Essays on Speculation

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Adam Wladyslaw Slawski

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Narayana Kocherlakota

August, 2009

Acknowledgments

I am grateful to Narayana Kocherlakota for his valuable advise. I also thank Chris Phelan and Jan Werner for constant help and encouragement throughout this project, Beth Allen, Nabil Al-Najjar, Marco Bassetto, Frederico Belo, Larry Blume, Kim Sau Chung, Larry Jones, Patrick Kehoe, Tim Kehoe, Christoph Mueller, David Rahman, Itai Sher, Ctirad Slavik, Radek Stefanski, Christian Stoltenberg, Tomasz Strzalecki, Aleh Tsyvinski, and all the participants of the workshops and seminars at the University of Minnesota, the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Chicago, the University of Chicago, Northwestern University, the University of British Columbia, Australian National University, Penn State University, Bocconi University, and Collegio Carlo Alberto for providing helpful feedback on various stages of the development of this work.

Separately, I would like to thank my former adviser from Warsaw School of Economics, prof. Honorata Sosnowska for making me pursue the scientific carrier.

Dedication

To my parents and to my friends in Minnesota, specially Neng.

Abstract

I examine the question whether a disagreement based speculative trade can persist in an environment in which agents are learning. The speculative trade is understood as an equilibrium outcome, in which agents who buy the asset pay higher price for it, than they would be willing to had they not had the opportunity to re-trade it in the future. The corresponding notion of speculative bubble, which is analyzed in Essay 2, is understood as an excess of the equilibrium price over the current market fundamental. Harrison and Kreps (QJE 1978) provide an example of a persistent speculative trade. The agents start with different prior beliefs, which are concentrated on disjoint sets. This prevents them from learning. The resulting permanent disagreement leads to a speculative bubble, which is independent on any observed history of data. I extend this example by adding learning. In Essay 1, I set up a general equilibrium model of trading with differential beliefs and learning. The dividend process follows a general hidden Markov process. Using recursive techniques I develop tools to compute and analyze equilibria in this environment. I also provide conditions under which a speculative trade arises. In Essay 2, I apply those techniques to analyze the dynamics of a speculative bubble in a very special case of a Markov dividend process and the prior beliefs concentrated on two transition matrices. Those matrices represent two possible theories considered by agents. Agents put positive probabilities on both theories, hence they are learning. The resulting speculative bubble arises whenever the data does not convincingly favor any of the theories used by agents. I give conditions for the data generating process to lead to persistent speculative bubble. I also show that even though the speculative bubble reappears infinitely often it also happens very rarely on a typical sample path. In fact, the average time in between the periods of high bubble is infinity.

Contents

Chapter 1

Introduction

1.1 Speculation vs. Speculative Bubbles

This dissertation provides a theoretical insight into the possibility of speculation in asset markets, driven purely by heterogeneous prior beliefs about the relevant dividend process.

The question of speculative trade is inevitably linked to so called *market bubbles*. Historically people give that name to a particular behavior of the price of a commodity or an asset, which was characterized by a significant growth of the price without any apparent fundamental reason, often followed by a collapse. Formally, a bubble is typically understood in the economic literature as the departure of an asset price from its fundamental value. What is problematic with this definition, however, is that there is no broadly accepted definition of the fundamental value. In this work I provide one definition without claiming any superiority over the others. In my opinion it naturally matches the environment of this particular model but would not work in many other environments.

The most famous and the most cited historical example was the Dutch tulip mania in 1636-37. Since then there have been many other notable examples of bubbles, with the most recent ones including the Japanese asset pricing bubble and the dot-com bubble.

All of these bubbles were called speculative bubbles, following the folk belief that they were generated by some traders buying an asset or commodity in order to profitably resell it in the future, without actually believing it has any significant intrinsic value.

In this dissertation I will be analyzing both speculative trade and the bubbles in a context of a model with risk neutral agents having heterogeneous beliefs, and facing short selling constraints. Both speculation and the bubble, or the fundamental value, will be given a precise meaning and it will be clear that both phenomena are very closely related.

By speculative trade I will understand a dynamic general equilibrium outcome, in which there is a contingency where an agent who currently holds the asset is paying a price that is higher then his subjective net present value of holding it forever. This means this particular trader buys the asset with the intention to re-sell it at some future date.

The bubble will be defined as an equilibrium outcome, in which the current price of an asset is higher than the fundamental value, which is defined as the maximum of any agent's subjective present value of holding it forever.

Here, I will be solely focused on the risk neutral environment, in which agents need not trade for any insurance purposes. With risk neutrality, when agents have homogeneous beliefs, even with short selling constraints, there cannot be any essential trade (agents are indifferent in terms of trading or not). This is the case even with heterogeneous signaling, as is implied by the "no-trade theorem," cf. Milgrom and Stokey (1982). Once we introduce heterogeneous beliefs the situation is not so clear. One could expect the heterogeneous priors would just make the most optimistic agent hold the asset forever. Obviously what might happen is that there is not a permanent optimist, but even if there is one, there is still some disagreement about the exact probabilities of various future events, not just fundamentals. Agents might be trying to exploit that disagreement to bet against each other, with the only tool to do that being to trade the asset. It seems very natural and following the popular understanding of the word *speculation* to call such a trade speculative. And my definition indeed exactly captures this type of trade in the risk neutral environment. I will discuss the issues associated with heterogeneous priors in more detail in the next subsection.

