

Endogenous Participation, Risk, and Learning in the Stock Market

Michael Shin *

April 18, 2018

Abstract

A simple asset pricing model with endogenous participation 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. We 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.

Keywords: stock market participation, adaptive learning, excess volatility, risk, asset pricing

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

*Department of Economics, University of California, Irvine. I am deeply indebted to Bill Branch for his extensive input on this paper. I would also like to thank John Duffy, David Hirshleifer, Chong Huang, Fabio Milani, Guillaume Rocheteau, Eric Swanson, Pat Testa, members of the UCI Macro workshop, and participants at Midwest Macro for their valuable feedback. All errors are my own. Email: mikeshin@uci.edu.

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. Stock Market Participation Rates from 1995 - 2013. Data was extracted from the Survey of Income and Program Participation.¹

prices. Table 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. We propose a theory of endogenous fluctuations in participation rates and demonstrates that it can be an important driver of stock price volatility.

We 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. We 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 provide two different feedback mechanisms that contribute to price fluctuations with learning about risk being quantitatively more important. We 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, we implement a cost function which captures features beyond fixed participation costs while keeping the model

¹In particular, the values are taken from official data tables produced by the U.S. Census which are only periodically documented. These rates include the percentage of households in the survey that stated ownership in either stock or mutual fund holdings. We also find a similar pattern when adding retirement accounts such as IRA and 401K holdings. Similar patterns are found in the SCF and PSID.

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, we 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. We argue that belief-driven learning dynamics are key in explaining the interplay between participation and stock price volatility. Hence we take a step down from RE and implement an adaptive learning rule.

We first characterize the steady-state equilibrium and do comparative statics which give insights on participation without learning. We 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.

We 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. We 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.

We 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.

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 but 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. Our 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, we focus on aggregate participation and how it jointly impacts asset prices and expected returns in the time-series.

Second, we 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 we know this is the first paper to combine an asset pricing model with endogenous participation and learning.

2 Model

Time is measured in discrete periods $t = 1, 2, \dots$ 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 we take 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. We view 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. We 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.² After the initial old is endowed with the shares, subsequent S follow an exogenous process:

$$S_t = S + \epsilon_t^S \tag{1}$$

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).

We 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). Each agent lives for two periods, has initial endowment w normalized

²In order to focus on the interactions between learning and participation, we abstract from seriously modeling dividends and asset supply, both of which are better approximated by persistent or non-stationary processes.

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). We 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.

Our 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 Sargent and Ljungqvist (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.³ Because of the Law of Large Numbers, e also corresponds to the aggregate participation rate.

3 Equilibrium

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:

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

³Their 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).

$$\text{subject to } c_{\chi t} = \begin{cases} R + x_t(p_{t+1} + D_{t+1} - Rp_t) & \text{if } \chi = 1 \\ R & \text{if } \chi = 0 \end{cases} \quad (2)$$

where x_t is the asset holding decision and p_t is the price of the risky asset. Equation (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, we arrive at the following first-order conditions:

$$x_t = \frac{E_t(p_{t+1} + D_{t+1}) - Rp_t}{\rho\sigma_p^2} \quad (3)$$

$$\Phi'(e_t) = \max\{E_t u(c_{1t}) - u(c_{0t}), 0\} \quad (4)$$

where $\sigma_p^2 \equiv \text{Var}_t(p_{t+1} + D_{t+1})$ 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:

$$e_t = \min\{\Phi'^{-1}[E_t u(c_{1t}) - u(c_{0t})], 1\} \quad (5)$$

Equation (3) is the standard mean-variance asset demand function which is downward sloping in the price and Equation (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.

3.2 Steady-State

To better understand the participation decision, it is illustrative to analyze the steady-state. We now 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 we now get:

$$e_t = \min\{A\Gamma(p_t), 1\} \quad (6)$$

where:

$$\Gamma(p_t) = e^{-\rho R} - e^{-\rho R - \frac{[E_t(p_{t+1} + D_{t+1}) - Rp_t]^2}{2\sigma_p^2}} \quad (7)$$

Equation (7) follows from the transformation of an exponential function with respect to normal random variables and is derived in the appendix. $\Gamma(p)$ is the expected utility difference between the two states which can be interpreted as the expected excess utility return of entering the stock market. The market-clearing condition is as follows:

$$e_t x_t = S_t$$

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

$$p_t = R^{-1} \left[E_t(p_{t+1} + D_{t+1}) - \frac{S_t}{e_t} \rho \sigma_p^2 \right] \quad (8)$$

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 (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 we 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, we look at the steady-state equilibrium. We 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.

