

Imperfect Knowledge, Liquidity and Bubbles

William A. Branch
University of California, Irvine

July 6, 2012
Preliminary and Incomplete

Abstract

This paper demonstrates that insufficient liquidity, in the form of a shortage of safe assets that are useful as collateral in facilitating exchange, can lead to substantial movements in asset prices. There is a single asset that yields a positive payoff stream and can be traded in a centralized market. The asset can also be used to facilitate exchange in decentralized, or over-the-counter, trade and if the asset is in sufficiently short supply the fundamental asset price includes a liquidity premium. Traders, though, have imperfect information about the future price at which the asset will trade and so they behave like a Bayesian who estimates an econometric forecasting model for the asset price that is updated in real-time via discounted least-squares. The paper has three primary results: first, a permanent decrease in the supply of safe assets can lead to a substantial over-shooting of the asset price from its fundamental value; second, an increase in collateral requirements can lead to a substantial over-shooting of asset prices; third, when asset prices include a liquidity premium there can be recurrent bubbles and crashes that arise as endogenous responses to economic shocks.

The world has a shortage of financial assets. Asset supply is having a hard time keeping up with the global demand for store of value and collateral by households, governments...The equilibrium response of asset prices and valuations to these shortages has played a central role in ... the recurrent emergence of speculative bubbles, the historically low real interest rates ... all fall into place once one adapts this asset shortage perspective.

– Ricardo J. Caballero in “On the Macroeconomics of Asset Shortages.”

The shrinking set of assets perceived as safe ... can have negative implications for global financial stability. It will increase the price of safety and ... could lead to more short-term volatility jumps, herding behavior, and runs on sovereign debt.

–I.M.F. Global Financial Stability Report April 2012.

1 Introduction

This paper studies asset pricing in an environment where the asset has a dual role as a store-of-value and in providing liquidity services. The liquidity role is formalized using search frictions that are familiar in monetary theory. In a stationary equilibrium, the asset carries a liquidity premium, above its discounted payment flows, whenever the supply of the asset is not sufficient to support all of the trade in over-the-counter transactions. However, the liquidity premium in the model does not generate the kinds of price dynamics – such as bubbles, crashes and excess volatility – that are typically observed in practice. This paper proposes a search-based asset pricing model with imperfect knowledge and adaptive learning as a means of generating bubbles and crashes in asset prices.

It has long been recognized that financial assets have important roles beyond a store-of-value including the provision of liquidity services. Assets that can be considered safe are increasingly viewed to be in short supply: safe, liquid assets are used as collateral in over-the-counter transactions and bilateral agreements while a rise in global demand by investors, governments and central banks, and changes to macro-prudential policies, are likely to exacerbate supply imbalances. When financial assets play a similar liquidity role as money, variations in the supply of assets can affect asset prices. For example, Krishnamurthy and Vissing-Jorgensen (2012) show that changes in the (relative) supply of treasury debt, corporate and agency bonds affect the price of these assets. Caballero, Farhi, and Gourinchas (2008) link global capital flows to a shortage in the supply of assets. Holmstrom and Tirole (2011) highlight a possible role for insufficient liquidity in the subprime crisis. Recently, the I.M.F.'s Global Financial Stability Report (2012) predicts that imbalances in the supply of safe assets could lead to bubbles, crashes, and excess volatility in asset prices.

In a frictionless environment, an asset's price should reflect the discounted, expected payment flows and be independent of changes in the supply of assets. A recent literature, building on insights from monetary theory, demonstrates that assets can carry a liquidity premium when overcoming trading frictions (such as limited

enforcement and lack of record keeping) that prevent the use of unsecured credit.¹ Search-and-matching models of asset pricing are useful environments for studying prices in economies with a shortage of safe assets since they make explicit the liquidity properties of assets: safe assets can serve as collateral to facilitate bilateral, or over-the-counter, trade when limited enforcement or imperfect recognizability preclude unsecured credit arrangements as a means of payment. In search models, when the amount, or supply, of the safe asset is sufficiently low, the asset price will reflect its dual roles as a store of value and as a provider of liquidity services. The component of the asset price attributable to a liquidity premium, is sometimes referred to as a rational bubble.² Although, search-based models have been useful in explaining certain empirical properties of asset prices, such as the risk-free rate and equity premium puzzles (see, Lagos (2010)), to date, they have not been successful in generating other salient features of asset prices such as the rapid price appreciations and depreciations typically attributed to speculative bubbles.

This paper presents a search-based asset pricing model that is capable of generating asset price bubbles and crashes. The economic environment is based on Nosal and Rocheteau (2011) and Rocheteau and Wright (2011): there is a single asset, similar to a Lucas Tree, that pays an i.i.d. dividend and is traded in a centralized, competitive market. The supply of this asset is subject to occasional, small iid shocks that captures asset float and other exogenous factors that affects asset supply. Absent trading frictions, this asset would price at the discounted present value of the dividend flow. However, the economy also consists of a decentralized market where buyers and sellers are bilaterally matched and buyers submit to sellers a “buyer-takes-all” offer. Unsecured credit is not available in these pairwise meetings because of limited enforcement. Instead, the “safe” asset can serve as collateral for secured credit giving rise to an endogenous liquidity role for financial assets. In a stationary (rational expectations) equilibrium, the asset price consists of two components: the expected present-value of future dividends and a liquidity premium. The liquidity premium arises only when the supply of the asset is too small to support the efficient level of bilateral trade.

In the model, the asset price is determined, in part, by the expected future price of the asset. The departure point of this paper is to replace rational expectations with

¹See Nosal and Rocheteau (2011) for an extensive survey of search-based monetary and asset-pricing theory. Search-based models of asset pricing and liquidity include Duffie, Garleanu, and Pedersen (2005), Geromichalos, Licari, and Suarez-Lledo (2007), Lagos (2010), Lagos and Rocheteau (2009), Weill (2008), Lagos and Wright (2005), Lester, Postlewaite, and Wright (2012), Rocheteau and Wright (2011).

²There is a long history of interpreting fiat monetary equilibria as a rational bubble and extending that interpretation to assets, more broadly, since fiat money is an asset with a constant, zero payment forever. See, Tirole (1985).

price expectations formed from an adaptive learning rule as in Evans and Honkapohja (2001).³ The imperfect knowledge environment under consideration assumes that individuals understand a lot about the economic environment, but they do not know – or harbor some doubt about – the particular values of the dividend process, the asset supply process and other values/coefficients that determine asset prices. As a result, individuals draw inferences about the asset price process from recent data by adopting an econometric forecasting model whose reduced-form nests the rational expectations equilibrium. These agents are Bayesian and, because of uncertainty about their model, they place a prior on structural change in their econometric model. This imperfect knowledge framework implies that individuals forecast via an AR(1) econometric model whose parameters are updated in real time with a form of discounted least squares (“constant gain learning”). The priors for this Bayesian model are specified in such a manner that beliefs are, on average, close to rational expectations. We show that the dynamic properties of an economy with a shortage of safe assets are altered in interesting ways once agents must adaptively learn about the price process.

There are several channels through which imperfect knowledge and adaptive learning affects asset prices. First, although over time beliefs tend to converge toward rational expectations, the combination of constant gain learning and a positive liquidity premium can lead individuals to temporarily believe that asset prices follow a random walk without drift. Under these beliefs, individuals will interpret recent innovations to price as permanent shifts in the long-run value of the asset, the resulting increase in asset demand will lead to higher asset prices. Random walk beliefs arise for a very intuitive reason. Imperfect knowledge about the price process lead individuals to estimate the mean asset price from historical data. As a thought experiment, suppose there is a slight (temporary) upward drift to asset prices. Individuals’ econometric models will pick up that drift, leading to higher expectations about future asset prices that feed back onto higher asset prices. This speculative bubble-like dynamic is self-reinforcing and in some cases can lead individuals in the market to believe that asset prices follow a random walk.

Second, random-walk beliefs, as will be shown below, are nearly self-fulfilling and, consequently, such beliefs tend to persist for a substantial length of time. Furthermore, these beliefs generate excess volatility in asset prices, characterized by significant bursts and collapses in asset prices that are reminiscent of speculative bubbles and crashes. During a bubble episode, buyers demand greater amounts of the asset, whose supply is exogenous and in short supply, and because of the higher anticipated return to the asset it becomes more liquid in over-the-counter markets as sellers are willing to part with more goods in exchange for the asset; a bubble leads to greater

³Baranowski (2012) is the first paper to study learning in a monetary search model.

economic activity. A collapse of the bubble, or an asset price crash, will often exhibit price dynamics that substantially under-shoot the long-run price. During a crash episode, the asset becomes less liquid and, as a result, there is less economic activity. Importantly, these changes in the liquidity property of the assets along a bubble or crash path arise as an endogenous response to the fundamental economic shocks.

Third, a decline in the quantity of safe assets will introduce just the type of drift in asset prices that can lead to random-walk beliefs and cause a substantial overshooting of the new stationary equilibrium price. In one numerical example, the asset price will overshoot the new equilibrium price by nearly 100% before collapsing and converging at its new long-run value. These asset price dynamics confirm the intuition of Caballero (2006) and the I.M.F.'s Global Financial Stability Report. Fourth, structural changes to the economic environment that increase the demand for collateral can also lead to a drift in asset prices that lead to random-walk beliefs and an overshooting of the equilibrium price. This final result can be interpreted in the context of recent macro-prudential policies in the Dodd-Frank Act and the Basel III accord that requires more third-party clearing of bilateral or over-the-counter transactions. Transactions that are run through clearinghouses typically require more collateral, and less unsecured credit, than purely bilateral transactions between buyers and sellers. We interpret the greater demand for collateral in terms of increased frequency of trade in the decentralized market and demonstrate that such a structural change in the trading environment can lead to an asset price bubble.

The results in this paper relate to a recent, and growing, literature that estimates time series variation in liquidity premia. The model presented in this paper generates bubbles, crashes and excess volatility because imperfect knowledge and adaptive learning lead to significant swings in the liquidity premium that arise in over-the-counter markets. Dick-Nielsen, Feldhutter, and Lando (2012) find that excess volatility in estimated liquidity premia for investment grade corporate debt, and that this estimated premium grew substantially after the onset of the financial crisis. Their results can be interpreted in the context of the present model as a decline in the supply of safe assets. Moreover, Bao and Pan (2012) find excess volatility in monthly bond and CDS returns that result from variation in the illiquidity of bonds. Most closely related to the results here, they find that bond prices deviate from fundamental prices because of variation in the liquidity premium in over-the-counter markets.

