

Speculative Trading with Bayesian Learning¹

Preliminary version

Adam Slawski
Department of Economics
University of Minnesota

May 1, 2008

¹I am grateful to Narayana Kocherlakota, Jan Werner, Itai Sher, David Rahman, Christian Stoltenberg, Kevin Wiseman, Christoph Mueller, Radek Stefanski for discussing the project and helpful feedback. I would like to also thank the participants of Macroeconomics and Public Policy as well as Math Econ workshops at University of Minnesota.

Abstract

The paper analyzes conditions for speculative bubbles to arise in a dynamic general equilibrium model with heterogeneous prior beliefs and the short sale constraints. The question of interest is how learning and convergence of beliefs affects speculation. There is a single asset, which pays a stochastic dividend each period. The dividend process follows a hidden Markov chain. There are two types of risk neutral traders who formulate different prior beliefs about the parametric specification of the process. I develop a systematic way of analyzing recursive equilibria with learning. I introduce the concept of belief based recursive equilibrium, which uses the current belief profile as a state variable. I show that any sequential market equilibrium price must be at least as big as the recursive equilibrium one. I propose a simple functional operator, which pins down the recursive equilibrium price as its fixed point. The operator is shown to be a monotone contraction. This gives uniqueness and provides a simple analytical criterion for existence of speculative bubbles in equilibrium.

1 Introduction

Stock prices often seem higher than the fundamental warrants. We could see that during the famous “Japanese asset price bubble” between 1986 and 1990 and also in the US in the late 90’s during the so-called “dot-coms” bubble when the share prices of internet and telecommunication stocks reached unprecedented high levels in March 2000 to lose huge part of their values within the next couple of months.

One explanation for this kind of speculation was proposed by Harrison and Kreps (1978). In their model the speculation is driven by differences in subjective beliefs accompanied by a short sale constraint.

They consider an example with two groups of risk neutral agents. The agents trade one risky asset of supply 1. There is a short sale constraint. Dividends follow a 2-state Markov chain and can take value 0 or 1 in each period. The agents differ in their beliefs about the transition matrix. Agents of type one assign relatively high probabilities of switching between the states (dividends) while agents of type two perceive the states as more persistent. Both types are certain that their transition matrix is correct. The numerical values are chosen so that in each state the present expected value of the stream of all future dividends (the fundamental value) is higher according to the agents of group two.

This would suggest that it should be the agents of type two who permanently hold the asset in equilibrium and also, given risk neutrality, the price should be exactly their fundamental value. Surprisingly, only if the dividend is 1 agents of group two hold the asset. If the dividend is 0 agents of type one buy it. Why is it so? The mechanism is quite simple. Type one agents buy the asset in state 0 (when the dividend is 0), because they think there is a good chance of switching to the other state next period. Further, they (rationally) expect that the price will be high if the dividend is 1, since type two agents assign a high fundamental value to the asset in that state. Moreover we get that the equilibrium price is higher than any agent’s fundamentals in each state. Throughout Sections 3 and 4 I will go over the Harrison and Kreps model in more detail, using it as an illustrative example for my model.

The result of Harrison and Kreps is simple and quite beautiful but unrealistic in some respects. The asymmetry of agents’ subjective beliefs persists in spite of commonly observed histories. It is this feature which makes the speculation so highly persistent (indeed it lasts forever).

The main goal of this paper is to build a model in which differences in subjective beliefs can be a highly persistent source of differences between equilibrium prices and fundamentals but where we can capture the idea of convergence of subjective beliefs upon observing the same information. I also want to capture the possibility of prices growing above the fundamentals even after a period of apparent convergence as a result of some events which were considered unlikely by the agents.

The model developed here is based on the example of Harrison and Kreps. I consider two risk neutral agents who trade an asset and face a short sale constraint. I introduce a two layer probabilistic structure by an adding additional (unobservable) process, called regime, which we assume has a Markov structure. The dividend process (now separated from the state) is independent over time but its distribution is determined by the underlying regime. Agents observe the dividends, infer information about the underlying regime, and

update their priors. Using this 2-level probabilistic structure captures the idea of eventual disappearance of the asymmetry in the subjective beliefs but allows a lot of freedom in the dynamics and structure of that learning.

With this Markov structure I introduce a systematic way of dealing with asymmetric beliefs economies, using the notion of belief-based recursive equilibrium. To my knowledge this idea is pretty original. Intuitively speaking, beliefs-based recursive equilibrium means that the agents' trading decision in each period is a function of the current beliefs rather than the whole history of events together with the initial beliefs. The current beliefs represent what is essential in the current state of the world in the sense that they capture the probabilistic behavior of all future events.

This allows me to map the consumers' problems of such economies into the stochastic dynamic programming techniques a'la Stokey et al. (2004). The main result showing that mapping is Theorem 1. This provides a powerful tool for numerical computation of such equilibria and offers a useful analytical tool for checking up if a given environment features speculation.

I also construct examples to illustrate the way these techniques work. In one example I prove that there is no speculation in any recursive equilibrium, no matter what initial beliefs are. In another setup I show that, for some particular initial belief, the equilibrium price has to feature speculation.

An interesting feature of this example is that it is significantly different from the one of Harrison and Kreps. Here agents have exactly the same transition matrix. They only differ in their beliefs about the current position of the regime. The fact that this leads to speculation is pretty significant, because learning about the current regime is much more difficult than learning about the transition matrix. The thing is that the regime changes over time, while the transition matrix stays the same. This will lead to the proposed persistence of the speculative behavior.

Section 2 reviews some of the preceding literature. Section 3 introduces the mathematical model of the economy, describes its probabilistic and informational structure, the agents' problems, and the trading mechanism, defines equilibrium, and speculation.

Section 4 switches to a recursive analysis. First I introduce the concept of recursive equilibrium, and then in Theorem 1 I prove that any recursive equilibrium is an equilibrium in the sequential sense from Section 3. Then using the linearity of preferences and a no arbitrage condition I am able to simplify the condition for the recursive equilibrium price system to be the fixed point of some well-behaved operator, which is not only a contraction but also features some monotone property which enables me to develop an easy criterion for the existence of speculation in a given economy. This is done in Proposition 2. In the last section I present the two examples mentioned before.

2 Related Literature

There is a broad literature on the topic of speculation. The main approaches to model it were bubbles (as in the overlapping generation model) or using asymmetric information among agents. I will go over these few papers which are important to put my model into a context and to underline how it extends the existing literature.

Here, I will be only interested in models with asymmetric information. These models can be divided them into two major groups. The criterion for the division is the source of the information heterogeneity. The first group features the heterogeneity resulting from observing different signals by different agents. It means that the agents are facing different information partitions. In the second group the agents' information partitions are the same, but they have different prior assumptions about the distribution of unobservable parameters of a model. This is the group my model belongs to.

Let us start with differentiated signals models. The general problem with that approach is the well-known no-trade theorem by Milgrom and Stokey (1982). The theorem essentially says, that rational, infinitely lived agents operating on a frictionless market, who know that other agents get different signals will always take that knowledge into account and the equilibrium price will be always reflecting the fundamental value of an asset. Various researchers were using different techniques to deal with that problem. Usually these models are adding some frictions to the market mostly by analyzing the market microstructure. Probably the most influential model within that group is the one by Kyle (1985). In that paper some group of additional agents is introduced. These agents are called *the noise traders*. Their behavior is exogenous to the model. They just provide some stochastic demand (or supply) for the asset which is not related to the price. The mechanism of getting around the no-trade theorem here is that, vaguely speaking, the rational agents are now playing a positive-sum game (they "make use" of the noise traders).