We characterize the steady-state equilibrium where $S_t = S$ and $p_{t+1} = p_t = \bar{p}$. Once we solve for the steady-state we get the following form:

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

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

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

where:

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

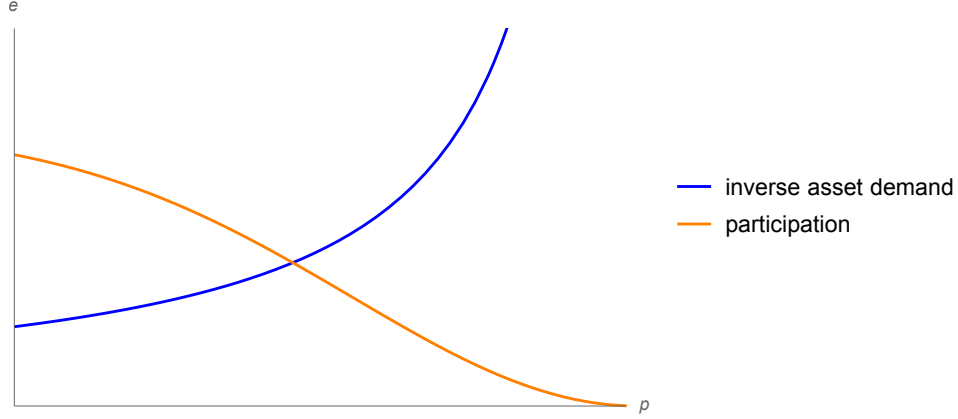


Figure 1. Steady-state Equilibrium.

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

Proofs are provided in the appendix. Figure 1 depicts Proposition 1 graphically for a set of parameters. Given that the steady-state equilibrium exists and is unique, we derive the expression for the steady-state price. We 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:

$$\bar{p} = \frac{\mu - S\rho\sigma_p^2}{R - 1} \quad (9)$$

The steady-state equation for our model is:

$$\bar{p} = \frac{\mu - \frac{S}{\bar{e}}\rho\sigma_p^2}{R - 1} \quad (10)$$

where:

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

When $\bar{e} = 1$, our 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. We graph Equations (9) and (10) in Figure 2 to describe the relationship between the two models.

In Figure 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, we can plug Equation (10) into the steady-state participation function to find \bar{e} as an implicit function of the fundamentals:

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

We 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.*

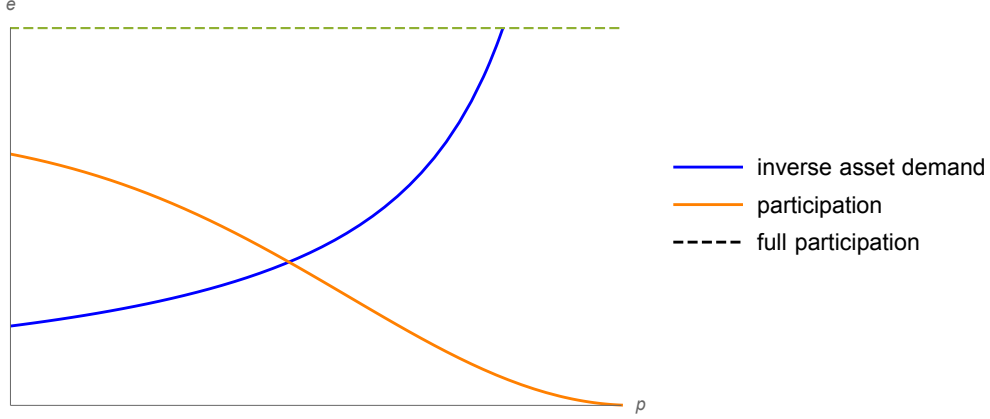


Figure 2. Steady-state Equilibrium: Limited and Full Participation.

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 (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.

We 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.

3.3 Endogenous Risk

We 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. We 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.

We now endogenize $\sigma_p^2 \equiv \text{Var}_t(p_{t+1} + D_{t+1})$. Then:

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

Solving out we have:

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

Solving for equilibrium risk σ_p^2 leads to:

$$\sigma_p^2 = \frac{\bar{e}^2 \pm \bar{e}\sqrt{\bar{e}^2 - 4R^{-2}\rho^2\sigma_S^2\sigma_D^2}}{2R^{-2}\rho^2\sigma_S^2} \quad (11)$$

Equation (11) 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 (11) which correspond to low and high risk steady-states. Branch and Evans (2011) show that the low risk steady-state is unstable under learning. We 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, we are unable to provide analytical

solutions for the case with endogenous risk. Instead, we rely on numerical analysis under learning.