Is it reasonable to assume that individuals might have imperfect knowledge about the asset price process? The answer is yes, for a variety of reasons. First, we adhere to the *cognitive consistency principle*, as articulated by Sargent (1993) and Evans and Honkapohja (2013), that economic agents should be assumed to behave like a good econometrician who forecasts future economic variables using time-series econometric

methods. This approach is reasonable since neither individuals or economists know the model and instead formulate and estimate models that are frequently revised in light of new data. Second, it is plausible to assume that the economy is subject to occasional structural change – such as a decline in the supply, or an increase in the global demand for safe assets – that is only revealed after some time has elapsed and a sufficient quantity of data bears evidence of the change. Third, the imperfect knowledge assumption assumes agents know the form of the law of motion for the state variables and only attempt to estimate the true parameter value by adjusting their estimates in light of recent data. On average, their beliefs are close to the rational expectations equilibrium values and are determined endogenously with the state variables, thereby, preserving the cross-equation restrictions that are a salient feature of equilibrium models.

An important feature of the analysis is that imperfect knowledge can generate instability in asset prices even though the departure from rational expectations is, on average, small. The framework employed here is related to an extensive literature that employs adaptive learning in macroeconomics. Most closely related are papers that incorporate constant gain learning in studies of monetary policy and asset pricing; see, for example, Branch and Evans (2011); Sargent (1999); Adam, Marcet, and Nicolini (2010); Orphanides and Williams (2005); Cho, Williams, and Sargent (2002); Williams (2004); Cho and Kasa (2008); McGough (2006).) Branch and Evans (2011), in particular, find that risk-averse agents in an OLG asset pricing model forecast both the risk and return of stock prices using a forecasting model whose parameters are updated using constant gain least squares then traders may also come to believe that stocks follow a random walk. These nearly self-fulfilling random walk beliefs lead to recurrent bubbles and crashes in stock prices. While there are similarities in the mechanism generating bubbles in Branch and Evans (2011) and in this paper, there are important differences. The present environment includes a liquidity services role for assets and this liquidity demand is essential for generating bubbles. While Branch and Evans (2011) emphasizing variations in the perceived riskiness of assets.

There are alternative explanations for bubbles and variations in liquidity premia. Rocheteau and Wright (2011), using a very similar search-based model, allow for endogenous firm entry that depends, in part, on the asset price. They demonstrate the possibility for multiple equilibria, and cycling between those equilibria, to generate bubbles and crashes. Guerrieri and Shimer (2012) generate endogenous illiquidity from an adverse selection problem that leads sellers of high quality, safe assets to be unwilling to sell if prices are too low. A closely related economic environment is in Kiyotaki and Moore (2008) who generate endogenous collateral constraints. A number of papers demonstrate the efficiency enhancing properties of bubbles in overlapping generations models, including Tirole (1985), Grandmont (1985), and Santos

and Woodford (1997). Bubbles can arise from agency problems in Allen and Gale (2000), Barlevy (2011) and Farhi and Tirole (2011). This paper is complementary to these, and other related papers, and instead emphasizes the liquidity role of assets in an imperfect knowledge environment where endogenous expectations arise from a real-time adaptive learning rule.

This paper proceeds as follows. Section 2 describes the economic environment. Section 3 illustrates the main economic mechanism behind the results. Section 4 presents the main results. Section 5 includes further discussion. Section 6 concludes.

2 A Search-based Asset Pricing Model with Imperfect Knowledge

2.1 Environment

The economic environment is adapted from Rocheteau and Wright (2011) and Rocheteau and Wright (2005).⁴ Each time period consists of two subperiods: in the first subperiod agents gather in a decentralized, or over-the-counter, market, where buyers and sellers meet in bilateral matches and exchange specialized goods; in the second subperiod, agents interact in a centralized market where each agent is free to consume and produce a general good using a linear production technology and trade in a financial asset (claims to a Lucas tree). Following Rocheteau and Wright (2005) agents are heterogeneous with a continuum of buyers and sellers, each with measure one, in the decentralized market. The financial assets pay an iid dividend y_t to holders of the asset and have an exogenous supply A_t . This asset can be interpreted as a “safe” asset as all individuals in the economy have perfect information about the (stochastic) dividend process. There is a stochastic process for A_t meant to capture exogenous variation in the supply of safe assets. This variation could also be considered “asset float” – i.e., IPO lock-up expirations, stock splits, repurchase agreements, etc. – changes in the supply of safe government debt, or an increase in global demand for safe assets that is external to the economy.⁵ Changes in asset supply are implemented via lump-sum transfers to buyers at the beginning of the centralized market. The exogeneity of the asset supply process is made for technical

⁴Rocheteau and Wright (2011) extend the monetary framework of Rocheteau and Wright (2005) and Lagos and Wright (2005) to a model where the medium of exchange are claims to a Lucas tree and there is endogenous firm entry.

⁵An alternative interpretation of the stochastic component to A_t is that only a fraction of the asset can be used as collateral in the decentralized market (see Kiyotaki and Moore (2008)).

convenience.

The model without decentralized trading is equivalent to the frictionless, risk-neutral Lucas asset pricing model. With all trade taking place in the centralized market the asset will be priced as the discounted expected capital return, reflecting the “store of value” role of assets. Under rational expectations, the (unique) equilibrium asset pricing sequence has the property that the asset price equals the present value of future expected dividends. With decentralized trading trade takes place through an over-the-counter market with bilateral meetings and bargaining. Limited enforcement, or a lack of commitment, precludes the use of unsecured credit arrangements. Instead, buyers purchase goods with financial assets as a means of payment or, equivalently, with the assets serving as collateral to secure loans extended by sellers. This friction captures the role “safe” assets can play in facilitating trade. This formalizes the liquidity properties of the financial asset. When the supply of assets is in short supply, i.e. does not provide sufficient liquidity to facilitate the efficient level of trade, then the fundamental price of the asset will carry an additional liquidity premium. This paper demonstrates that the liquidity premium is essential for generating the departures from rational expectations that leads to bubbles, and subsequent crashes, in asset market prices.

2.2 Model

Buyers choose sequences of the generalized good, the specialized good, labor, and asset holdings to maximize

$$\hat{E}_0 \sum_{t=0}^{\infty} \beta^t (U(x_t) + u(q_t) - l_t)$$

subject to the constraints

$$x_t + p_t a_t = l_t + (p_t + y_t) a_{t-1} + T_t$$

where x_t is the generalized good, q_t is the differentiated good purchased in the decentralized market, l_t is labor hours used to produce the generalized good according to the production function $x_t = l_t$, a_t are asset holdings traded at price p_t , y_t is a (stochastic) dividend, and T_t are the lump-sum transfers that distributes, without loss of generality, the changes in asset supply to buyers. \hat{E} is the (possibly) non-rational expectations operator (to be specified below). As is standard in these models, define $U'(x^*) = 1$ and let $U(x^*) = x^*$. For simplicity, assume that the specialized good is produced in the decentralized market according to the cost function $c(q) = q$. The

efficient quantity of the decentralized good is, therefore, $u'(q^*) = c'(q^*) = 1$. Finally, assume that agents rank alternative bundles of the specialized good by $u(q) = \frac{q^{1-\sigma}}{1-\sigma}$.⁶

The decentralized, or over-the-counter, market operates as follows. Buyers and sellers meet in bilateral matches where the probability of a buyer and seller meeting is given by the constant probability α . The parameter α captures the search friction present in over-the-counter markets. Because of a lack of commitment, or imperfect credit enforcement, trade involves issuing debt backed by collateral in the form of claims to the asset, or equivalently, a *quid pro quo* transfer of shares in the asset.

The search and bargaining friction highlights two specific roles of the asset: as an asset that yields a capital return in the form of $p_t + y_t$ and as a liquid instrument that enables trade that would otherwise not occur. Once matched a buyer and seller bargain over the terms of trade. To keep the analysis simple, it is assumed that buyers make a “take it or leave it” offer.⁷

The timing of the model assumes that buyers make asset holding decisions during the centralized market and these assets can be used for trade in the following period. Thus, asset demand will depend critically on expectations about the future price of the asset. Individuals must have full information about the distribution of the endogenous variables in order to form rational expectations. An alternative to rational expectations is to assume that individuals behave like econometricians who hold a (correctly) specified model of the economy, but they must recover the parameters in real time from data. The imperfect knowledge assumption in this paper builds on this approach.

Of course, along a learning path, the decisions that individuals make will be informed, in part, by their forecasting model for price; as beliefs adjust, so will the decisions made by agents. Because of the close interaction between beliefs and individuals’ choices, it is important to be clear about several key assumptions regarding the timing of decisions, outcomes, and the updating of beliefs. When solving the buyers’ intertemporal optimization problem a central assumption regards the extent to which individuals take the future evolution of beliefs into account when forming their expectations. An individual that recognizes that the learning process will imply

⁶With search frictions, with positive probability there may not be trade and the utility function $u(q)$ may not be defined at zero. The literature deals with this possibility typically by altering the utility function to be $u(q) = \frac{(q+b)^{1-\sigma} - b^{1-\sigma}}{1-\sigma}$ so that it is defined at $q = 0$. For the present paper, it is sufficient to assume that $u(q)$ is locally CRRA, but defined at $q = 0$. Therefore, throughout we set b arbitrarily close to zero.

⁷The qualitative results in this paper do not hinge on the specifics of the bargaining between buyer and seller. What is needed is that the asset is in sufficiently short supply that it carries a premium above its discounted payment flow.

a probability distribution over all possible sequences of future beliefs, and takes this uncertainty into account when formulating decision rules, is a fully Bayesian decision-maker. On the contrary, an individual that takes his/her beliefs as given – that is, that the coefficients in the econometric model are fixed and will not be updated again in the future – when calculating the decision rule is called an anticipated utility maximizer (see Kreps (1998) and Cogley and Sargent (2008)). The present environment assumes an anticipated utility framework for decision making. This assumption is made for technical convenience and because it strikes us as a more plausible description of individuals’ behavior when they take actions conditional on their beliefs. An individual that maximizes an infinite horizon optimization problem by accounting for all of the potential sequences of beliefs will be faced with a more complicated task than even agents with rational expectations. We note, however, that the qualitative results below do not hinge on the anticipated utility assumption. A second assumption that we make is that all individuals (buyers and sellers) are assumed to have the same beliefs: they observe the same information, have the same learning rule, and its common knowledge that they hold identical expectations. Finally, it is also assumed that beliefs and endogenous state variables are not determined simultaneously. In particular, we assume that that current price of the asset is not observable when agents form expectations. This breaks the simultaneity of beliefs and outcomes that is a feature of rational expectations models but is not consistent with agents who form expectations as out-of-sample forecasts given the available data. Define $\Omega_t = \left(\{p_j\}_{j=0}^{t-1}, \{A_j, y_j\}_{j=0}^t \right)$ as the information set available to agents when they make decisions at time t .