Another recent model of speculation driven by heterogeneous signals is the one of Bacchetta and Wincoop (2006), which has been shown to be a good tool to explain the exchange rate puzzle (cf. Bacchetta and Wincoop (2007)). In their model they consider infinitely many generations of finitely lived agents, who inherit their information within each generation. Each of the generations is observing a different signal each period, which conveys information about future T dividends. The market stock supply is also stochastic. In order to find the price clearing the market, the market maker must take into account the expected value of future expected values of dividends for agents (higher order expectations). Using CARA utility function and a normal environment they are able to get an explicit solution to their model, which features speculative behavior. In their model the heterogeneity of agents comes purely from differentiated signals. The main reason for which they can go around the no-trade theorem is that even though the agents pass information within each generation, they don't care about future generations' utility — they just solve a 1-period problem.

Now let us turn to the models, where the common prior beliefs assumption is waived. This trend in the literature was started by the paper by Harrison and Kreps (1978). This paper I discussed in detail in the introduction and also throughout the paper in examples. It is perhaps worth underlining here that the force responsible for speculation there is not only asymmetric beliefs but also the short sale constraint.

Morris (1996) considers a special case of Harrison and Kreps, with a iid dividend process. This enables the agents to learn over time. Using some parametric classes of invariant distributions for the priors (like β -distributions) he gets a nice explicit formula for the learning dynamics. Also he gets a nice criterion for having speculation in equilibrium. He indeed gets some speculative behavior but the numerical experiments show that the speed of the convergence of the equilibrium price to the fundamental value is very fast, so this model can be only used to model speculation which arises only after some new asset is introduced.

An interesting attempt to control the convergence of equilibrium prices to rational expectations is Bossaerts (1995). Here there is no dividend bearing, infinitely lived asset. There are only 1-period future contracts with risky return. The payoffs of these contracts are iid over time. There are countably many generations of agents. The beliefs are shared and updated within generations. Each period a new generation of agents joins the market and stays there forever. The new generation comes with its own initial beliefs, which are immediately updated by the up to date stream of returns. He assumes that the returns are normal with mean zero and the unknown variance, and that the beliefs about the variance are inverted gamma-2 distributions. This specification allows for an easy analytical treatment and a flexible control over the equilibrium price dynamics. It is easy to get that the initial price exceeds the rational expectation one. The conditions for the convergence of the prices to the rational expectation values are given. For some choice of beliefs we can get no convergence, which gives a powerful tool to control the rate of convergence. All of this is done at the expense of having new coming generations with more and more biased initial beliefs. Also an important role is played by the fact that the agents' problem is not dynamic (the future contract is only for 1 period).

The most recent paper in that spirit is the one of Scheinkman and Xiong (2003). In their model the (cumulative) dividend follows a diffusion process, with drift f_t , which is called a fundamental variable and is not observed by the agents, they only know it follows another diffusion process. Even though they use continuous time diffusion process techniques, their model can be treated almost as a special case of mine (after appropriate discretization of their setup or redesigning mine to cover the continuous time case). There is one crucial difference though. Given the normal environment of them, which is easy to deal with analytically, it is pretty hard to obtain a speculation in the case when the agents only observe the cumulative dividend as a common signal. To fix that problem Scheinkman and Xiong consider additional signal processes s_t^A and s_t^B , which are both diffusion processes with f_t as their drift part. As for the innovations part, A believes that the one of s_t^A is correlated with the one for the process f_t while agent B thinks that it is the one of s_t^B . So agent A even though he can observe both signals, he thinks his signal has a better quality than the signal of agent B and vice versa. The model stated in that way can be explicitly solved analytically and features speculation. The problem is that it requires an additional signalling structure (besides dividends), and also even though agents are updating their beliefs about the underlying fundamental process, f_t , they are not learning about the informativeness of the signals and always use their own one for updating.

The main contribution of my paper is to extend the existing asymmetric beliefs literature by introducing and exploring the concept recursive equilibrium, which is pretty novel in a presence of learning. Also, using hidden Markov process for dividends makes the speculation relatively persistent even though the beliefs are converging.

3 The Model

3.1 Economy

There are 2 agents, who are endowed with zero units of consumption good at each time period, $t = 0, 1, \dots$ they are both risk neutral and have a discount factor β .

There is 1 unit of risky asset in this economy, which agent can trade each period in general equilibrium fashion with no short sales allowed. The asset gives to its owner a dividend d_t each period.

There is an underlying regime process, a_t , taking value in some state space \mathcal{A} . We assume a_t is Markov and it cannot be directly observed by the agents.

The dividend, d_t , is generated independently each period from a distribution which depends on the current regime, $a_t \in \mathcal{A}$. The distribution associated with regime a_t we denote Φ_{a_t} .

Now let us turn to describing these processes formally.

First, to fix ideas, I will denote (Ω, \mathcal{F}) an abstract measurable space over which all the random variables in this paper will be defined.

Let the set of possible regimes, \mathcal{A} , has a structure of the Polish space and the regime process, $(a_t)_{t=0,1,\dots}$ be a stationary Markov process with the transition function $q : \mathcal{A} \rightarrow \Delta(\mathcal{A})$ assumed Borel-measurable, with $\Delta(\mathcal{A})$ denoting the linear space of all Borel probabilistic measures over \mathcal{A} endowed with weak*-topology.

Let $\mathcal{D} \subseteq \mathbb{R}$ denote the set of possible dividends (assume it is Borel-measurable), and let $\{\Phi_a\}_{a \in \mathcal{A}}$ be a family of probability distributions over \mathcal{D} (i.e. $\Phi_a \in \Delta(\mathcal{D})$ for each $a \in \mathcal{A}$), such that $\Phi : \mathcal{A} \rightarrow \Delta(\mathcal{D})$ is Borel-measurable (with respect to weak*-topology on $\Delta(\mathcal{D})$). We will also need to assume (in order to be able to use Bayes' rule) that Φ_a has a density function with respect to some regular measure on \mathbb{R} , μ (usually either discrete or Lebesgue). Denote this density by ϕ_a .

Having specified $\bar{a}_0 \in \mathcal{A}$, the family $\Phi \in \mathcal{B}(\mathcal{A}, \{\zeta \in \Delta(\mathcal{D}) | \zeta \ll \mu\})$, and the transition function $q \in \mathcal{B}(\mathcal{A}, \Delta(\mathcal{A}))$ (I follow the convention of denoting the linear space of Borel-measurable functions by $\mathcal{B}(\cdot, \cdot)$) we denote by $\Pr^{\bar{a}_0, \Phi, q}$ a probability measure over (Ω, \mathcal{F}) which is consistent with the Markov structure of the process a_t and with the described structure of the process d_t (i.e. such that $a_0 = \bar{a}_0$ with probability 1, a_t is Markov with the transition function q and d_t is drawn independently each period from the distribution Φ_{a_t}). Formally, for each $A_0 \in \mathcal{B}(\mathcal{A}), \dots, A_t \in \mathcal{B}(\mathcal{A}), D_1 \in \mathcal{B}(\mathcal{D}), \dots, D_t \in \mathcal{B}(\mathcal{D})$ we have:

$$\begin{aligned} \Pr^{\bar{a}_0, \Phi, q}(a_0 \in A_0, \dots, a_t \in A_t, d_1 \in D_1, \dots, d_t \in D_t) &= \\ &= \int_{A_0, \dots, A_t} \Phi_{a_1}(D_1) \dots \Phi_{a_t}(D_t) \delta_{\{\bar{a}_0\}}(da_0) q(a_0, da_1) \dots q(a_{t-1}, da_t) \\ &= \int_{\substack{A_0, \dots, A_t \\ D_1, \dots, D_t}} \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t) \delta_{\{\bar{a}_0\}}(da_0) q(a_0, da_1) \dots q(a_{t-1}, da_t) \mu(dd_1) \dots \mu(dd_t) \end{aligned}$$

Now let's turn to agent's information structure. They both can observe dividend d_t each period and none of them can observe a_t . Denote by $(\mathcal{F}_t^d)_t$ the filtration generated by the process d_t .