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, we are 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. We characterize the noisy steady-state REE and then analyze its stability under learning. We do this by first taking the risk σ_p^2 as exogenous to clearly understand the dynamic properties of the participation decision. We then analyze the numerical properties when σ_p^2 is endogenous.

4.1 Rational Expectations Equilibrium

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

$$p_t = R^{-1} \left[E_t p_{t+1} + \mu - \frac{S_t}{e_t} \rho \sigma_p^2 \right] \quad (12)$$

The noisy steady-state is characterized by a function $p(\epsilon_t^S)$ that solves Equation (12). Then the noisy steady-state takes the following form:

$$p(\epsilon_t^S) = R^{-1} \left[E_t p(\epsilon_{t+1}^S) + \mu - \frac{S}{e(E_t p(\epsilon_{t+1}^S), p(\epsilon_t^S))} \rho \sigma_p^2 - \frac{\epsilon_t^S}{e(E_t p(\epsilon_{t+1}^S), p(\epsilon_t^S))} \rho \sigma_p^2 \right] \quad (13)$$

where:

$$e = \min \left\{ A e^{-\rho R} - A e^{-\rho R - \frac{[E_t p(\epsilon_{t+1}^S) + \mu - R p(\epsilon_t^S)]^2}{2\sigma_p^2}}, 1 \right\} \quad (14)$$

Plugging Equation (13) into Equation (14) we get:

$$e(\epsilon_t^S) = \min \left\{ A e^{-\rho R} - A e^{-\rho R - \frac{\left[\frac{S}{e(\epsilon_t^S)} \rho \sigma_p^2 + \frac{\epsilon_t^S}{e(\epsilon_t^S)} \rho \sigma_p^2 \right]^2}{2\sigma_p^2}}, 1 \right\}$$

Then:

$$p(\epsilon_t^S) = R^{-1} \left[E_t p(\epsilon_{t+1}^S) + \mu - \frac{S}{e(\epsilon_t^S)} \rho \sigma_p^2 - \frac{\epsilon_t^S}{e(\epsilon_t^S)} \rho \sigma_p^2 \right] \quad (15)$$

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

$$p_t = R^{-1} [\bar{p} + \mu] + \eta_t \quad (16)$$

where $\eta_t \equiv -R^{-1} \left[\frac{S}{e(\epsilon_t^S)} \rho \sigma_p^2 + \frac{\epsilon_t^S}{e(\epsilon_t^S)} \rho \sigma_p^2 \right]$. Equation (12) has two noisy steady-state REE: a fundamentals REE as in Equation (16) and a bubbles REE. As is standard in the asset pricing literature, we restrict our attention to the fundamentals REE, which is the unique non-explosive REE. It is also well known that the bubbles REE is unstable under learning. Moreover, we can see from Equation (15) that supply shocks are driving the noisy steady-state. We 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 (16) is the unique noisy steady-state REE solution to Equation (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 we have fully characterized the noisy steady-state REE of our model and now we 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 our numerical simulation we use empirical moments and robustness checks to insure that the system is locally stable.

4.2 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 we adhere to the Cognitive Consistency Principle (Sargent 1993) which requires agents and econometricians to be on equal footing. In this regard we 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. We make a small deviation from RE where agents implement a learning rule and run least-squares regressions on the perceived pricing function.⁴

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. We can rewrite the sample average recursively where expectations formation

⁴We 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.

take the following form:

$$p_{t+1}^e = p_t^e + t^{-1}[p_{t-1} - p_t^e]$$

where p_t^e 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. First thing to note is that p_t is assumed to be unknown to the agent during the time of the forecast. Unlike RE, having beliefs and outcomes be determined simultaneously with learning is not reasonable. Hence, we have agents forecast both 1-period (p_t^e) and 2-period (p_{t+1}^e) ahead prices. For instance, $\Gamma(p)$ is as follows:

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

Then p_t is unknown to the agent at time t . If p_t is unknown at time t , then the agent's best forecast is p_t^e . 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 $p_t^e = a_{t-1} = p_{t+1}^e$. The equation then becomes:

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

This formulation has the same steady-state as before which allows us to analyze the learning dynamics with respect to Equation (12). We say that the REE is locally stable if the model converges to the REE under decreasing-gain learning. We check the properties of the model with decreasing-gain learning and show that the REE is in fact stable under learning. Since there is a kink at $e_t = 1$ some of the analysis may not hold at the corner. Hence, we focus on parameterizations that keep e_t away from the corner i.e. interior solutions.