In the centralized market, individuals hold expectations about the next period’s price, conditional on all of the previously realized state variables, before they observe the current price. Thus, we interpret the optimization problem as determining a demand schedule that the agents turn into the Walrasian auctioneer in the centralized market and the auctioneer sets the price to clear the market. More specifically, the timing is as follows:

- At the beginning of the decentralized and centralized markets in time t , buyers and sellers hold expectations $\hat{E}[p_t + y_t | \Omega_t]$, and y_t, A_t are observable but $p_t \notin \Omega_t$.
- At the beginning of the centralized market, buyers submit their asset demand schedule to the auctioneer based on $\hat{E}[p_{t+1} + y_{t+1} | \Omega_t]$. The auctioneer clears the market.
- At the end of period t , after observing the realized price p_t , buyers and sellers update their information set Ω_{t+1} to include p_t .

To solve the model for an equilibrium asset price, we proceed sequentially beginning with the bargaining solution in the decentralized market. During the decentralized market, at the beginning of time t , a buyer makes a take-it-or-leave-it offer in the form of a pair (q_t, d_t) that specifies the exchange of q units of the good in exchange for d units of the asset.⁸ This offer solves

$$(q_t, d_t) = \arg \max_{q_t, d_t} \left[u(q_t) + \hat{E} [p_t + y_t | \Omega_t] (a_{t-1} - d_t) - \hat{E} [p_t + y_t | \Omega_t] a_{t-1} \right]$$

subject to the seller's participation constraint

$$-q_t + \hat{E} [p + y | \Omega_t] d_t \geq 0$$

The buyer's offer maximizes his surplus from an offer where the term $\hat{E} [p_t + y_t | \Omega_t] (a_{t-1} - d_t) - \hat{E} [p + y | \Omega_t] a_{t-1}$ is the expected consumption foregone in the centralized market after transferring d units of the asset to the seller. Because p_t , the price of the asset in the centralized market in time t , is not contemporaneously observable, the bargaining terms between buyer and seller depend on their (possibly) non-rational expectations of the value of the asset. The seller will participate so long as the anticipated consumption in the centralized market, $\hat{E} [p + y | \Omega_t] d_t$, is greater than their cost of producing q_t . This is how learning and beliefs can affect liquidity. The solution to this bargaining problem is

$$q_t = \begin{cases} q^* & \text{if } \hat{E} [p_t + y_t | \Omega_t] a_{t-1} > q^* \\ \hat{E} [p_t + y_t | \Omega_t] a_{t-1} & \text{else} \end{cases}$$

If buyers have sufficient holdings of the asset they purchase the efficient quantity q^* , otherwise they turn over all of their holdings of the asset and receive $\hat{E} [p_t + y_t | \Omega_t] a_{t-1}$ in return.

The value function for a buyer in the decentralized market, given (q_t, d_t) is given by the expression

$$V_t(a_{t-1}) = \alpha [u(q_t) + W_t(a_{t-1} - d_t)] + (1 - \alpha)W_t(a_{t-1})$$

where W_t is the value function for a buyer in the centralized market:

$$W_t(a) = \max_{x_t, l_t, a_t} u(x_t) - l_t + \beta \hat{E} [V_{t+1}(a_t) | \Omega_t]$$

⁸The take-it-or-leave-it offer is a special case of proportional bargaining, where the buyer captures the entire surplus from trade. Proportional bargaining has certain theoretical properties, such as surpluses that increase along with the bargaining set, that are more attractive than other forms of bargaining such as Nash bargaining. See Nosal and Rocheteau (2011) for details. The qualitative results do not hinge on the proportion of the surplus assigned to the buyer.

Combining these expressions, and making use of the quasi-linearity, leads to the following equation

$$W_t(a_{t-1}) = (p_t + y_t)a_{t-1} + T_t + \max_{x^*} [U(x^*) - x^*] \\ + \max_{a_t \geq 0} \left\{ -p_t a_t + \beta \hat{E} [\alpha \{u(q_{t+1}) - (p_{t+1} + y_{t+1})d_{t+1}\} + (p_{t+1} + y_{t+1})a_t | \Omega_t] \right\}$$

With these assumptions, asset demand a_t is the solution to

$$\max_{a_t \geq 0} \left\{ \left(\beta \hat{E}_t [p_{t+1} + y_{t+1} | \Omega_t] - p_t \right) a_t \right. \\ \left. + \max_{d_{t+1} \in [0, a_t]} \alpha \beta \left[u(\hat{E} [q_{t+1} | \Omega_t]) - \hat{E} [p_{t+1} + y_{t+1} | \Omega_t] d_{t+1} \right] \right\} \quad (1)$$

To derive this expression for asset demand, we impose that (i.) agents use point expectations, i.e. $\hat{E} [u(q_{t+1}) | \Omega_t] = u(\hat{E} [q_{t+1} | \Omega_t])$, and, (ii.) expectations obey a law of iterated expectations, i.e. $\hat{E} [\hat{E} (z | \Omega_{t+1}) | \Omega_t] = \hat{E} [z | \Omega_t]$ for any variable z . That agents use point expectations is a behavioral assumption that essentially holds that decisions only depend on the mean of their subjective beliefs. This assumption is made for technical convenience and is a standard restriction imposed in many rational expectations and adaptive learning models (see Evans and Honkapohja (2001)). The law of iterated expectations, as expressed in (ii.), is a consequence of the anticipated utility framework. This assumption does not impact the main results.

Notice that in (1), when $\alpha = 0$, i.e. there is no decentralized market, the buyer's demand for the asset is equivalent to the risk-neutral Lucas asset pricing model. The first expression in (1) shows that a part of the demand for the asset depends on the expected return on the asset. The second expression is the liquidity demand for the asset and here it depends on the expected surplus from trading in the decentralized market.

There are three cases to consider:

1. When $\beta \hat{E} [p_{t+1} + y_{t+1} | \Omega_t] > p_t$, then households desire an infinite amount of the asset and the optimization problem does not have a solution.
2. When $\beta \hat{E} [p_{t+1} + y_{t+1} | \Omega_t] = p_t$ then households hold enough to purchase $q^* = 1$, $d_{t+1} = 1/\hat{E} [p_{t+1} + y_{t+1} | \Omega_{t+1}]$, and any $a_t \geq d_{t+1}$ is a solution to the optimization problem. In this case, there is no liquidity premium and the asset is priced as the discounted expected payment flow of the asset.

3. When $\beta \hat{E} [p_{t+1} + y_{t+1} | \Omega_t] < p_t$, then the household is liquidity constrained and $q_{t+1} = \hat{E} [p_{t+1} + y_{t+1} | \Omega_{t+1}]$, and a_t solves

$$\max_{a_t} \left\{ \left(\beta \hat{E}_t [p_{t+1} + y_{t+1} | \Omega_t] - p_t \right) a_t + \alpha \beta \left[u(\hat{E} [p_{t+1} + y_{t+1} | \Omega_t] a_t) - \hat{E} [p_{t+1} + y_{t+1} | \Omega_t] a_t \right] \right\}$$

The first-order condition from the buyer's problem combined with a market clearing condition yields an expression for the equilibrium price.⁹

It remains to specify the stochastic processes for dividends and for the supply of the asset. For simplicity, and without loss of generality, assume that dividends follow the process $y_t = y + \eta_t$ where $y > 0$ and η_t is white noise with variance σ_η^2 . From a theoretical perspective, and for the learning results presented below, the precise details of the process followed by y_t are not important. Since we are interpreting the asset as a "safe asset" it is natural to assume that y_t is known with certainty or subject to small iid shocks. Assume also that the supply of the asset is given by the process $\log A_t = \log A - \frac{1}{\sigma} \log \hat{\varepsilon}_t$ where $A > 0$ and $E \hat{\varepsilon}_t = 1$ with a small compact support. The stochastic process for the supply of shares is meant to proxy for exogenous changes in asset supply. Supply variation could arise because of changes in government debt issuance (as in Krishnamurthy and Vissing-Jorgensen (2012)) or asset float where the tradeable supply of shares may vary because of repurchase agreements, stock splits, lock-up expirations etc. Asset float has been shown to be an important factor in asset pricing (see Cochrane (2005), Baker and Wurgler (2000)). Because the stochastic component of the transfers are unpredictable, buyers will not anticipate receiving transfers in period $t + 1$ when deciding on their asset demand at time t . However, the variation in the outside supply of shares will affect the equilibrium price at time t through the market clearing condition.

With these assumptions in hand, it is straightforward to solve for the following equilibrium price

$$p_t = \begin{cases} (1 - \alpha) \hat{E}_t (p_{t+1} + y_{t+1}) + \alpha \beta \left[\hat{E}_t (p_{t+1} + y_{t+1}) \right]^{1-\sigma} A_t^{-\sigma} & \text{if } A_t < \frac{q^*}{\hat{E}_t(p_{t+1} + y_{t+1})} \\ \beta \hat{E}_t (p_{t+1} + y_{t+1}) & \text{else} \end{cases} \quad (2)$$

where we now make use of the simplifying notation $\hat{E}_t z = \hat{E} [z | \Omega_t]$. Recall, also that $\hat{E}_t y_{t+1} = y$. This law of motion for the equilibrium price can be written compactly as

$$p_t = G(\hat{E}_t p_{t+1}, A_t) \quad (3)$$

⁹The liquidity premium implies that there is a holding cost to the asset. Since sellers do not use the liquidity services of the asset, they will choose not to buy the asset in the competitive market.

2.3 Rational Expectations Equilibria

A *rational expectations equilibrium* is a sequence $\{p_t\}$ that is a (bounded) solution to (2). Because of the non-linear nature of the expectational difference equation (2) a complete characterization of the set of equilibria is not possible. In the deterministic version of this model, for some parameter values, there can exist cycles, complicated dynamics and sunspot equilibria. (see Lagos and Wright (2005) and Rocheteau and Wright (2011)). However, it is straightforward to verify that there exist solutions to (2) that take the form of a noisy steady-state. In the analysis below, the model will be parameterized so that it is locally determinate.