In particular the agents don't know the initial regime, a_0 and also they don't know the family of distributions Φ . or the transition function q . We will assume that Φ . and q can take values in some Borel sets of admissible values, $\Phi \subseteq \mathcal{B}(\mathcal{A}, \{\zeta \in \Delta(\mathcal{D}) | \zeta \ll \mu\})$, and $\mathcal{Q} \subseteq \mathcal{B}(\mathcal{A}, \Delta(\mathcal{A}))$, respectively. Hence the agents formulate beliefs about the value of $(a_0, \Phi., q)$.

Let \Pr^π be a measure that a player with beliefs $\pi \in \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})$ assigns to (Ω, \mathcal{F}) . This is clearly given by:

$$\Pr^\pi(F) = \int_{\mathcal{A} \times \Phi \times \mathcal{Q}} \Pr^{\bar{a}_0, \Phi, q}(F) \pi(d\bar{a}_0, d\Phi, dq)$$

for any $F \in \mathcal{F}$. Let E^π be the expected value operator associated with that measure. Denote the initial beliefs of player i by $\pi_0^i \in \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})$.

The timing is as follows. In the beginning of period t the new position in the Markov process a_t is established, then dividend d_t is generated and paid to the current owner of the asset. After the dividend is paid the agents can trade on the centralized market for price p_t , subject to the short sale constraint.

Now we are ready to define (competitive) equilibrium.

Definition. An *equilibrium* (given initial beliefs, π_0^1, π_0^2) consists of the processes: an allocation, $(\hat{c}_t^i)_t$, asset holdings, $(\hat{\gamma}_t^i)_t$, and prices $(p_t)_t$ such that prices are \mathcal{F}_t^d -adapted, and:

- For $i = 1, 2$, taking $(p_t)_t$ as given, $(\hat{c}_t^i)_t$ and $(\hat{\gamma}_t^i)_t$ solve:

$$\begin{aligned} \max E^{\pi_0^i} \sum \beta^t c_t^i & \quad (1) \\ \text{s.t. } \hat{c}_t^i + p_t \gamma_{t+1}^i & \leq p_t \gamma_t^i + \gamma_t^i d_t \\ \text{s.t. } \hat{c}_t^i, \gamma_{t+1}^i & \text{--- } \mathcal{F}_t^d\text{-measurable} \\ \text{s.t. } \gamma_0^i = 0, \gamma_t^i & \geq 0 \end{aligned}$$

- Asset market clears: $\gamma_t^1 + \gamma_t^2 = 1$

It is worth noting at this point, that it is the implicit feature of this general equilibrium environment, that agents are facing prices as functions of all potential histories not beliefs. Even though the equilibrium prices, in order to clear the market, have to convey the information about all the agents beliefs, agents do not need to know the beliefs of the others. That information, however, can be often inferred from the prices. In either case in this Walrasian type of equilibrium, where agents take prices as given the structure of higher order beliefs seems to be irrelevant.

3.2 Speculation

In this setup, it is the most natural way to define speculation as the situation in which the equilibrium price exceeds the fundamental valuation of the asset for the agent for whom it is the highest. By fundamental valuation of agent i we understand the discounted stream of all the future dividends expected by agent i . This is the highest price he would be willing to pay for the asset if he was forced to keep it forever after the purchase.

Definition. The fundamental value of the asset for agent i at time t , given history d^t is:

$$V_t^i(d^t) \equiv E^{\pi^i} \left(\sum_{\tau=t+1}^{\infty} \beta^\tau d_\tau | d^t \right)$$

Definition. We say that an equilibrium features speculation if there exist a time, t , and a history, d^t , such that $p(d^t) > \max_i V_t^i(d^t)$.

Example. In order to illustrate the notation introduced I will show how the example given by Harrison and Kreps fits into this notation.

In their model the dividend the dividend itself follows a 2-state Markov chain (can be either 0 or 1). Agents differ in what they think the transition matrix is. Agent 1 thinks the transition matrix is:

$$Q^1 = \begin{bmatrix} 1/2 & 1/2 \\ 2/3 & 1/3 \end{bmatrix}$$

agent 2 thinks it is:

$$Q^2 = \begin{bmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{bmatrix}$$

To map it into my notation it is enough to take $\mathcal{A} = \mathcal{D} = \{0, 1\}$, $\Phi = \{\Phi\}$ with $\Phi_0 = \delta_{\{0\}}$ $\Phi_1 = \delta_{\{1\}}$ (i.e. the current regime coincides with the current dividend and both agents agree about it). We also have $\mathcal{Q} = \{Q^1, Q^2\}$. As for a_0 we may assume the agents know it as it coincides with the dividend hence the agent's beliefs about this one coincide and put the whole measure on its true value. Hence the initial beliefs are just measures over $\mathcal{A} \times \mathcal{Q}$, specifically π_0^1 puts all the measure on (\bar{a}_0, Q^1) and π_0^2 puts all the measure on (\bar{a}_0, Q^2) (where a_0 is the true value of a_0), i.e. $\pi_0^i = \delta_{\bar{a}_0} \otimes \delta_{Q^i}$

We can see that since agents have disjoint supports in their beliefs they will not be learning the transition matrix over time. We will see in the next section that this lack of dynamics in the beliefs is crucial for having an explicit solution for equilibrium prices in the Harrison and Kreps example. In this case it is also straightforward to compute the fundamental values. It is clearly only the function of last period dividend (for each agent) because it is the sole factor to determine the future distribution of dividends: $V^i(d^t) = V^i(d_t)$. Denoting $V^i \equiv [V^i(0), V^i(1)]'$ and using recursiveness and Markov property we get that it has to satisfy:

$$V^i = \beta Q^i V + \beta \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

so

$$V^i = \beta (I - \beta Q^i)^{-1} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

In case of $\beta = 3/4$ we get:

$$V^1 = \begin{bmatrix} 4/3 \\ 11/9 \end{bmatrix} = \begin{bmatrix} 1.33 \\ 1.22 \end{bmatrix} \quad V^2 = \begin{bmatrix} 16/11 \\ 21/11 \end{bmatrix} = \begin{bmatrix} 1.45 \\ 1.91 \end{bmatrix} \quad (2)$$

4 Recursive Equilibrium

Since we assume rationality of agents (given their own initial beliefs) they must do learn from the observed signal in the Bayesian way. This creates certain dynamics of beliefs. I want to make new beliefs be only dependent on the last period beliefs and the current period dividends (rather than the whole history). To achieve that we introduce some new notation for the updated beliefs, and understand the current beliefs to be a distribution over the current position of the Markov process, a_t rather than the initial one. This will lead me to the notion of recursive equilibrium, which will appear to be a powerful tool in analysis of the equilibria in this model. That doing so makes the model consistent with the Bayesian learning requires some formal argument. It will be done in this section by proving that a recursive equilibrium can be translated to an equilibrium.

First let us define the learning operator, λ_t , which is crucial for applying the stochastic dynamic programming techniques.