4.3 Stability Under Learning

We show analytically that the REE solution is locally stable under learning. To do this we 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 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:

$$p_t = R^{-1} \left[a + \mu - \frac{S_t}{e_t} \rho \sigma_p^2 \right]$$

$$e_t = \min \left\{ A e^{-\rho R} - A e^{-\rho R - \frac{[(1-R)a + \mu]^2}{2\sigma_p^2}}, 1 \right\}$$

Plugging into the learning rule, we get:

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

$$T(a) = R^{-1} \left[a + \mu - \frac{S_t}{e_t(a)} \rho \sigma_p^2 \right]$$

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 real parts, then the REE is locally stable. The fixed point of the ODE is:

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

$$e = \min \left\{ A e^{-\rho R} - A e^{-\rho R - \frac{S^2 \rho \sigma_p^2}{2e^2}}, 1 \right\}$$

where (a, e) correspond to the steady-state values. We 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 our 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 our model is locally stable under learning.

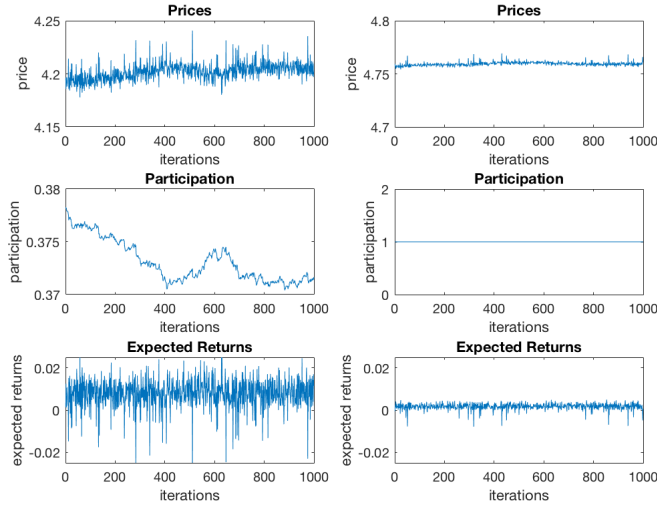


Figure 3. Constant-Gain Learning. For 1000 iterations, $\gamma = 0.05$, $\sigma_S^2 = 0.435$, $A = 1.15$, $\sigma_p^2 = 0.46$, $\mu = 1$, $R = 1.007$, $\rho = 0.45$. Left figure is limited participation, right figure is full participation.

4.4 Constant-Gain

So far we have demonstrated the model properties under decreasing-gain learning. In our simulations we implement constant-gain learning, where agents weigh each observation with geometrically declining weights. This is appropriate because our 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. We 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 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 is endogenized. We 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 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.

4.5 Learning about Risk

We 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.⁵ Then the learning rule for σ_p^2 is:

$$\sigma_{p,t+1}^2 = \sigma_{p,t}^2 + \gamma[(p_t - p_{t-1}^e + \epsilon_t^D)^2 - \sigma_{p,t}^2]$$

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, we find that the low risk steady-state is numerically stable under learning while the high risk steady-state is not. Figure 4 depicts the simulation with learning about risk.

We 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 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 our framework.

⁵Alternatively, 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.

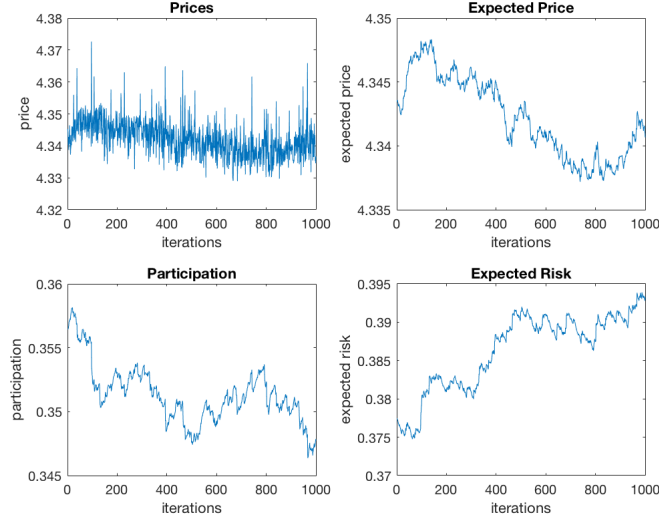


Figure 4. Endogenous Risk. For 1000 iterations, $\gamma_1 = 0.05$, $\gamma_2 = 0.0005$, $\sigma_S^2 = 0.435$, $A = 1.15$, $\sigma_D^2 = 0.28$, $\mu = 1$, $R = 1.007$, $\rho = 0.45$