The expression in (2) also demonstrates that asset prices will include a liquidity premium when $\alpha > 0$ and the supply of the asset A is sufficiently low so that there is an inefficient quantity traded in the decentralized market. These are rational bubbles much like monetary equilibria are bubbles or as in rational bubbles in Tirole (1985) which arise in dynamically inefficient economies. Because the liquidity premium arises out of a fundamental property of the asset – that is, its ability to facilitate bilateral exchange – we refer to this as the fundamental price.

Definition 1 *The “fundamental,” or stationary, equilibrium price is the steady-state $\bar{p} = G(\bar{p}, A)$.*

Remark. Of course, when A is sufficiently high (or $\alpha = 0$) then there is no liquidity premium and $\bar{p} = \beta y / (1 - \beta)$, which is the expected present value of the dividend flow.

Definition 2 *A noisy steady-state rational expectations equilibrium is a function $p(A_t)$ defined so that $p(A_t) = G(\hat{p}, A_t)$ with \hat{p} such that $\hat{p} = EG(\hat{p}, A_t)$, where the expectation is taken with respect to the distribution of A_t .*

The following result is a direct application of a theorem in Evans and Honkapohja (1995).

Proposition 3 (Evans and Honkapohja (1995)) *Consider a family of distribution functions for A_t , indexed by α , with $F_\alpha(-\alpha) = 0$, $F_\alpha(\alpha) = 1$ and F_α (weakly) converges as $\alpha \rightarrow 0$ to $F_0(A) = 1$. Define $\hat{p}(\alpha) = EG(\hat{p}(\alpha), A_t(\alpha))$ and $\bar{p} = G(\bar{p}, A)$ is the fundamental steady-state. Then there exists a noisy steady-state $p(A_t) = G(\hat{p}(\alpha), A_t)$ with $\hat{p}(\alpha)$ arbitrarily close to \bar{p} , for sufficiently small α .*

In a noisy steady-state equilibrium, prices are small iid deviations from the fundamental, or stationary, equilibrium price. Even though, the rational expectations equilibrium features iid fluctuations around the fundamental price, under learning the iid fluctuations are sufficient to generate substantial, temporary departures from the fundamental price.

3 Liquidity and Beliefs

This section presents results on the types of beliefs that can arise in an equilibrium and along a typical learning path. The main insight is that imperfect knowledge can introduce (nearly) self-confirming serial correlation into the model that would not exist under full information.

3.1 Beliefs

Because there exist rational expectations equilibria that are noisy steady-states one possible learning rule would be to simply recursively estimate the sample mean of the asset price. Since it is not possible to rule out, in general, the existence of other stochastic equilibrium paths it is not reasonable to expect that agents will know the complete underlying economic structure and be able to form rational expectations. In response, many modelers assume that agents behave like a good Bayesian who holds priors about the perceived model of the economy and updates those priors as new data becomes available. This adaptive learning approach typically assumes that agents have a correctly specified model with unknown parameters and use a reasonable estimator to update their parameter estimates. In many environments, these beliefs converge to rational expectations.¹⁰

In practice, however, econometricians often misspecify their models. In particular, even though the actual data generating process may be non-linear, econometricians and professional forecasters typically estimate linear models such as vector autoregressive models. This section takes this approach seriously by imposing that agents form their expectations via a linear AR(1) model of the asset price. Although this forecasting model is misspecified, we will require that it be optimal within the class of linear forecasting rules. In a *stochastic consistent expectations equilibrium (SCE)* agents' forecasting model is optimal within the class of misspecified models, i.e. the optimal linear projection, so that, within the context of their perceived model,

¹⁰See Evans and Honkapohja (2001) for extensive treatment of adaptive learning and expectational stability.

they are unable to detect their misspecification. The projection parameters and the equilibrium stochastic process for asset prices are jointly determined so that a SCE preserves many of the cross-equation restrictions that are a salient feature of rational expectations models. It is important to note that, although the AR(1) model may be misspecified for some AR coefficients, the forecasting model nests the (unique) noisy steady-state equilibrium. Thus, an AR(1) model is a natural and reasonable assumption on agents' beliefs.

Specifically, assume that agents form expectations from the forecasting model

$$p_t = c + b(p_{t-1} - c) + \varepsilon_t$$

Thus,

$$\hat{E}_t p_{t+1} = c + b^2(p_{t-1} - c) \quad (4)$$

Plugging (4) into (2) leads to the actual law of motion, given by

$$p_t = G(y + c + b^2(p_{t-1} - c); A_t) \quad (5)$$

where $G(\hat{E}_t(p_{t+1} + y_{t+1}); A_t)$ is given by the expression (2). Linear beliefs, such as those in (4), can be justified when the non-linear environment is not completely known since agents would be unable to exploit the non-linear structure for the purpose of forecasts. Indeed, in a SCE, as will be seen below, agents are unable to detect their misspecification within the context of their perceived model. The next sections discuss, in detail, how the coefficients (c, b) are determined in equilibrium and how they are updated by the learning rule.

The forecasting model and beliefs in (4), formalize the nature of individuals' imperfect knowledge. These agents have considerable understanding of their economic environment. They have an imperfect understanding of the process that determines the market asset price, p_t , and they specify a linear econometric forecasting model that nests the noisy rational expectations equilibrium price (i.e. where $c = \bar{p}, b = 0$). Given these beliefs, they determine their optimal asset demand and bargain over terms of trade with sellers (who share the same beliefs) in the over-the-counter market.

3.2 Stochastic Consistent Expectations Equilibria

This subsection presents insights on the nature of beliefs in an SCE. Branch and McGough (2005) characterize an equilibrium where agents hold linear beliefs, as in (5), and the state variable follows a non-linear reduced-form, as in (2), such that the belief parameters c, b are linearly consistent with the associated equilibrium dynamics. To state a precise definition of a stochastic expectations equilibrium (SCEE), the

following adapts Branch and McGough (2005) to the present environment. Define the following notation: for any initial distribution λ_0 on a compact set, with the initial condition p_0 chosen with respect to this distribution, and for $t \geq 1$, let $\lambda_t(\lambda_0)$ be the unconditional distribution of p_t , and let $\Lambda_t(\lambda_0)$ be the unconditional joint distribution of (p_t, p_{t-1}) , as determined by (5).

Definition 4 *The triple $(\{p_t\}, c, b)$ is a stochastic consistent expectations equilibrium (SCE) provided the following hold:*

1. p_t is generated by (5);
2. there exists a unique distribution λ so that for initial distribution λ_0 , the distribution $\lambda_t(\lambda_0)$ converges weakly to λ ;
3. for any λ_0 , $\lim_{t \rightarrow \infty} E_{\lambda_t(\lambda_0)}(p_t) = c$ and $\lim_{t \rightarrow \infty} \text{corr}_{\Lambda_t(\lambda_0)}(p_t, p_{t-1}) = b$.

An SCE occurs when there is a unique distribution to which p_t converges weakly, for any initial condition, and the asymptotic mean and autocorrelation coincide with the beliefs of agents. It is in this sense that agents are unable to detect their misspecification within the context of their model as a check of regression residuals would not reveal any first order autocorrelation that would lead the forecaster to reject the econometric model.

Notice that if $\bar{p} = G(\bar{p}; A)$ is a steady-state of the model (2) then in an SCE $c = \bar{p}$. In an SCE, the mean asset price will coincide with the mean price under rational expectations. Thus, showing existence of an SCE is straightforward: if \bar{p} is a fixed point of G then the pair $(\bar{p}, 0)$ characterizes an SCE. Branch and McGough refer to an SCE with zero autocorrelation as a ‘trivial SCE,’ though in the present context it accords with the (locally unique) rational expectations equilibrium. Showing existence of non-trivial SCE is challenging and many of the sufficient conditions in Branch and McGough (2005) are violated in the present environment. In a linear model (which arises here when $\sigma = 1$), Hommes and Zhu (2011) show that a non-trivial SCE does not exist for iid stochastic shocks but do exist if the shocks are serially correlated. Moreover, Branch and McGough (2005) showed that non-trivial SCE, when $\bar{p} \neq 0$, will be unstable under learning.

Studying asset pricing properties in an SCE may nevertheless yield important intuition for understanding asset pricing dynamics under adaptive learning. In particular, numerical simulations under learning show that bubbles and crashes can arise when agents’ approximating model becomes close to a random walk, i.e. $b \approx 1$. This subsection and the next present an argument that the onset of such beliefs is intuitive and expected in this environment.

To illustrate the possible types of equilibria it is useful to define the map

$$T(b) = \lim_{t \rightarrow \infty} \text{corr}(p_t, p_{t-1})$$


as the asymptotic first-order serial correlation, given b , between p_t and p_{t-1} . If \bar{p} is a fixed point of (2) then an SCE is a pair (\bar{p}, \bar{b}) where $\bar{b} = T(\bar{b})$ is a fixed point of the map T . It is straightforward to numerically compute an SCE. Figure 1 plots two different examples of SCE. In each plot the parameters are set to the baseline parameterization in Table 1. The value of β accords with a 2% real interest rate, while α was chosen to match the velocity of collateral estimate in Singh (2011). All values of $\sigma < 2$ yield a determinate model. In the baseline parameterization $\sigma = 1.75$, though the subsequent analysis considers alternative values for σ . The values of A and y were chosen so that there was a liquidity premium in the fundamental price, and the value of σ_ε^2 ensures that the supply shocks are “small”. The two plots in Figure 1 differ by the utility curvature parameter σ . This is a key parameter governing how strongly the expectational feedback in the liquidity premium impacts asset prices. The right panel is for $\sigma = 1.75$ and the left panel plots $\sigma = 0.05$. Where the line $T(b)$ crosses the 45° line is a SCE. In the case of a small σ , so that there is a negative feedback in the liquidity premium, then there exists a unique SCE at $b = 0$, which corresponds to the fundamental equilibrium. For a value of $\alpha = 0.05$, where there is now positive expectational feedback in the liquidity premium, then there exists two SCE one with $b = 0$ and the other at $b = 1$.