Definition. Given $\pi \in \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})$, and $d^t \in \mathcal{D}^t$ define a measure $\lambda_t(d^t|\pi) \in \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})$, by

$$\lambda_t(d^t|\pi)(A \times \Phi \times Q) \equiv \Pr^\pi(a_t \in A \wedge \phi \in \Phi \wedge q \in Q | d^t)$$

for each measurable $A \subseteq \mathcal{A}, \Phi \subseteq \Phi, Q \subseteq \mathcal{Q}$

As a matter of fact I will only need the one-period update, $\lambda_1(\cdot|\cdot)$, hence, throughout the paper, I skip subscript '1' and denote it simply with $\lambda(\cdot|\cdot)$.

Example of Harrison and Kreps (1978) –cont. In the example of Harrison and Kreps we would have in this notation:

$$\lambda_t(d^t|\pi^i) = \delta_{d_t} \otimes \delta_{Q^i} \tag{3}$$

in words: agents each time update their beliefs about the current regime to be the true value of it and remain stuck with what they initially believed about the transition matrix.

Definition. A (symmetric) *recursive equilibrium* consists of a value function $V : [0, 1] \times \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2 \rightarrow \mathbb{R}$, policy function, γ' , and pricing function, $p : \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2 \rightarrow \mathbb{R}$, such that:

$$V(\gamma, \pi^1, \pi^2) = \max_{\gamma' \geq 0} \{ (\gamma - \gamma')p(\pi^1, \pi^2) + \beta E^{\pi^1} (V(\gamma', \lambda(d_1|\pi^1), \lambda(d_1|\pi^2)) + \gamma' d_1) \}$$

and for each γ, π^1, π^2, d we have:

$$\gamma'(\gamma, \pi^1, \pi^2) + \gamma'(1 - \gamma, \pi^2, \pi^1) = 1$$

For notational simplicity it is useful to denote the beliefs process by π_t^i , which is defined recursively via:

$$\begin{aligned} \pi_1^i &\equiv \lambda(d_1|\pi_0^i) \\ \pi_{t+1}^i &\equiv \lambda(d_{t+1}|\pi_t^i) \end{aligned}$$

Theorem 4.1. *Given some initial beliefs, $\pi_0^1, \pi_0^2 \in \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2$, if $V : [0, 1] \times \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2 \rightarrow \mathbb{R}_+$, $\gamma', p : \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2 \rightarrow \mathbb{R}$ constitute a recursive equilibrium then the processes: $p_t^* \equiv p(\pi_t^1, \pi_t^2)$, $\gamma_{t+1}^{*i} \equiv \gamma'(\pi_t^i, \pi_t^{-i})$, $c_t^{*i} \equiv p_t(\gamma_t^i - \gamma_{t+1}^i) + \gamma_t^i d_t$ are an equilibrium, given beliefs π_0^1, π_0^2 id additionally the transversality condition holds:*

$$\lim_{t \rightarrow \infty} \beta^t E^{\pi_0^i} (V(\gamma_t^i, \pi_t^i, \pi_t^{-i}) + \gamma_t^i d_t) = 0$$

The proof will follow from the following lemma:

Lemma 4.2. *For each $d^{t+s} \in \mathcal{D}^{t+s}$, $A_0, \dots, A_s \subseteq \mathcal{A}$, $\Phi \subseteq \Phi$, $Q \subseteq \mathcal{Q}$, such that A_0, \dots, A_s, Φ, Q are measurable subsets, we have:*

$$\begin{aligned} & \Pr^\pi(a_t \in A_0 \wedge \dots \wedge a_{t+s} \in A_s \wedge \phi \in \Phi \wedge q \in Q | d^t) \\ &= \Pr^{\lambda_t(d^t | \pi)}(a_0 \in A_0 \wedge \dots \wedge a_s \in A_s \wedge \phi \in \Phi \wedge q \in Q) \end{aligned}$$

This lemma states that all the future distribution of all the processes, at time t from the perspective of player i is completely described by the updated measure $\pi_t^i = \lambda_t(d^t | \pi_0^i)$. This legitimates the introduction of the recursive equilibrium in this environment.

The proof of theorem 4.1 as well as that of lemma 4.2 are relegated to the appendix as they are purely technical.

4.1 Characterization of recursive pricing rule

Using linearity of preferences and no short sales condition, we can argue that a pricing function $p : \Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2 \rightarrow \mathbb{R}$ is the pricing rule of some recursive equilibrium iff it satisfy the following first order condition to the Bellman equation:

$$p(\pi^1, \pi^2) = \max_{i=1,2} \beta E^{\pi^i} (p(\lambda(d' | \pi^1), \lambda(d' | \pi^2)) + d')$$

Also note, that in order to reflect the switch to the recursive analysis, instead of using time subscript '1' in d , I use d' to denote the next period dividend.

In order to organize the notation let's define the following operators, $T, T^1, T^2 : \mathcal{B}(\Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2, \mathbb{R}) \rightarrow \mathcal{B}(\Delta(\mathcal{A} \times \Phi \times \mathcal{Q})^2, \mathbb{R})$, with:

$$\begin{aligned} T^i p(\pi^1, \pi^2) &\equiv \beta E^{\pi^i} \{d' + p(\lambda(d' | \pi^1), \lambda(d' | \pi^2))\} & i = 1, 2 \\ Tp &\equiv \max_{i=1,2} T^i p \end{aligned}$$

With this notation, the equation for prices becomes:

$$p = Tp \tag{4}$$

so we are just looking for a fixed point of T .

Also note, that the the fixed point of operator T^i is the fundamental value for agent i , V^i .

Example of Harrison and Kreps (1978) – cont. Now we can see how much the lack of learning facilitates the solution of the functional equation. Using (3) the formula for operator T^i becomes (note that without loss of generality we can treat prices as a function of a current dividend since so are the beliefs):

$$\begin{aligned} T^i p(d) &= \beta E^{\pi^i} \{d' + p(d')\} \\ &= \begin{cases} \beta \{(1/2)p(0) + (1/2)(1 + p(1))\} & \text{for } d = 0, i = 1 \\ \beta \{(2/3)p(0) + (1/3)(1 + p(1))\} & \text{for } d = 0, i = 2 \\ \beta \{(2/3)p(0) + (1/3)(1 + p(1))\} & \text{for } d = 1, i = 1 \\ \beta \{(1/4)p(0) + (3/4)(1 + p(1))\} & \text{for } d = 1, i = 2 \end{cases} \end{aligned}$$

This allows us for almost immediate guess the solution to (4), which in the case of $\beta = 3/4$ is $p^*(0) = 24/13 = 1.85$ and $p^*(1) = 27/13 = 2.04$, which is clearly higher than the corresponding maximal fundamental values derived in (2).

Let us investigate some properties of T .

Lemma 4.3. *Operators T, T^i are all β -contractions w.r.t. the sup-norm.*

Proof. A direct application of Blackwell's sufficient conditions. □

Lemma 4.4. *If p is continuous, then Tp is continuous.*

As a corollary we get:

Theorem 4.5. *If $\beta < 1$ then there is the unique solution, p^* to the functional equation (4). Moreover, p^* is continuous and $p^* = \lim_{t \rightarrow \infty} T^t p$ (in the sup norm).*

It is worth noting, that there is no hope for more general regularity conditions for the price system beyond continuity. The following examples will show the lack of differentiability, while any kind of convexity seems meaningless in this setup (at least in general). Nevertheless, the following monotone property of operator T appears to be useful.