5 Quantitative Analysis

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

5.1 Parameter Values

The parameters are calibrated according to the values in Table 2. The risk aversion ρ is calibrated to a value within the range of studies found in Babcock, Choi, and Feiernerman (1993) at 0.45. The historical average real interest rate in the U.S. is 2.7% so we 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 we 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, we choose the cost parameter $A = 1.15$ as a benchmark which corresponds to the 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. 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, we are interested in the unconditional moments so we take a long, transient simulation of two million iterations and burn-in the first one million.

Parameters	Meaning	Calibration	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
μ	Mean Dividend	1	Calibrated to reasonable Mean-SD ratio
A	Cost Parameter	1.15	Calibrated to SIPP participation 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 2. Parameter Values.

5.2 Moments

The moments we are 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. We do a quarterly interpolation of the SIPP stock market participation data including retirement accounts. There is a small trend component at the beginning of the series due to structural changes with tax policies so we take the HP filter of the series. We 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. 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.

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

5.3 Results

Table 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 25% of the volatility in stock prices. Since we abstract from many important features of the stock market, we can interpret the findings as endogenous participation independently accounting for 25% of the excess volatility in stock prices. We 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 we 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.

We can also match half of the mean excess returns at 1.032 and we do much better at matching the autocorrelation of stock prices at 0.959 while without learning about risk, the autocorrelation is 0.625. We are unable to match the standard deviation of excess returns at 0.047. We 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.⁶

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

⁶Suitably extended versions of the model can explain returns such as in Adam et al (2017).

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.

6 Conclusion

We 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 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.

7 Appendix

Derivation of Equation (7): The expected value of the utility of c_{1t} is:

$$E_t u(c_{1t}) = E_t[-e^{-\rho c_{1t}}] = -e^{-\rho E_t c_{1t} + \frac{\rho^2}{2} V_t c_{1t}} \quad (17)$$

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

$$\begin{aligned} E_t c_{1t} &= R + x_t [E_t(p_{t+1} + D_{t+1}) - R p_t] \\ V_t c_{1t} &= x_t^2 \sigma_p^2 \end{aligned}$$

We know:

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

Plugging this into Equation (17) we 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: We 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 we 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: We 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: We 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 $\frac{\partial \bar{p}}{\partial R} < 0$. $\frac{\partial \bar{x}}{\partial \sigma_p^2} < 0$ and $\frac{\partial \bar{e}}{\partial \sigma_p^2} < 0$. $\frac{\partial \bar{p}}{\partial \rho} = - \left[\frac{\sigma_p^2}{\bar{e}} - \frac{\rho\sigma_p^2 \frac{\partial \bar{e}}{\partial \rho}}{\bar{e}^2} \right] \frac{S}{R-1}$ 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 $\frac{\partial \bar{e}}{\partial \sigma_p^2} < \frac{\bar{e}}{\sigma_p^2} \iff \frac{\sigma_p^2 A \rho^2}{2\bar{e}^3 e^{R\rho + \frac{\rho^2 \sigma_p^2}{2\bar{e}^2}} + 2A\rho^2 \sigma_p^2} < 1$

$\Leftrightarrow \frac{\sigma_p^2 A \rho^2 (-2\bar{e}^3 e^{R\rho + \frac{\rho^2 \sigma_p^2}{2\bar{e}^2}} - 1)}{2\bar{e}^3 e^{R\rho + \frac{\rho^2 \sigma_p^2}{2\bar{e}^2}} + 2A\rho^2 \sigma_p^2} < 0$ 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 we use to prove our 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. We 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 we 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: $\frac{da}{d\tau} = R^{-1} \left[a + \mu - \frac{S}{e_t} \rho \sigma_p^2 \right] - a$ which can be rewritten as $\frac{da}{d\tau} = a(R^{-1} - 1) + R^{-1}(\mu - \frac{S}{e_t} \rho \sigma_p^2)$. Furthermore, since e is a function of a , we 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 our proposition.