Table 1: Baseline parameterization

β	σ	α	A	y	σ_ε^2
0.98	1.75	0.10	0.40	0.01	0.004

There are two quick conclusions to draw from Figure 1. First, for small values of σ there can exist an equilibrium with self-fulfilling serial correlation. The equilibrium with $b = 1$ has non-zero serial correlation, even though the shocks and rational bubbles equilibrium are iid, entirely because agents’ perceive serial correlation, that is reinforced through the self-referential property of the asset-pricing model. In the non-trivial SCE it is the case that $E_t p_{t+1} = p_{t-1}$ so that agents’ perceived model of the asset price process is a random walk. Second, we can expect that the $b = 0$ equilibrium will be stable under learning since at $b = 1$ the slope of $T(b)$ is greater than one. This is not unexpected as Branch and McGough (2005) showed that non-trivial SCE are unstable in models with a non-zero steady-state. The analysis below, however, will demonstrate that along a transition path to the stable SCE, the learning

Figure 1: Stochastic Consistent Expectations Equilibria. Left panel is for the baseline parameterization. Right panel sets $\beta = 0.99$, $A = 0.1$, $\sigma = 0.05$.



dynamics will bring beliefs close to the non-trivial SCE so that there can be (nearly) self-fulfilling serial correlation for finite stretches of time. This will be the key intuition for why bubbles and crashes arise in this model.

3.3 Learning

The previous subsection examined an equilibrium where the belief parameters (c, b) are linearly consistent with the sample mean and first-order autocorrelation coefficients of the actual data generating process. A natural question is whether individuals will learn to coordinate on a SCE. This question is addressed by assuming agents recursively estimate the coefficients of their (linear) perceived law of motion and use these estimates to form expectations. These expectations generate new data via the reduced form model (5), and agents use these new data to again update their estimates.

Given the perceived law of motion

$$p_t = c + b(p_{t-1} - c) + \varepsilon_t$$

the actual law of motion is given by (5). Given parameters c and b it is possible to define the map $T : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}^2$ as follows

$$\begin{aligned} T_c(c, b) &= \lim_{t \rightarrow \infty} Ep_t \\ T_b(c, b) &= \lim_{t \rightarrow \infty} corr(p_t, p_{t-1}) \end{aligned}$$

and $T(c, b)' = (T_c(c, b), T_b(c, b))$. The map T can be interpreted as follows: given fixed beliefs (c, b) , the actual law of motion is given by (5) and the corresponding asymptotic mean and first-order autocorrelation is given by $T(c, b)$. Unsurprisingly, an equilibrium is a fixed point of the T-map.

Under real-time learning the parameters (c, b) are not fixed, and instead are adjusted gradually over time using least-squares to update their values in response to the changing data. The learning literature, e.g. Evans and Honkapohja (2001), has shown that the T-map can be used to compute a stability condition, known as E-stability, which often governs whether or not equilibrium parameters are locally stable under learning and that the differential equation, used to define E-stability, also provides information on the global dynamics under learning. The mathematical theorems underlying the E-stability principle rely on the stochastic approximation approach, and those theorems could be applied to the present non-linear environment. However, the form of the T-map is sufficiently complicated that general results are not available. It is possible to numerically solve for the E-stability dynamics and present available analytic results for the special case $\sigma = 1$.

The E-stability principle states that locally stable rest points of the ordinary differential equation

$$\frac{d(c, b)'}{d\tau} = (T(c, b) - (c, b))' \quad (6)$$

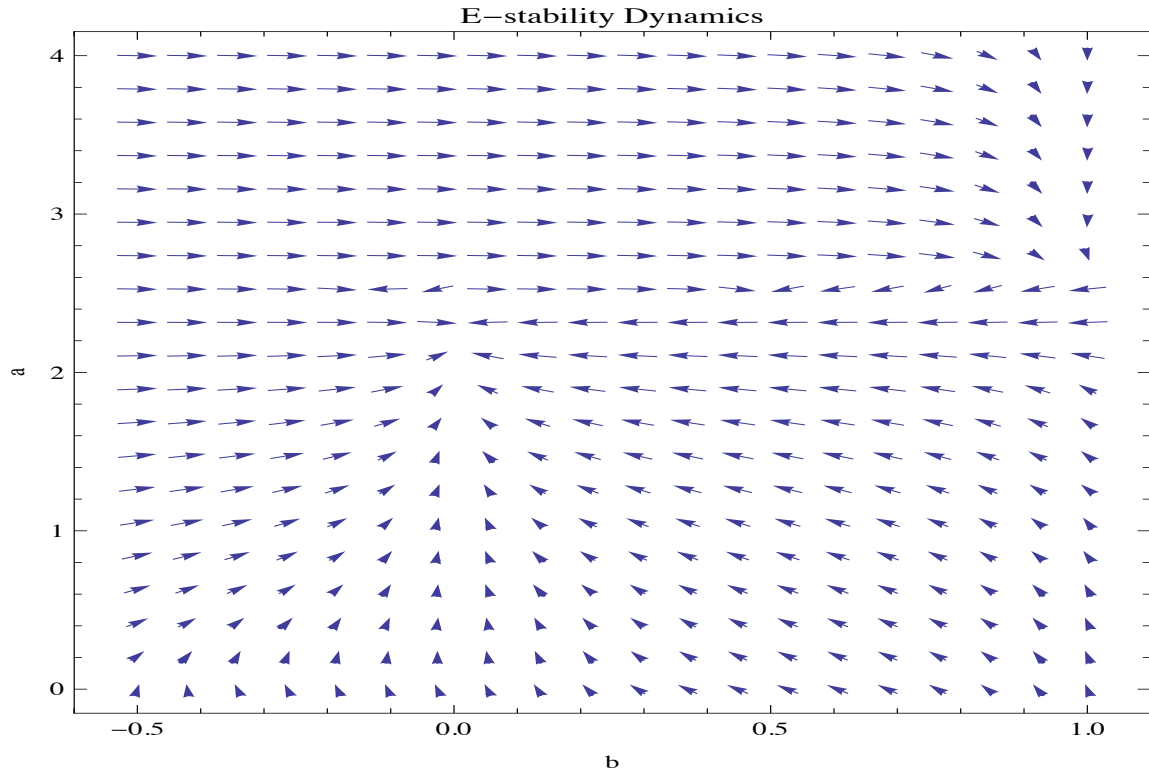
will be attainable under least squares and closely related learning algorithms.¹¹ That the E-stability principle governs stability under learning is intuitive since under (6) the parameters (c, b) are adjusted in the direction of the asymptotic moments implied by the actual law of motion generating the data given (c, b) . Local stability of (6) then answers the question of whether, under these E-stability dynamics, a small displacement of (c, b) from a SCE would return to the equilibrium.

Analytic results on E-stability of SCE are not available because there is not a closed-form expression for the T-map.¹² However, it is straightforward to demonstrate

¹¹Here τ denotes “notional” time

¹²Analytic E-stability results are available for alternative learning rules. For example, in the monetary version of this model (e.g. $y = 0$) Baranowski (2012) shows that the stationary equilibrium is E-stable when agents simply estimate the conditional mean (i.e. omit the lag from their forecasting

Figure 2: E-stability Dynamics



E-stability by numerically calculating the T-map for a specific parameter choice. A numerical investigation revealed that only the trivial SCE, i.e. the fundamental equilibrium, is E-stable as anticipated by the results in Branch and McGough (2005). Figure 2 demonstrates the E-stability dynamics for the baseline parameterization in Table 1. In this case, the mean price, i.e. the fundamental equilibrium, is $\bar{p} \approx 2.3$.

Figure 2 shows that the SCE corresponding to the fundamental equilibrium $(\bar{p}, 0)$ is E-stable. The figure plots the resting points of the E-stability ODE and the associated vector field. There are rest points at $(0, 2.3)$ and $(1, 2.3)$ and the arrows indicate the direction of adjustment in (6). The figure shows that the non-trivial SCE with $b = 1$ is unstable under the E-stability dynamics. In contrast, the fundamental equilibrium with $c = \bar{p}$ and $b = 0$ is a sink under learning.

Figure 2 illustrates three further features. First, the fundamental equilibrium is E-stable and its basin of attraction includes all initial conditions. The price that

model). Baranowski demonstrates the E-stability arises across a number of specifications for the Lagos and Wright (2005) model, including Nash and proportional bargaining.

would arise without a liquidity premium, i.e. discounted present value of dividends, which lies below \bar{p} , is not an equilibrium and would not arise under learning. Thus, learning dynamics drive asset prices towards the price that includes the liquidity premium. Third, most analyses of asset pricing under learning focus on the E-stability properties of a particular equilibrium. The figure also demonstrates that the transitional dynamics might be of independent interest. The vector field indicates that some transitional paths may include non-linear paths to the fundamental equilibrium.

3.4 A Special Case: $\sigma = 1$

The E-stability dynamics govern the stability of the rational expectations equilibrium. However, they do not give the full picture of global learning dynamics. This section details the learning dynamics and illustrates how the liquidity properties of the asset can alter the qualitative nature of learning dynamics. The central idea is the following: agents are aware of the form of the asset pricing process but the specifics, such as the frequency of transactions in the decentralized market, the illiquidity of the asset, etc., are unknown. An agent in this setting would be wise to remain alert to potential changes in the liquidity of the asset. Such an agent will then place a prior probability on drifting coefficients in their forecasting model. There are two central ingredients to the results that come below: the asset is in a sufficiently low supply (i.e. illiquid) that is imperfectly known by agents, and a prior belief of possible structural change. The results from this section will be instructive for the asset price dynamics featured in subsequent sections.

3.4.1 Asymptotic Learning Dynamics

In the case that $\sigma = 1$ the form of the asset pricing equation is simplified and additional analytic results are available on the global learning dynamics. In this case, the law of motion for asset prices is conditionally linear:

$$p_t = (1 - \alpha)\beta\hat{E}_t(p_{t+1} + y_{t+1}) + \alpha\beta A^{-1}(1 + \varepsilon_t) \quad (7)$$

where ε_t is white noise with variance σ_ε^2 . For fixed belief parameters (a, b) , the law of motion in (7) is conditionally linear and the analysis in Branch and Evans (2011) can be used to gain analytic insight into the nature of real-time learning dynamics. When $\sigma = 1$ there is a unique SCE that corresponds to the fundamentals equilibrium. The next section presents simulation results for $\sigma > 1$.