Lemma 4.6. *If $Tp \geq p$ for some price system p , then $T^2 p \geq Tp$. Hence also $p^* \geq p$*

Proof. For any beliefs π^1, π^2 we have:

$$\begin{aligned} T^2 p(\pi^1, \pi^2) &= \beta \max_{i=1,2} \int_{\mathcal{D}} p^{\pi^i}(d')(d' + Tp(\lambda(d'|\pi^1), \lambda(d'|\pi^2))) dd' \\ &\geq \max_{i=1,2} \int_{\mathcal{D}} p^{\pi^i}(d')(d' + p(\lambda(d'|\pi^1), \lambda(d'|\pi^2))) dd' \\ &= Tp(\pi^1, \pi^2) \end{aligned}$$

with $p^{\pi^i}(d')$ denoting probability density of next period dividend according to an agent with beliefs π^i i.e.

$$p^{\pi^i}(d') = \int_{\mathcal{A} \times \Phi \times \mathcal{Q}} \int_{\mathcal{A}} \phi_{a'}(d') q(a, da') \pi^i(da, d\phi, dq)$$

□

As a corollary we get a useful fact.

Theorem 4.7. *If we define the fundamental pricing rule by $p^f = \max_{i=1,2} V^i$ then if for some beliefs (π^1, π^2) , $Tp^F(\pi^1, \pi^2) > p^F(\pi^1, \pi^2)$, then $p^*(\pi^1, \pi^2) > p^F(\pi^1, \pi^2)$, i.e. we have speculation in equilibrium for those initial beliefs.*

Proof. It is clear, that for fundamental pricing we have $Tp^F \geq p^F$. From the previous lemma we get, that $T^2p^F \geq Tp^F$ hence using this lemma again we get that $p^* \geq Tp^f$. Hence by our assumption $p^*(\pi^1, \pi^2) \geq Tp^F(\pi^1, \pi^2) > p^F(\pi^1, \pi^2)$ \square

This theorem gives us an easy tool to check if we have a speculation in a given economy. Just take an initial guess for pricing system to be $p^f = \max_{i=1,2} V^i$ (the highest fundamental value). Then iterate it once. Obviously we must have $Tp^f \geq p^f$ (if the agents are promised to be able to resell the asset at the highest fundamental price next period then in the current period they must be willing to pay at least their fundamental values). If we get $Tp^f = p^f$ then p^f is the solution to the functional equation (4) i.e. $p^* = p^f$ and we have no speculation. Otherwise there are some beliefs, π^1, π^2 for which $Tp^f(\pi^1, \pi^2) > p^f(\pi^1, \pi^2)$, which by proposition 2 implies that $p^*(\pi^1, \pi^2) \geq Tp^f(\pi^1, \pi^2) > p^f(\pi^1, \pi^2)$, which means that we have a speculation in equilibrium.

5 Examples

In this section I will consider the environment in which the agents' only potential disagreement is about the current regime.

I present two examples. The first one has 2-state regime process where the states are stable in the sense that in each of them probability of staying in it is bigger than moving out of it. In that case I prove that no speculation can exist.

In the second example I consider 3-state regime process. There are two "good" states and one "bad" state. In good states agents get relatively high probability of dividend (around 2/3) but if the state is bad the probability of dividend is 0. The "high" regimes are different in terms of probability of switching to the "bad" regime. These probabilities are both low but one is smaller than the other. Agents initially think they are in a "good" regime but of a different type. Here I will be able to prove that I can construct arbitrary small transition probabilities so that the equilibrium price features speculation. In that example we will also see that the speculation may persist arbitrarily long even though the agents are learning.

5.1 A Simplified Environment

Here the agents agree upon the value of the transition function as well as upon the distribution of dividends. They just disagree about the current regime. Formally this means that we take consider the class of models in which: $\mathcal{A} = \{a_1, \dots, a_n\}$, $\mathcal{D} = \{0, 1\}$, $\Phi = \{\phi\}$ with

$$\phi_{a_j}(1) = 1 - \phi_{a_j}(0) \equiv \phi_j, \mathcal{Q} = \{q\}, \text{ with } q = \begin{bmatrix} q_{11} & \cdots & q_{1n} \\ \vdots & \ddots & \vdots \\ q_{n1} & \cdots & q_{nn} \end{bmatrix} \text{ (Markov chain transition matrix).}$$

In words this setup means that there is an underlying \mathcal{A} -valued Markov chain in the economy, $a_0, a_1, \dots, a_t, \dots$ with the known transition matrix, q . Each period dividend d_t is paid with probability ϕ_j which is determined solely by the current position in the Markov chain, $a_t = a_j$. The agents formulate beliefs about the initial position of the Markov chain, a_0 . They both know the value of ϕ and q , and over time they beliefs evolve in the Bayesian fashion. The current beliefs of agent i are $\pi^i = [\pi_1^i, \dots, \pi_n^i] \in \Delta(\mathcal{A}) \simeq \Delta^{n-1}$.

In this setup the updating dynamics as well as finding the fundamental values become relatively easy. It is straightforward to see that the fundamental value of agent i , $V^i = \pi^i \cdot V$, where $V = [V_1, \dots, V_n]'$ is the vector of fundamental values for each initial position in the Markov chain. V is the solution to

$$(I - \beta q)V = \beta q \phi$$

The fundamental price becomes:

$$p^f = \max_{i=1,2} \pi^i V$$

Given the initial beliefs π^i and the current period dividend, d' the new beliefs are (using Bayes' rule):

$$\lambda(d'|\pi^i) = \begin{cases} \left[\frac{\phi_1 \sum_{j=1}^n q_{j1} \pi_j^i}{\sum_{j'=1}^n \phi_{j'} \sum_{j=1}^n q_{jj'} \pi_j^i}, \dots, \frac{\phi_n \sum_{j=1}^n q_{jn} \pi_j^i}{\sum_{j'=1}^n \phi_{j'} \sum_{j=1}^n q_{jj'} \pi_j^i} \right] & \text{if } d' = 1 \\ \left[\frac{(1-\phi_1) \sum_{j=1}^n q_{j1} \pi_j^i}{\sum_{j'=1}^n (1-\phi_{j'}) \sum_{j=1}^n q_{jj'} \pi_j^i}, \dots, \frac{(1-\phi_n) \sum_{j=1}^n q_{jn} \pi_j^i}{\sum_{j'=1}^n (1-\phi_{j'}) \sum_{j=1}^n q_{jj'} \pi_j^i} \right] & \text{if } d' = 0 \end{cases} \quad (5)$$

Also it will be useful to have the explicit expression for the probability of dividend next period, $d' = 1$, given the current beliefs are π^i :

$$\Pr^{\pi^i} \{d' = 1\} = \pi^i q \phi \quad (6)$$

The formula for operator T becomes:

$$Tp(\pi^1, \pi^2) = \beta \max_{i=1,2} \left[\Pr^{\pi^i} \{d' = 1\} (1 + p(\lambda(1|\pi^1), \lambda(1|\pi^2))) \right. \\ \left. + \left(1 - \Pr^{\pi^i} \{d' = 1\} \right) p(\lambda(0|\pi^1), \lambda(0|\pi^2)) \right]$$

5.2 Example with no speculation

Here we will see a situation in which we will actually solve the functional equation. The solution will be (as one may expect) the fundamental valuation by the agent for whom it's the highest (given beliefs).

Proposition 3. If $\mathcal{A} = \{h, l\}$ and $q = \begin{bmatrix} q_{hh} & q_{hl} \\ q_{lh} & q_{ll} \end{bmatrix} = \begin{bmatrix} 1 - \epsilon^1 & \epsilon^1 \\ \epsilon^2 & 1 - \epsilon^2 \end{bmatrix}$ satisfies $\epsilon^1 + \epsilon^2 \leq 1 < 1$, then the equilibrium price is equal to the fundamental price: $p^* = p^F$ for all beliefs.

