round 20, what is the manager's profits? How many points would you earn at the end of round 21? (Hint: $p_t + 3 - 1.05p_{t-1} - k_i$, where k_i is the manager's transaction fee)

Appendix C

Appendix for Chapter 3

Here is list the analytical forms for the RPE. First let:

$$b = R^{-1}$$
$$r = \frac{ED_t S_t}{ED_t^2}$$
$$\tilde{r} = \frac{ED_t S_t}{ES_t^2}$$

Then the RPE is:

$$A_1 = A_2 = A_3 = \frac{\mu - Sa\sigma_p^2}{R-1}$$
$$\psi = 1 - b\eta + b(-\rho + \phi b\rho\phi - br\rho\phi\tilde{r})n_2 + b^2\rho\phi(-1 + r\tilde{r})n_2^2 + b\rho n_1(-1 + b\phi + b\phi(-1 + r\tilde{r}))n_2$$
$$B_1 = \frac{b\rho(1 - b\phi + b(-ar\sigma_p^2 + \phi)n_2)}{\psi}$$
$$B_2 = \frac{b(\rho - ar\sigma_p^2 - b\rho\phi + abr\rho\sigma_p^2 n_1 + b\rho\phi n_2 - br\rho\phi\tilde{r}(-1 + n_1 + n_2))}{\psi}$$
$$C_1 = \frac{b(-b\rho\phi\tilde{r}(-1 + n_1 + n_2) + a\sigma_p^2(-1 + b\rho n_1 + b\rho n_2))}{\psi}$$
$$C_3 = \frac{b(a\sigma_p^2(-1 + b\rho n_1 + b\rho n_2) - \rho\tilde{r}(-1 + b\phi n_1 + abr\sigma_p^2 n_2))}{\psi}$$