To formulate a least-squares updating rule, it is convenient to formulate the fore-

casting rule as a linear AR(1) model:

$$p_t = c + bp_{t-1} + \varepsilon_t$$

so that

$$\hat{E}_t p_{t+1} = c(1+b) + b^2 p_{t-1} \quad (8)$$

Plugging these expectations into (7) leads to the actual law of motion

$$p_t = (1-\alpha)\beta(y + c(1+b)) + \alpha\beta A^{-1} + \beta b^2 p_{t-1} + \alpha\beta\varepsilon_t \quad (9)$$

$$\equiv T(a,b)'X_{t-1} + \alpha\beta\varepsilon_t \quad (10)$$

A key assumption of the approach in this paper is that the agents must jointly learn the mean and serial correlation of asset prices. Let $\theta' = (c, b)$, $X' = (1, p)$. Agents are assumed to update their parameter estimates according to the following recursive algorithm

$$\theta_t = \theta_{t-1} + \gamma S_{t-1} X_{t-1} (p_t - \theta' X_{t-1})' \quad (11)$$

$$S_t = S_{t-1} + \gamma (X_t X_t' - S_{t-1}) \quad (12)$$

The equations in (11)-(12) are the updating equations for recursive least squares where the data are discounted by a constant “gain” γ . Here S_t is an estimate of $EX_t X_t'$, the second moment matrix of the regressors. Least-squares updating arises when the constant gain γ is replaced by a decreasing sequence $\gamma_t = t^{-1}$. Sargent and Williams (2005) demonstrate that constant gain learning equations (11)-(12) arise from an approximate Bayesian learning process in which the prior on parameter drift, a common assumption in applied econometric work, is proportional to the ratio of observation noise variance to the covariance of the regressors, with the speed of drift controlled by the constant gain γ . An alternative interpretation of (11)-(12) is that agents use least squares modified to discount past data due to a concern about (possible) structural change of an unknown form. In the stochastic simulations below, we set $\gamma = 0.10$ which equates to an effective sample size of approximately 80 years.¹³

The asymptotic behavior of θ_t is non-trivial because the model (7) is self-referential. It turns out that for small gains γ it is possible to obtain results on the asymptotics by studying a continuous time approximation to the recursive algorithm. More specifically, Evans and Honkapohja (2001) demonstrate that asymptotically the dynamics

¹³Empirical estimates of constant gains typically fall in the range of 0.02 – 0.10, see Branch and Evans (2006). Bubbles and crashes arise more often in stochastic simulations for larger values of γ but also introduce a greater amount of volatility.

are governed by the “mean dynamics” ordinary differential equation (ODE)

$$\frac{d\theta}{d\tau} = S^{-1}M(\theta)(T(\theta) - \theta) \quad (13)$$

$$\frac{dS}{d\tau} = M(\theta) - S \quad (14)$$

where $\tau = \gamma t$, $M(\theta)$ is the unconditional covariance matrix of the regressors holding θ fixed. The ODE (13) guiding the evolution of θ is identical to the E-stability differential equation with the exception that it includes weighting terms that depend on estimates of the covariance matrix. It is straightforward to see that the fundamental equilibrium is a locally stable rest point provided that $(1 - \alpha)\beta < 1$.

Under decreasing gain learning γ is replaced with $1/t$ and it can be shown that in the limit as $t \rightarrow \infty$ the learning dynamics converge with probability one to the bubble equilibrium. This paper focuses on constant gain learning, in which parameter estimates weight recent data more heavily than past. We next summarize the analytical results for constant gain learning by directly applying the results in Branch and Evans (2011).

The first result establishes that for a sufficiently small constant gain the perceived coefficients θ_t will be an approximately normal random variable with a mean equal to its fundamental equilibrium value and a variance that depends on both the constant gain and other parameters of the model including the illiquidity parameter A . The second result shows that from a given initial condition (θ_0, S_0) the solution to the “mean dynamics” of the ODE (13-14) give the expected transition path to the fundamental equilibrium values.

Proposition 5 *The belief parameters θ_t are approximately distributed as $\theta_t \sim N(\bar{\theta}, \gamma V)$ for small $\gamma > 0$ and large t , where $\bar{\theta} = (\bar{p}, 0)'$ and for appropriately defined V .*

Proposition 6 *Define $\phi_t = (\theta_t, \text{vec}(S_t))'$. For any ϕ_0 within a suitable neighborhood of the fundamental equilibrium, define $\tilde{\phi}(\tau, \phi_0)$ as the solution to the differential equation (13)-(14), with initial condition ϕ_0 . Fix $T > 0$. The mean dynamics of (11)-(12) satisfy $E\phi_t \approx \tilde{\phi}(\gamma t, \phi_0)$ for γ sufficiently small and $0 \leq t < T/\gamma$.*

There are important conclusions to draw from these propositions. First, the fundamental equilibrium provides a natural benchmark in the sense that the coefficients for the forecast rule under learning are centered on these equilibrium values. Second, for $\gamma \rightarrow 0$, the learning dynamics are arbitrarily close to their fundamental equilibrium values with high probability. Third, we can gain insight into the global learning

dynamics, for finite stretches of time, by studying the solution paths to the mean dynamics differential equation given initial conditions. Moreover, these initial conditions can be drawn from the asymptotic distribution for θ computed in Proposition 5.

3.4.2 Learning Dynamics and Random-Walk Beliefs


Under constant gain learning there can be occasional, temporary departures from rational expectations. These departures can arise as an endogenous response to the exogenous liquidity shocks. This section illustrates these possibilities for the special case $\sigma = 1$.

Proposition 6 shows that for any initial condition, and finite period of time, that the expected learning path to the fundamental equilibrium will be the solution path to the mean dynamics equations (13)-(14). One can think of constant gain learning, which respond more strongly to recent shocks, as re-initializing the mean dynamics. Figure 3 plots representative mean dynamics where initial values for $c, b > 0$ are drawn from the asymptotic distribution in Proposition 6. To generate this figure the parameter values are chosen according to Table 1 except $\sigma = 1$. The initial values are $c = 3.73$ and $b = .4$ correspond to an increase in the perceived mean and serial correlation of the asset price. The top panel plots the perceived value for the intercept, c , while the bottom panel plots the perceived lag coefficient b .

Figure 3 illustrates two aspects of the expected learning path. First, the bubble equilibrium is a stable rest point of the mean dynamics implying that learning paths converge to the bubble equilibrium. Second, the transition path to the bubble equilibrium is interesting. At first the path for (c, b) moves back toward the bubble equilibrium then abruptly reverses course with $c \approx 0$ and $b \approx 1$ for a finite period of time, before converging to the bubble equilibrium values. The mean dynamics, therefore, show that the expected learning path has agents perceiving, for stretches of time, that the asset price follows a random walk without drift.

As discussed in the above subsection on stochastic consistent expectations equilibria, random-walk beliefs play a key role in the learning dynamics. First, they can be nearly self-fulfilling and the additional serial correlation introduced through beliefs has important implications for the dynamic nature of asset prices. Second, random-walk beliefs can arise through learning as temporary deviations from the bubble equilibrium. Third, as the next section will demonstrate, random walk beliefs can lead to speculative bubbles and crashes. In essence, agents come to believe that recent innovations in asset prices are permanent shifts and not mean-reverting fluctuations. They trade on these beliefs and bid asset prices up.

Figure 3: Mean Dynamics.




3.4.3 Implications of Asset Shortages

Having established the possibility of random-walk beliefs emerging under learning, we turn to the comparative effects of asset shortages. One can proxy for asset shortages via the parameter A which measures the average asset float in the economy. Of course, if A is sufficiently large then there is no liquidity premium. Smaller values of A then yield a larger liquidity premium and the asset can be said to be illiquid since there is not sufficient supply to support an efficient quantity of trade in the decentralized market. The onset of random-walk beliefs depends on a complicated interaction of the entire set of structural parameters. It is natural, though, to wonder what role the supply of the asset has in leading to random-walk beliefs.

To address this question, Figure 4 plots the mean dynamics for $A = 0.4$ and $A = 0.75$, keeping the other parameter values the same as in Figure 3 and Table 1. As above, the learning dynamics are initialized from the 95% confidence ellipse around the stationary equilibrium. For each value of A , the learning dynamics converge to the stationary equilibrium. The transitional dynamics differ between the two learning paths. For a larger asset supply – or, smaller liquidity premium – there is a slow drift in the estimated intercept coefficient, c , and the learning dynamics do not feature

Figure 4: Mean Dynamics.



random walk beliefs that arise when the asset is in shorter supply (e.g. $A = 0.4$).

When the asset is in short supply, it directly affects asset price dynamics through the mean price, i.e. a larger liquidity premium, and by making asset prices more volatile. The results in Figure 4 also show that a smaller supply can have an indirect effect via the learning dynamics. Only when the liquidity properties are sufficiently high – and asset prices are volatile – do the learning dynamics feature random walk beliefs which, as will be seen below, are key to generating bubbles and crashes.

4 Asset Shortages, Bubbles and Crashes

When the asset is in short supply, its stationary equilibrium price consists of two components the discounted dividend flow and a liquidity premium. The results from the previous section show that with a large liquidity premium, random walk beliefs can arise in transitional learning dynamics. This section presents three implications of the model with adaptive learning: first, a permanent change in the (mean) asset supply leads to an asset price path that overshoots its new higher level; second, an increase in the demand for collateral – arising, for example, from macroprudential

policies that emphasize secured over unsecured credit – will lead to an asset price path that overshoots its new higher level; third, occasional temporary shocks to the supply of assets can lead to bubbles and crashes.

4.1 A Change in Asset Supply

An imbalance in the supply of safe assets can be interpreted, within the context of the model, as a decrease in the (mean) asset share supply A . Since A is the per-capita supply of the asset, values of $A < 1$ imply a demand imbalance that manifests as a liquidity premium in the stationary equilibrium price. A decrease in A is a further shortening of asset supply.


This subsection considers the following experiment of a decrease in the (mean) asset supply. The model parameters are chosen according to the baseline parameterization with an initial asset share supply $A = 0.40$. The model is initialized in a stationary equilibrium and individuals' beliefs are set to their rational expectations equilibrium values. Then A is lowered permanently to $A = 0.2$, a change that reflects a decrease in the supply of safe assets. Although, the economy is initially in a rational expectations equilibrium, the agents in the economy have imperfect information about the change in the supply of assets and the greater liquidity premium that will eventually arise. Figure 5 plots the resulting belief and price dynamics.¹⁴ At time 0, the asset supply A decreases, raising the liquidity premium and the asset price without a corresponding increase in price expectations (which are determined by an adaptive learning rule). Initially, the asset price is below the new stationary equilibrium price. The increase in the asset price is tracked by agents' econometric model as an increase in the persistence of prices, reflected in an increase in the estimated value for b_t . As the mean dynamics predict (for the special case), eventually agents' beliefs are that prices follow a random walk. At this point, there is a burst in prices as the asset price increases to nearly double its new long-run value before converging to the new equilibrium price.