This proposition states that if we have only 2 regimes there will be no speculation. It seems that only two regimes cannot give enough room for disagreement if we have only 2 signals (the dividend either paid or not).

Proof. Since we have only 2-state Markov chain, the beliefs can be just represented by one number: probability of being in a given state. To fix ideas let it be state h. Hence the beliefs are: $\pi^i \in [0, 1]$, $i = 1, 2$ and satisfy: $\Pr^{\pi^i} \{a_0 = h\} = \pi^i$.

As usual, denote: $\phi = [\phi_h, \phi_l]'$ to be the vector of the probabilities of the dividend in regime.

I will show that the equilibrium price is just the fundamental price, i.e. $Tp^f = p^f$ (for all beliefs).

Without loss of generality assume that the fundamental vector $V = [V_h, V_l]'$ satisfies $V_h > V_l$ (otherwise just relabel the states). We have then

$$\begin{aligned} p^f(\pi^1, \pi^2) &= \max_{i=1,2} \{ \pi^i V_h + (1 - \pi^i) V_l \} \\ &= (\max_{i=1,2} \pi^i) V_h + (1 - \max_{i=1,2} \pi^i) V_l \end{aligned}$$

by symmetry of the agents, wlog I can always assume $\pi^1 \geq \pi^2$ (otherwise just renumber them), which leads us to

$$p^f(\pi^1, \pi^2) = \pi^1 V_h + (1 - \pi^1) V_l \quad (7)$$

Now I will show that $\lambda(d|\pi)$ is increasing in π for $d = 0, 1$ (in words: agent who was more optimistic in the first period will always remain more optimistic in the next period, no matter which dividend occurred). By (5) we have:

$$\lambda(d|\pi) = \begin{cases} \frac{\phi_h [(1-\epsilon^1)\pi + \epsilon^2(1-\pi)]}{\phi_h [(1-\epsilon^1)\pi + \epsilon^2(1-\pi)] + \phi_l [\epsilon^1\pi + (1-\epsilon^2)(1-\pi)]} & \text{for } d = 1 \\ \frac{\phi_l [(1-\epsilon^1)\pi + \epsilon^2(1-\pi)]}{\phi_l [(1-\epsilon^1)\pi + \epsilon^2(1-\pi)] + \phi_h [\epsilon^1\pi + (1-\epsilon^2)(1-\pi)]} & \text{for } d = 0 \end{cases} \quad (8)$$

A bit of algebra gives:

$$\frac{\partial}{\partial \pi} \lambda(d|\pi) = \frac{1 - \epsilon^1 - \epsilon^2}{(\text{appropriate denominator expression} \neq 0)^2} > 0 \quad d = 0, 1$$

so indeed for $\epsilon^1, \epsilon^2 < 1/2$ $\lambda(d|\pi)$ is increasing in π . Hence $\pi^1 \geq \pi^2$ implies that also $\lambda(d|\pi^1) \geq \lambda(d|\pi^2)$.

Another thing I am going to need is $\Pr^{\pi^1}(d' = 1) \geq \Pr^{\pi^2}(d' = 1)$ (always assuming $\pi^1 \geq \pi^2$), which is easy to verify using (6). Also I will need $\lambda(1|\pi) > \lambda(0|\pi)$, which is intuitively obvious and straightforward to check from (8).

This gives us:

$$\begin{aligned}
Tp^f &= \beta \max_{i=1,2} \left\{ \Pr^{\pi^i} \{d = 1\} [1 + p^f (\lambda(1|\pi^1), \lambda(1|\pi^2))] \right. \\
&\quad \left. + \left(1 - \Pr^{\pi^i} \{d = 0\}\right) p^f (\lambda(0|\pi^1), \lambda(0|\pi^2)) \right\} \\
&= \beta \max_{i=1,2} \left\{ \Pr^{\pi^i} \{d = 1\} (1 + \lambda(1|\pi^1)V_h + (1 - \lambda(1|\pi^1))V_l) \right. \\
&\quad \left. + \left(1 - \Pr^{\pi^i} \{d = 0\}\right) (\lambda(0|\pi^1)V_h + (1 - \lambda(0|\pi^1))V_l) \right\} \\
&= \beta \Pr^{\pi^1} \{d = 1\} (1 + \lambda(1|\pi^1)V_h + (1 - \lambda(1|\pi^1))V_l) \\
&\quad + (1 - \Pr^{\pi^1} \{d = 0\}) (\lambda(0|\pi^1)V_h + (1 - \lambda(0|\pi^1))V_l) \\
&= T^1 V^1 \\
&= V^1 \\
&= \max_{i=1,2} \{V^1, V^2\} \\
&= p^f
\end{aligned}$$

where the second line comes from (7), the third uses the fact that $\Pr^{\pi^1}(d' = 1) \geq \Pr^{\pi^2}(d' = 1)$ and $\lambda(1|\pi) > \lambda(0|\pi)$. Then we use the fact that V^i is the fixed point of operator T^i . this allows us to conclude that p^f is the equilibrium price for all beliefs. \square

5.3 Example with speculation

Let $\mathcal{A} = \{h_1, h_2, l\}$, $\phi = [\Phi_1, \Phi_2, 0]'$ and

$$q = \begin{bmatrix} q_{h_1 h_1} & q_{h_1 h_2} & q_{h_1 l} \\ q_{h_2 h_1} & q_{h_2 h_2} & q_{h_2 l} \\ q_{l h_1} & q_{l h_2} & q_{ll} \end{bmatrix} = \begin{bmatrix} 1 - \epsilon^1 & 0 & \epsilon^1 \\ 0 & 1 - \epsilon^2 & \epsilon^2 \\ 0 & 0 & 1 \end{bmatrix}$$

This setup means that we have 2 'good' regimes, h_1, h_2 , and one 'bad' regime, l . In regime h_1 the probability of dividend is $\Phi_1 > 0$ and the probability of switching to the bad regime is ϵ^1 . In regime h_2 the probability of dividend is Φ_2 and the probability of switching to the bad regime is ϵ^2 . In bad regime l there are no dividends ($\Phi_3 = 0$) and this state is absorbing.

In this setup we can readily get an interesting result.

Proposition 5.1. *If $\Phi_1, \epsilon^1, \Phi_2, \epsilon^2$ are such that $V_1 = V_2$, $\pi^1 = (1, 0, 0)$, $\pi^2 = (0, 1, 0)$, and $\Phi_1 \neq \Phi_2$ (the agents' valuations are exactly the same but the beliefs differ) then we have speculation, namely $p^*(\pi^1, \pi^2) > p^F(\pi^1, \pi^2) = V_1 = V_2$.*

This proposition says, that whenever the parameters are such that both agents' valuations of the asset are exactly the same but their beliefs about the probabilistic structure of dividends differ in any way, then there must be some speculation going on. The intuition behind this result is that if agents agree upon the discounted present value of the stream

of the future dividends, then in order to have different probabilistic structure of them one of them must have a higher probability of dividend in 'his' good state (Φ_i), which must be compensated by a lower higher probability of switching into the low state (ϵ_i). Also, a simple algebra shows that the agent with the higher Φ_i must also have a higher probability of seeing a dividend the next period. This means that his willingness of buying the asset today and selling it for its fundamental value tomorrow must be higher than that of the other agent. We know that any agent's willingness to buy it today with the option of resell it tomorrow must be at least his own fundamental value, V_i . Finally, using that $V_1 = V_2$ we conclude that the agent with the higher Φ_i must be willing to pay for the asset more than $V_1 = V_2$, hence we must have speculation.