8 References

- [1] Adam, K., Marcet, A., Beutel, J. 2017. *Stock Price Booms and Expected Capital Gains*. American Economic Review.
- [2] Adam, K., Marcet, A., Nicolini, J. 2016. *Stock Market Volatility and Learning*. The Journal of Finance
- [3] Allen, F., Gale, D. 1994. *Limited Market Participation and Volatility of Asset Prices*. The American Economic Review.
- [4] Arrondel, L., Calvo-Pardo, H., Tas, D. 2014. *Subjective Return Expectations, Information and Stock Market Participation: Evidence from France*. Working paper.
- [5] Babcock, B., Choi, E., Feinerman, E. 1993. *Risk and Probability Premiums for CARA Utility Functions*. Journal of Agricultural and Resource Economics.
- [6] Baker, M., Wurgler, J. 2000. *The Equity Share in New Issues and Aggregate Stock Returns*. The Journal of Finance.
- [7] Branch, W., Evans, G. 2006. *A Simple Recursive Forecasting Model*. Economic Letters.

- [8] **Branch, W., Evans, G.** 2011. *Learning about Risk and Return: A Simple Model of Bubbles and Crashes*. American Economic Journal: Macroeconomics.
- [9] **Branch, W., Evans, G.** 2013. *Bubbles, Crashes, and Risk*. Economics Letters.
- [10] **Bertaut, C., Starr-McCluer, M.** 2002. *Household Portfolios in the United States*. in Guiso, Haliassos, Jappelli, eds.: *Household Portfolios*.
- [11] **Campbell, J.** 2006. *Household Finance*. The Journal of Finance.
- [12] **Evans, G., Honkapohja, S.** 1995. *Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models*. Econometrica.
- [13] **Evans, G., Honkapohja, S.** 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- [14] **Fagereng, A., Gottlieb, C., Guiso, L.** 2017. *Asset Market Participation and Portfolio Choice Over the Life-Cycle*. The Journal of Finance.
- [15] **Gomes, F., Michaelides, A.** 2005. *Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence*. The Journal of Finance.
- [16] **Guiso, L., Haliassos, M., Jappelli, T.** 2002. *Household Portfolios*. MIT Press.
- [17] **Guiso, L., Jappelli, T.** 2002. *Household Portfolios in Italy*. in Guiso, Haliassos, Jappelli, eds.: *Household Portfolios*.
- [18] **Guiso, L., Sodini, P.** 2013. *Household Finance: An Emerging Field*. The Handbook of the Economics of Finance.
- [19] **Lucas, R.** 1978. *Asset Prices in an Exchange Economy*. Econometrica.
- [20] **Malmendier, U., Nagel, S.** 2011. *Depression Babies: Do Macroeconomic Experiences affect Risk-taking?* The Quarterly Journal of Economics.
- [21] **Malmendier, U., Nagel, S.** 2015. *Learning From Inflation Experiences*. The Quarterly Journal of Economics.
- [22] **Mankiw, G., Zeldes, S.** 1991. *The Consumption of Stockholders and Nonstockholders*. Journal of Financial Economics.
- [23] **Marcet, A., Sargent, T.** 1989. *Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models*. Journal of Economic Theory.
- [24] **Merton, R.** 1971. *Optimum Consumption and Portfolio Rules in a Continuous-Time Model*. Journal of Economic Theory.
- [25] **Nakov, A., Nuño, G.** 2015. *Learning from Experience in the Stock Market*. Journal of Economic Dynamics and Control.
- [26] **Orosel, G.** 1998. *Participation Costs, Trend Chasing, and Volatility of Stock Prices*. The Review of Financial Studies.
- [27] **Rogerson, R.** 1988. *Indivisible Labor, Lotteries, and Equilibrium*. Journal of Monetary Economics.
- [28] **Sargent, T.** 1993. *Bounded Rationality in Macroeconomics*. Oxford University Press.
- [29] **Sargent, T., Ljungqvist, L.** 2011. *A Labor Supply Elasticity Accord?* American Economic Review.

- [30] **Timmermann, A.** 1993. *How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices.* The Quarterly Journal of Economics.
- [31] **Vissing-Jorgensen, A.** 2002. *Limited Asset Market Participation and the Elasticity of Intertemporal Substitution.* Journal of Political Economy.