Thus, a decrease in the supply of safe assets can lead to price dynamics that resemble speculative bubbles. This bubble-like path arises because the initial upward drift in prices leads to a nearly self-fulfilling belief that the asset price follows a random walk.¹⁵ The price dynamics are consistent with the observation that a shortage in safe assets will lead to asset bubbles.

¹⁴Figures 5 and 6 are generated as the average time-path across 10,000 stochastic simulations of length 1,000.

¹⁵These results are similar to McGough (2006) who examines changes to the natural rate of unemployment in the model of policymaker learning developed in Sargent (1999).

Figure 5: A Decrease in Asset Supply: $A = 0.4 \rightarrow A = 0.2$.




4.2 Implications of Changes in Collateral Demand

An imbalance in the demand for assets can arise from a shortage of safe assets, as in the previous subsection, or from a greater demand for the use of assets as collateral in over-the-counter markets. Recently, policymakers such as the IMF, and other market observers, have identified several structural and policy changes that have the potential to increase the global demand for safe assets.¹⁶ Among these changes, are new financial regulations that require an increasing number of over-the-counter transactions to be cleared through central clearinghouses. Additionally, many central banks of use safe assets as collateral in repurchase agreements. This subsection investigates the potential impact of increased demand for collateral in over-the-counter transactions.

The model itself, of course, consists of bilateral trade and does not readily feature a third party clearing house. Incorporating a tri-party market is beyond the scope of the present study. However, the model lends itself to the following interpretation: as the probability α of a match between buyer and seller increases, buyers will

¹⁶See Iorgova, Al-Hassan, Chikada, Fandl, Morsy, Pihlman, Schmieder, Severo, and Sun (2012).

Figure 6: An Increase in the Demand for Collateral: $\alpha = 0.01 \rightarrow \alpha = 0.2$.



find themselves in more frequent need of collateral to secure trade in over-the-counter transactions. Hence, an increase in α can be interpreted as a structural increase in the demand for collateral.

Essentially, a higher probability α makes the asset more liquid as it becomes easier with the lower search friction to exchange the asset for goods and services. This subsection considers the following experiment. The model parameters are chosen as before, except initially $\alpha = 0.01$, signifying an illiquid market with a low probability of over-the-counter trade. The model, and agents' beliefs, are initialized at their stationary equilibrium values. Then α is increased to $\alpha = 0.2$. Figure 6 illustrates the results. As the asset becomes more liquid, there is an upward drift in prices. Agents' econometric model of the price quickly picks that drift up as a greater serial correlation in the price process. A positive feedback loop results where agents believe that asset price innovations are serially correlated and so recent price increases will persist, leading to higher price dynamics, and higher estimated degrees of persistence. Eventually, agents come to believe that the asset price is a random-walk.

4.3 Real-time Asset Price Dynamics

Figures 5 and 6 demonstrate that changes in the demand imbalance for the asset can result in asset price dynamics that overshoot their stationary values. These bubble like dynamics arise because agents come to believe that asset prices follow a random-walk and, as argued before, these beliefs are nearly self-confirming. Random-walk beliefs can only persist for a finite period of time before, as the mean dynamics predict the learning process brings the economy back near its stationary equilibrium. These results suggest that real-time learning dynamics might exhibit recurrent bubbles and crashes. This is the topic of the present subsection.


This subsection uses stochastic simulations of the real-time learning dynamics to demonstrate that bubbles and crashes – i.e., substantial deviations from the stationary equilibrium values – can arise as an endogenous response to the fundamental asset supply shocks.¹⁷ For example, Figure 7 plots a real-time simulation of price dynamics. As before, the figure is computed with the parameter values in Table 1 except with $A = 0.15$ and a constant gain of $\gamma = 0.10$. To generate this figure the model is initialized at the stationary rational expectations equilibrium, expectations are generated according to (8) with parameters updated via constant gain least-squares, and price is determined by (3). The left panels plot the asset price and the quantity traded in the decentralized market. The right panels plot the estimated coefficients (c_t, b_t) .

Under constant gain learning, the economy hovers near its stationary rational expectations equilibrium price most of the time with small iid deviations. Since these learning dynamics are near the noisy rational expectations equilibrium, the figure illustrates that, on average, beliefs are close to their rational expectations values. At about period 200, there is an abrupt qualitative change in the dynamics with a bubble in the asset price. This bubble features a price that increases over 3 times above its fundamental value. The bubble persists only for a finite length of time before returning to a neighborhood of the stationary equilibrium. The pattern of beliefs correspond with what was observed in Figure 3 and in Proposition 6, in that for finite stretches of time agents believe that inflation follows a random walk. In simulations, these large deviations from rational expectations are recurrent.

Figure 7 provides an intuitive story for the existence of asset price bubbles. Because of imperfect knowledge about the economy, individuals learn about the price process via an econometric forecasting rule that remains robust to structural change and model misspecification by weighting recent data more heavily than past data.

¹⁷To prevent explosive dynamics, we impose the following restrictions on the learning dynamics: (1.) agents only update their estimates of b provided that it lies below $1/\beta$; (2.) agents forecast a non-negative return on the asset, i.e. $\hat{E}_t p_{t+1} + y \geq 0$.


Figure 7: Bubble.



Occasional economic shocks can give the asset price process an upward drift that is captured by agents' econometric model as serial correlation. Eventually this serial correlation becomes self-fulfilling and, for a finite stretch of time, agents come to believe that asset prices follow a random-walk. With random-walk beliefs, buyers will interpret recent price deviations as permanent shifts in the price process and will demand more of the asset. Eventually, though, as the asset price becomes sufficiently high, the liquidity premium vanishes and the learning dynamics return to a neighborhood of the rational expectations equilibrium. Importantly, Figure 7 shows that the collapse in the bubble can be costly as there is a sharp decline in q_t , the quantity traded in the decentralized market. So, although a bubble can lead to a greater amount of decentralized trade, when bubbles eventually collapse there is an abrupt economic contraction. This figure is a key result of this paper.

It is worth briefly remarking on the qualitative nature of the bubble and crash in Figure 7. The bubble features an abrupt upward spike followed by a more gradual adjustment to the fundamental price. Many bubble episodes, in practice, feature more gradual run-ups in price that are followed by an abrupt crash. The model is capable of producing bubbles and crashes whose qualitative properties more closely match practical experience. However, to generate more empirically realistic bubbles

Figure 8: Crash.



requires abandoning the modeling assumptions – such as quasi-linear preferences, linear production technologies, etc. – that were made for technical convenience and analytic transparency.

Figure 8 demonstrates that the first deviation away from the equilibrium price can be a crash. In this case, the logic is symmetric to why bubbles arise. A sequence of positive asset supply shocks place a downward drift on price leading agents' economic model to pick up this trend with higher estimated values of b_t and lower values of a_t . In turn, asset price expectations decrease which leads to a lower demand for the asset and a smaller quantity of over-the-counter trade. A downward price spiral arises as there is a further downward drift in price, higher estimated values of b_t until the coefficients arrive at a random walk model which, as argued above, is nearly self-confirming. These beliefs only persist for a finite period of time and eventually the stability of the stationary equilibrium takes over and price dynamics return to their stationary equilibrium value.

Unlike Figure 5, the deviations away from the rational expectations equilibrium in Figure 7 are an endogenous response to fundamentals rather than to a change in the long-run asset share supply. It is possible, using the techniques in Cho, Williams, and

Sargent (2002), to examine which of the “escape paths” are most likely to drive the economy away from the stationary equilibrium and lead to random-walk beliefs by identifying the “most likely unlikely” sequence of shocks that move asset price a given distance away from the stationary equilibrium value. In principle, one can compute these escape paths analytically in special cases, but more typically it is necessary to resort to simulations. However, as explained above, it is intuitive that random walk beliefs arise for the right sequence of shocks.

5 Discussion

The results presented here demonstrate that bubbles and crashes can arise in response to changes in the supply of assets, changes in the over-the-counter liquidity of the asset, and as an endogenous response to transitory, economic shocks. There were two key features to the analysis: first, because of search frictions in over-the-counter market, and a shortage of assets that can serve as collateral, the financial asset’s price includes a liquidity premium; second, imperfect knowledge about the future asset price leads individuals to formulate, and estimate in real-time, an econometric forecasting model. The adaptive learning process of revising the estimates of the econometric model can lead individuals to temporarily believe that price follows a random walk without drift, which can be nearly self-confirming. Random walk beliefs lead asset prices to deviate from their fundamentals price as agents interpret recent price innovations as permanent, leading to a positive feedback loop that results in rapid price appreciation and increased trade in over-the-counter markets. As that liquidity premium vanishes, the learning dynamics return price to its fundamental value, however, the route back to equilibrium features a crash in asset prices and an abrupt decline in over-the-counter trade.

The results presented were in a simple, stylized search-based model. It is natural to wonder whether the results are sensitive to a number of assumptions that were made for technical convenience. First, all trade takes place in the decentralized market through “buyer-takes-all” bargaining. Since the buyer decides on the asset holdings based on the expected return of the asset as well as its liquidity value, the fact that the buyer captures all of the surplus evidently increases the expectational feedback by making the liquidity premium higher. As an alternative, one could relax this assumption by imposing a proportional bargaining game between buyers and sellers. Although, this would alter the quantitative results of the model, the qualitative finding of bubbles that arise from random-walk beliefs would not be affected. The quasi-linearity assumption implies that the store-of-value role of the asset economizes on labor in the competitive market. Nor does the asset carry a risk premium. A

careful quantitative analysis with the goal of matching key empirical moments would need to relax this assumption along the lines of Lagos (2010).

The asset supply process was taken to be exogenous for technical simplicity. But because of the important role that asset supply shortages make it would be worthwhile to extend the model to endogenize the supply of assets. In particular, one lesson from the housing bubble and financial crisis was that a shortage of safe assets led to clever financial engineering that created safe assets (or, at least the illusion of safe assets). This issue is addressed by Li, Rocheteau, and Weill (2012) who study an environment where some individuals can choose to create fraudulent assets that appear as safe assets. An alternative approach is Rocheteau and Wright (2011) where firms that want to enter the over-the-counter market must pay a fixed entry cost and that cost could be financed through issuing equity. Holmstrom and Tirole (2011) endogenize asset supply by firms who finance investment and have limited pledgeability from future income flows. The main purpose of this paper is to explore the asset pricing implications of asset shortages in an imperfect knowledge environment and so abstracts from these interesting issues. Similarly, Guerrieri and Shimer (2012) and Guerrieri, Shimer, and Wright (2010) show that adverse selection in asset markets can affect the distribution of safe assets and market liquidity.