Proof Without loss of generality assume $\Phi_1 > \Phi_2$. We have:

$$V = \begin{bmatrix} V_1 \\ V_2 \\ V_3 \end{bmatrix} = \begin{bmatrix} \frac{\beta(1-\epsilon^1)\Phi_1}{1-\beta(1-\epsilon^1)} \\ \frac{\beta(1-\epsilon^2)\Phi_1}{1-\beta(1-\epsilon^2)} \\ 0 \end{bmatrix}$$

Hence the condition $V_1 = V_2$ implies (after some rearrangements):

$$\frac{\Phi_2}{\Phi_1} = \frac{\frac{1}{1-\epsilon^2} - \beta}{\frac{1}{1-\epsilon^1} - \beta}$$

Then by $\Phi_1 > \Phi_2$ we get that:

$$\frac{1}{1-\epsilon^2} - \beta < \frac{1}{1-\epsilon^1} - \beta$$

which implies $\epsilon^1 > \epsilon^2$. Also note, that we can rearrange the condition $V_1 = V_2$ in another way to get:

$$\frac{(1-\epsilon^1)\Phi_1}{(1-\epsilon^2)\Phi_2} = \frac{1-\beta(1-\epsilon^1)}{1-\beta(1-\epsilon^2)}$$

But because $\epsilon^1 > \epsilon^2$, the RHS of the above is bigger than 1 hence we have:

$$(1-\epsilon^1)\Phi_1 > (1-\epsilon^2)\Phi_2$$

Now we will see that $T^1 p^F(\pi^1, \pi^2) > T^2 p^F(\pi^1, \pi^2)$. Using the definition of T^i we get:

$$\begin{aligned} T^i p^F(\pi^1, \pi^2) &= \beta \left[\Pr^{\pi^i}(d' = 1)(1 + \max\{\lambda^1(\pi^1)V, \lambda^1(\pi^2)V\}) + \Pr^{\pi^i}(d' = 0) \max\{\lambda^0(\pi^1)V, \lambda^0(\pi^2)V\} \right] \\ &= \beta \left[(1-\epsilon^i)\Phi_i(1+V_1) + (1-(1-\epsilon^i)\Phi_i) \max \left\{ \frac{(1-\epsilon^1)(1-\Phi_1)}{1-(1-\epsilon^1)\Phi_1} V_1, \frac{(1-\epsilon^2)(1-\Phi_2)}{1-(1-\epsilon^2)\Phi_2} V_2 \right\} \right] \\ &= \beta \left[(1-\epsilon^i)\Phi_i(1+V_1) + (1-(1-\epsilon^i)\Phi_i) \max \left\{ \frac{(1-\epsilon^1)(1-\Phi_1)}{1-(1-\epsilon^1)\Phi_1}, \frac{(1-\epsilon^2)(1-\Phi_2)}{1-(1-\epsilon^2)\Phi_2} \right\} V_1 \right] \end{aligned}$$

Since $\frac{(1-\epsilon^1)(1-\Phi_1)}{1-(1-\epsilon^1)\Phi_1}, \frac{(1-\epsilon^2)(1-\Phi_2)}{1-(1-\epsilon^2)\Phi_2} < 1$, and $(1-\epsilon^1)\Phi_1 > (1-\epsilon^2)\Phi_2$ then indeed we must have $T^1 p^F(\pi^1, \pi^2) > T^2 p^F(\pi^1, \pi^2)$. We also must have that $T^2 p^F(\pi^1, \pi^2) \geq V_2$ (if agent 2 is

promised to be paid at least his own valuation tomorrow, then his willingness to pay for the asset must be at least his own valuation today, which is V_2). Hence we have shown that $T^1 p^F(\pi^1, \pi^2) > T^2 p^F(\pi^1, \pi^2) \geq V_2 = V_1 = p^F(\pi^1, \pi^2)$. Now by Proposition 2 we get $p^*(\pi^1, \pi^2) > p^F(\pi^1, \pi^2)$ hence we have speculation in equilibrium. ■

One can expect that this speculation may persist very long. Each time the agents observe $d = 1$ they know they cannot be in the state l , so system is reset (we are back to the initial beliefs, which we know lead to speculation). Indeed this speculation will last till we are finally settled in state l .

6 Conclusion

In this paper I construct a model of speculation which looks like a promising tool in modeling long lasting speculative behavior when investors are learning from data. The idea is that investors, even though they learn from data, sometimes have to wait for some particular stream of signals to learn about certain aspects of the underlying regime. In the last example we saw that each time the agents observe dividend 1 the system is almost reset. Hence in order to achieve convergence of beliefs they need to observe a sufficiently long stream of zeros, so that both agents get convinced that a bad regime really occurred (once it happens their beliefs are pretty much the same because we have only one bad regime).

It is clear that when calibrating this model we can use more regimes in order to be able to capture some more sophisticated states which can be distinguished only after observing some very specific sequence of signals. For such signals we may need to wait very long. This creates the persistence of speculation, which is somehow hidden before that particular sequence occurs. This would be a good explanation for bursting of market bubbles. Doing so seems like a natural direction of future research.

References

- [1] Philippe Baccheta and Eric van Wincoop. Can information heterogeneity explain the exchange rate determination puzzle. *The American Economic Review*, 96(3), 2006.
- [2] Philippe Baccheta and Eric van Wincoop. Higher order expectations in asset pricing. *Working Paper*, 2006.
- [3] David Blackwell and Lester Dubins. Merging of opinions with increasing information. *The Annals of Mathematical Statistics*, 33(3), 1962.
- [4] Peter Bossaerts. The econometrics of learning in financial markets. *Economic Theory*, 11(1), 1995.
- [5] J.Michael Harrison and David M. Kreps. Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92(2), 1978.
- [6] Albert S. Kyle. Continuous auction and insider trading. *Econometrica*, 53(6), 1985.

- [7] Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 12, 1982.
- [8] Stephen Morris. Speculative investor behavior and learning. *The Quarterly Journal of Economics*, 111(4), 1996.
- [9] Jose A. Scheinkman and Wei Xiong. Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6), 2003.
- [10] Nancy L. Stokey, Jr. Robert L. Lucas, and Edward C Prescott. *Recursive Methods in Economic Dynamics*. 7 edition, 2004.

Appendix

Proof of Lemma 4.2. First note, that since we defined:

$$\begin{aligned} \lambda_t(d^t|\pi)(A \times \Phi \times Q) &= \Pr^\pi(a_t \in A \wedge \phi \in \Phi \wedge q \in Q|d^t) \\ &= \frac{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_t \in A} \int_{a_0, \dots, a_{t-1} \in \mathcal{A}^t} \pi(da_0, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)}{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_0, \dots, a_t \in \mathcal{A}^{t+1}} \pi(da_t, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)} \end{aligned}$$

for each measurable $A \subseteq \mathcal{A}, \Phi \subseteq \Phi, Q \subseteq \mathcal{Q}$, then each time we integrate with respect to the measure $\lambda_t(d^t|\pi)$, we can do the following replacement under any integral (the quotes will not be needed under an actual integral):

$$\lambda_t(d^t|\pi)(da_t, d\phi, dq) = \frac{\int_{a_0, \dots, a_{t-1} \in \mathcal{A}^t} \pi(da_0, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)}{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_0, \dots, a_t \in \mathcal{A}^{t+1}} \pi(da_t, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)}$$