Of course, in practice investors have an array of assets, with different liquidity, risk, and return characteristics, to choose from, or accept, in over-the-counter markets. An interesting extension of the present environment would be to expand the number of assets and examine the extent to which imperfect knowledge affects the cross-sectional liquidity characteristics of assets. Similarly, the model presented could easily be extended to include money and government bonds and used to study the optimal provision of liquidity and safe assets in an imperfect knowledge environment.

The form of imperfect knowledge assumes that individuals make optimal asset holdings given their beliefs. Individuals are assumed to be anticipated utility maximizers, taking the future evolution of their beliefs given. An alternative approach would be to assume the agents are fully Bayesian and decide on asset holdings by taking into account all of the possible future learning paths. It is possible to incorporate fully Bayesian decision making into the present environment because the quasi-linearity, and portfolio adjustment in the competitive market, essentially means that agents are only making decisions one period in advance. However, in a less stylized environment fully Bayesian decision making would require that individuals form optimal plans for each of the possible infinitely-long sequences of future beliefs. Thus, Bayesian decision making would impart unrealistic cognitive abilities onto agents with imperfect knowledge, and so this approach is not pursued in this paper.¹⁸

¹⁸Cogley and Sargent (2008) present a framework for comparing anticipated utility and Bayesian

6 Conclusion

This paper studied asset pricing in an environment where the asset has a dual role as a store-of-value and in providing liquidity services. The liquidity role is formalized using search frictions that are familiar in monetary theory. In a stationary equilibrium, the asset carries a liquidity premium, above its discounted payment flows, whenever the supply of the asset is not sufficient to support all of the trade in decentralized, or over-the-counter, transactions. However, the liquidity premium in the model does not generate the kinds of price dynamics – such as bubbles, crashes and excess volatility – that are typically observed in practice. This paper proposes a search-based asset pricing model with imperfect knowledge and adaptive learning as a means for generating bubbles and crashes in asset prices.

The primary results of this paper are as follows. First, although over time beliefs converge toward rational expectations, and price towards its fundamental value, the combination of adaptive learning and a shortage of assets can lead individuals in the economy to temporarily believe that asset prices follow a random walk without drift. Such beliefs are temporarily (nearly) self-confirming. When agents perceive the price process to be a random walk they interpret recent price innovations as permanent shifts in the fundamental price. These random walk beliefs arise for a very intuitive reason. Imperfect knowledge about the economic environment lead individuals to estimate, using recent data, an econometric forecasting model for the future price. If data lead to a slight upward drift in the price, agents' econometric model will pick up that drift, leading to higher price expectations that feed back into higher prices. This process is self-reinforcing and in some cases agents eventually come to believe that prices follow a random walk. Crucially, we have shown that these beliefs are nearly self-fulfilling.

A decrease in the supply of safe assets, by raising the liquidity premium, will introduce just the type of price drift that can lead to random walk beliefs. These random walk beliefs cause a substantial overshooting of the fundamental price. An increase in the demand for collateral – for example, changes in regulatory and macro-prudential policies that require more over-the-counter trade to post safe assets as collateral – can also introduce the type of drift in inflation. Finally, occasional “unlikely” sequences of shocks to asset supply can also introduce drift to the price process that trigger random-walk beliefs and large deviations from the fundamental equilibrium. Such departures from rational expectations can generate significant bubbles and crashes. These results shed light on the implications of demand imbalances in safe assets can

decision making. They find that anticipated utility provides a good approximation to Bayesian decision making.

have for economic stability.

References

- ADAM, K., A. MARCET, AND J. P. NICOLINI (2010): “Stock Market Volatility and Learning,” Working paper.
- ALLEN, F., AND D. GALE (2000): “Bubbles and Crises,” *Economic Journal*, 110, 236–255.
- BAKER, M., AND J. WURGLER (2000): “The Equity Share in New Issues and Aggregate Stock Returns,” *Journal of Finance*, 55(5), 2219–2257.
- BAO, J., AND J. PAN (2012): “Relating Equity and Credit Markets through Structural Models: Evidence from Volatilities,” working paper.
- BARANOWSKI, R. (2012): “Adaptive Learning and Monetary Exchange,” working paper.
- BARLEVY, G. (2011): “A Leverage-based Model of Speculative Bubbles,” working paper.
- BRANCH, W., AND G. W. EVANS (2006): “A Simple Recursive Forecasting Model,” *Economic Letters*, 91, 158–166.
- BRANCH, W. A., AND G. W. EVANS (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3(3), 159–191.
- BRANCH, W. A., AND B. MCGOUGH (2005): “Consistent Expectations and Misspecification in Stochastic Non-linear Economies,” *Journal of Economic Dynamics and Control*, 29, 659–676.
- CABALLERO, R. J. (2006): “On the Macroeconomics of Asset Shortages,” in *The Role of Money: Money and Monetary Policy in the Twenty-First Century*, ed. by A. Beyer, and L. Reichlin, pp. 272–283.
- CABALLERO, R. J., E. FARHI, AND P.-O. GOURINCHAS (2008): “An Equilibrium Model of “Global Imbalances” and Low Interest Rates,” *American Economic Review*, 98(1), 358–393.

- CHO, I.-K., AND K. KASA (2008): “Learning Dynamics and Endogenous Currency Crises,” *Macroeconomic Dynamics*, 12, 257–285.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1–40.
- COCHRANE, J. H. (2005): “Liquidity, Trading and Asset Prices,” *NBER Reporter*.
- COGLEY, T., AND T. J. SARGENT (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49.
- DICK-NIELSEN, J., P. FELDHUTTER, AND D. LANDO (2012): “Corporate Bond Liquidity Before and After the Onset of the Subprime Crisis,” *Journal of Financial Economics*, 103, 471–492.
- DUFFIE, D., N. GARLEANU, AND L. H. PEDERSEN (2005): “Over-the-Counter Markets,” *Econometrica*, 73, 1815–1847.
- EVANS, G. W., AND S. HONKAPOHJA (1995): “Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models,” *Econometrica*, 63, 195–206.
- (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2013): “Learning as a Rational Foundation for Macroeconomics and Finance,” in *Rethinking Expectations: The Way Forward for Macroeconomics*, ed. by R. Frydman, and E. S. Phelps. Princeton University Press.
- FARHI, E., AND J. TIROLE (2011): “Bubbly Liquidity,” *Review of Economic Studies*, forthcoming.
- GEROMICHALOS, A., J. M. LICARI, AND J. SUAREZ-LLEDO (2007): “Asset Prices and Monetary Policy,” *Review of Economic Dynamics*, 10(4), 761–779.
- GRANDMONT, J.-M. (1985): “On Endogenous Competitive Business Cycles,” *Econometrica*, 53, 995–1045.
- GUERRIERI, V., AND R. SHIMER (2012): “Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality,” working paper.
- GUERRIERI, V., R. SHIMER, AND R. WRIGHT (2010): “Adverse Selection in Competitive Search Equilibrium,” *Econometrica*, 78(6), 1823–1862.

- HOLMSTROM, B., AND J. TIROLE (2011): *Inside and Outside Liquidity*. MIT Press.
- HOMMES, C. H., AND M. ZHU (2011): “Learning under misspecification: a behavioral explanation of excess volatility in stock prices and persistence in inflation,” working paper.
- IORGOVA, S., A. AL-HASSAN, K. CHIKADA, M. FANDL, H. MORSY, J. PIHLMAN, C. SCHMIEDER, T. SEVERO, AND T. SUN (2012): “Safe Assets: Financial System Cornerstone?,” in *Global Financial Stability Report*. International Monetary Fund.
- KIYOTAKI, N., AND J. MOORE (2008): “Liquidity, Business Cycles and Monetary Policy,” mimeo.
- KREPS, D. M. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge University Press.
- KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2012): “The Aggregate Demand for Treasury Debt,” *Journal of Political Economy*, *forthcoming*.
- LAGOS, R. (2010): “Asset Prices and Liquidity in an Exchange Economy,” *Journal of Monetary Economics*, 57(8), 913–930.
- LAGOS, R., AND G. ROCHETEAU (2009): “Liquidity in Asset Markets with Search Frictions,” *Econometrica*, 77, 403–426.
- LAGOS, R., AND R. WRIGHT (2005): “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113, 463–484.
- LESTER, B., A. POSTLEWAITE, AND R. WRIGHT (2012): “Information, Liquidity, Asset Prices, and Monetary Policy,” *Review of Economic Studies*, *forthcoming*.
- LI, Y., G. ROCHETEAU, AND P.-O. WEILL (2012): “Liquidity and the Threat of Fraudulent Assets,” working paper.
- MCGOUGH, B. (2006): “Shocking Escapes,” *Economic Journal*, 116, 507–528.
- NOSAL, E., AND G. ROCHETEAU (2011): *Money, Payments, and Liquidity*. MIT Press.
- ORPHANIDES, A., AND J. C. WILLIAMS (2005): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford, chap. 5, pp. 201–234. University of Chicago Press.

- ROCHETEAU, G., AND R. WRIGHT (2005): “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73, 175–202.
- (2011): “Liquidity and Asset Market Dynamics,” *forthcoming in Journal Of Monetary Economics*.
- SANTOS, M., AND M. WOODFORD (1997): “Rational Asset Pricing Bubbles,” *Econometrica*, 65, 19–57.
- SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- SARGENT, T. J., AND N. WILLIAMS (2005): “Impacts of Priors on Convergence and Escapes from Nash Inflation,” *Review of Economic Dynamics*, 8, 360–391.
- SINGH, M. (2011): “Velocity of Pledged Collateral: Analysis and Implications,” Discussion Paper WP/11/256, IMF Working Paper.
- TIROLE, J. (1985): “Asset Bubbles and Overlapping Generations,” *Econometrica*, 53(6), 1499–1528.
- WEILL, P.-O. (2008): “Liquidity Premia in Dynamic Bargaining Markets,” *Journal of Economic Theory*, 140, 66–96.
- WILLIAMS, N. (2004): “Escape Dynamics in Learning Models,” Discussion paper, working paper, Princeton University.