Using this we have:

$$\begin{aligned} \Pr_t^\lambda(d^t|\pi)(a_0 \in A_0, \dots, a_s \in A_s, \phi \in \Phi, q \in Q) &= \\ &= \int_{\phi \in \Phi} \int_{q \in Q} \int_{(a_0, \dots, a_s) \in A_0 \times \dots \times A_s} \lambda_t(d^t|\pi)(da_0, d\phi, dq) q(a_0, da_1) \dots q(a_{s-1}, da_s) \\ &= \int_{\phi \in \Phi} \int_{q \in Q} \int_{(a_t, \dots, a_{t+s}) \in A_0 \times \dots \times A_s} \frac{\int_{a_0, \dots, a_{t-1} \in \mathcal{A}^t} \pi(da_0, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)}{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_0, \dots, a_t \in \mathcal{A}^{t+1}} \pi(da_t, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)} \\ &= \frac{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_0, \dots, a_{t-1} \in \mathcal{A}^t} \int_{(a_t, \dots, a_{t+s}) \in A_0 \times \dots \times A_s} \pi(da_0, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)}{\int_{\phi \in \Phi} \int_{q \in Q} \int_{a_0, \dots, a_t \in \mathcal{A}^{t+1}} \pi(da_t, d\phi, dq) q(a_0, da_1) \dots q(a_{t-1}, da_t) \phi_{a_1}(d_1) \dots \phi_{a_t}(d_t)} \\ &= \Pr^\pi(a_t \in A_0, \dots, a_{t+s} \in A_s, \phi \in \Phi, q \in Q|d^t) \end{aligned}$$

□

Proof of Theorem 4.1. It is straightforward to check that the proposed allocation satisfies feasibility, budget feasibility as well as measurability assumptions. The only thing which requires an argument is that the proposed agents' plans, $(\gamma_t^{*i})_t$ maximize their utilities, given prices. We shall do it only for agent 1 (the other follows by symmetry). The proof here follows along the lines of the proof of theorem 9.2 of Stokey et al. (2004) with an adjustment for slight change in their Markov environment (our environment is technically not Markov but thanks to Lemma 1 we may treat it as if it was).

Clearly in any solution to an agent's problem the budget constraint is satisfied with the equalities, therefore wlog we may assume agent 1 is choosing only $\gamma^1 \equiv (\gamma_1^1, \gamma_2^1(d^1), \gamma_3^1(d^2) \dots) \geq 0$ (following Stokey et al. (2004) we call it a plan) to maximize:

$$u(\gamma, \gamma_0^1, \pi^1, \pi^2) \equiv \lim_{T \rightarrow \infty} u_T(\gamma, \gamma_0^1, \pi^1, \pi^2)$$

where $u_T(\gamma, \gamma_0^1, \pi^1, \pi^2) \equiv \mathbb{E}^{\pi_0^1} \sum_{t=0}^T \beta^t [p_t(\gamma_t^1 - \gamma_{t+1}^1) + \gamma_t d_t]$ taken as given $\gamma_0^1 = 0$.

Denote Γ to be the set of feasible plans for asset holdings for agent 1 (i.e. satisfying $\gamma_t^1 \geq 0$, \mathcal{F}_t^d -measurability and such that u is well defined, potentially allowing for $\pm\infty$).

Following the notation of Stokey et al. (2004) we denote:

$$V^*(\gamma_0^1, \pi_0^1, \pi_0^2) = \sup_{\gamma \in \Gamma} u(\gamma, \pi_0^1, \pi_0^2) \quad (9)$$

for each $\gamma_0^1 \geq 0$.

We will show that $V(\gamma_0^1, \pi_0^1, \pi_0^2) = V^*(\gamma_0^1, \pi_0^1, \pi_0^2)$ and that proof will imply that γ^{*1} attains the sup in (9). First we prove that

$$V(\gamma_0^1, \pi_0^1, \pi_0^2) \geq u(\gamma, \gamma_0^1, \pi_0^1, \pi_0^2) \quad (10)$$

for all $\gamma \in \Gamma$, and then we will see that

$$V(\gamma_0^1, \pi_0^1, \pi_0^2) = u(\gamma^{*1}, \gamma_0^1, \pi_0^1, \pi_0^2) \quad (11)$$

We have for any $\gamma^1 \in \Gamma$,

$$\begin{aligned} V(\gamma_0^i, \pi_0^1, \pi_0^2) &= \max_{\gamma' \geq 0} \left\{ (\gamma_0^i - \gamma') p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\pi_0^1} (V(\gamma', \lambda(d_1|\pi_0^1), \lambda(d_1|\pi_0^2)) + \gamma' d_1) \right\} \\ &\geq (\gamma_0^i - \gamma_1^1) p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\pi_0^1} (V(\gamma_1^1, \lambda(d_1|\pi_0^1), \lambda(d_1|\pi_0^2)) + \gamma_1^1 d_1) \\ &= (\gamma_0^i - \gamma_1^1) p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\pi_0^1} \left(\max_{\gamma' \geq 0} \left\{ (\gamma_1^1 - \gamma') p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\lambda(d_1|\pi_0^1)} (V(\gamma', \lambda(d_1|\pi_0^1), \lambda(d_1|\pi_0^2)) + \gamma' d_1) \right\} + \gamma_1^1 d_1 \right) \\ &= (\gamma_0^i - \gamma') p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\pi_0^1} \left(\max_{\gamma' \geq 0} \left\{ (\gamma_1^1 - \gamma') p(\pi_0^1, \pi_0^2) + \beta (V(\gamma', \lambda(d_2|\pi_0^1), \lambda(d_2|\pi_0^2)) + \gamma' d_2) \right\} + \gamma_1^1 d_1 \right) \\ &\geq (\gamma_0^i - \gamma') p(\pi_0^1, \pi_0^2) + \beta \mathbb{E}^{\pi_0^1} \left((\gamma_1^1 - \gamma_2^1) p(\pi_0^1, \pi_0^2) + \beta (V(\gamma_2^1, \lambda(d_2|\pi_0^1), \lambda(d_2|\pi_0^2)) + \gamma_2^1 d_2) + \gamma_1^1 d_1 \right) \\ &= u_1(\gamma^1, \gamma_0^1, \pi^1, \pi^2) + \beta^2 \mathbb{E}^{\pi_0^1} (V(\gamma_2^1, \pi_1(d_2|\pi_0^1), \lambda(d_2|\pi_0^2)) + \gamma_2^1 d_2) \end{aligned}$$

Here, line 4 follows from Lemma 1 and the law of iterated expectations. Note some notational complication in line 3 caused by the fact that d_1 under the second expectation is a dummy variable for that expectation and is a different d_1 then that out of that expectation. Indeed d_1 under the second expectation refers to the period 2 from the perspective of initial beliefs, but it is the first period from the perspective of updated second period beliefs — actually thanks to Lemma 1 we can replace that d_1 with d_2 in line 4.

Now we may continue this process to obtain by induction that

$$V(\gamma_0^i, \pi_0^1, \pi_0^2) \geq u_T(\gamma^1, \gamma_0^1, \pi^1, \pi^2) + \beta^T \mathbb{E}^{\pi_0^1} (V(\gamma_T^1, \lambda(d_2|\pi_0^1), \lambda(d_2|\pi_0^2)) + \gamma_T^1 d_T)$$

for all T . Now, using the assumption that $V \geq 0$ we conclude that (10) holds. Having (10) we can go over the above derivation replacing in each line γ' with the respective γ_t^{*1} (now getting the equality in each line by the construction of γ_t^{*1} which comes from the policy function for V) to obtain (11). To do so we need to use the assumed transversality condition. But this means that the plan γ^{*1} attains the maximum for agent one's problem. \square