Note that in environments with risk aversion, one could easily construct examples in which the price is higher than the fundamental even with homogeneous beliefs. This is because the short-selling constraint makes the full insurance or smoothing impossible. Even though the price might be considered by agents higher than their marginal valuation, they cannot permanently reduce their asset holdings because of the binding short-selling constraints in the future. They also do not want to reduce the asset holding only in the current period because that would reduce the consumption smoothing. According to my definition however, this would also be a speculative trade, although clearly the mechanisms pushing agents to trade in this way would be different than in the case of heterogeneous belief-driven trade.

Another important remark is that the reader should not get the impression that *bubble* and *speculative trade* need to always be present together. Even in the context of the simple model of this dissertation it might actually happen that there is speculative trade but no bubble. To understand how this can happen, consider an economy with three types of agents' beliefs, say A, B, and C. First, assume that types A and B are chosen in a way that if they were the only traders there would be both speculative trade and the bubble in equilibrium. Now construct agent C's beliefs in a way such that his fundamentals are always equal to the equilibrium price resulting from the environment with only A and B. Clearly, if agent C is faced with these prices he will be exactly indifferent between trading and not trading, so towards constructing the equilibrium we might as well assume he is not trading. The existence of agent C who does not trade does not change anything in the environment faced by agents A and B. Hence, it is easy to see that the allocation from the previous equilibrium is also an equilibrium in the new one, with agent C never trading. Agents A and B are trading the asset, always paying a higher price than their fundamental value; hence, we have speculative trade, but according to my definition there is no bubble. The price is always equal to the valuation of the most optimistic trader, who, in this case, is never interested in trading. One way to fix this deficiency of my definition would be to define the market fundamental as the valuation of the most optimistic agent who ever trades in the future. I will leave aside these issues, however, and only focus on analyzing speculation through the resulting bubbles. In all the examples I will be using, this problem will not arise. The most optimistic agent will always be trading.

1.2 Disagreement and Learning

In this work to model speculative trade, I will follow the approach of Harrison and Kreps (1978). In their paper they established that heterogeneous beliefs can result in speculative trade and a bubble (called a *speculative premium* therein). The important feature of these models is that the source of belief heterogeneity is not the difference in information the agents receive but from having unexplained differences in priors.

Before I move to discuss the implication of having heterogeneous priors, I will first say a few words about the models with differentiated signals. In this class of models, as in Rational Expectation Equilibrium (REE) models, it follows from the famous *agreeing to disagree* (cf. Auman (1976)) and *no-trade* (cf. Milgrom Stokey (1982)) theorems that the price has to reflect any important private information, in particular ruling out agents' betting against their posterior beliefs, which in particular implies there cannot be any speculation. In other words and agent cannot take advantage of his private information because the rest of the market would learn that information instantly from the price.

The existence of speculative trade in the context of REE models is related to the value of information. There is a line of research that focuses on the value of information. The most influential attempt to give information a value was done by making a market game a positive sum game. Clearly if the game is of a positive sum, then the no trade theorem is not a problem anymore because the market participants do not need to use the private information against each other. They can coordinate in some fashion to exploit that information and share the surplus. Important papers in this spirit include Grossman and Stiglitz (1980), Hellwig (1980) and Kyle(1985). These papers consider a market game, in which there are noise traders, who trade for the reason external to the model. In other words, their utility function is not explicitly analyzed and they provide a stochastic supply, which is independent of the price. This gives that game a positive sum. If some of the agents have some information about the behavior of the noise traders they can use that information against them without making other strategic agents worse off.

Unsurprisingly, Grossman and Stiglitz (1982) and Hellwig (1982) obtain the result that as the noise goes to zero, the value of information goes to zero and the price becomes fully revealing. This just confirms that in the context of the models where heterogeneity is coming from differentiated signals, in order to get speculation one needs to significantly depart from the rationality of some agents.

As mentioned earlier, this dissertation is pursuing the line of Harrison and Kreps (1978), with agents having unexplained heterogeneous priors. This allows avoiding the problems related to the no-trade theorem (agents trading with each other because they know that their beliefs differ, not their signals). Using heterogeneous priors is a little bit tricky, however. Besides the obvious objection that if one is willing to assume heterogeneous priors one can explain pretty much everything, there is a more troubling point, called the market selection hypothesis. It is a conjecture which states that even if agents start with heterogeneous priors, after a sufficiently long period of trading, the agents who started with beliefs that are the farthest from the truth should either learn or be removed from the market. In other words, in well–established markets the traders should have essentially the same beliefs. As compelling as it may sound, the market selection hypothesis was never convincingly proved or disproved in the economics literature, with many papers showing examples going either way (see Sandroni (2000) or Blume and Easley (2006) for a detailed market selection hypothesis literature overview).

My research here is somehow parallel, but instead of long run survival I will be analyzing the persistence of speculative bubbles in the presence of learning. A naive guess would be that if agents are learning any essential difference in their prior beliefs should be removed, and hence, the speculative trade should eventually cease to exist. The main goal is to construct an example in which differences in subjective prior 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.

This work consists of two essays presented in the following two chapters.

- Essay 1 develops general tools to analyze dynamic equilibria with heterogeneous prior beliefs and learning in an environment with risk–neutral agents, short–selling constraints and the dividend process following a hidden Markov process. I introduce the concept of recursive equilibrium, which uses current beliefs as state variables. I give the precise definition of the speculative bubble as a difference between the price and the market fundamental, with market fundamental being the maximum across the agents of buy-and-hold forever individual valuation of the asset. Then I show that any sequential market equilibrium price must be at least of the magnitude of that of the recursive equilibrium. Since in this environment the fundamental value of the asset is not affected by the particular equilibrium outcome, then this implies that the magnitude of the speculative bubble in equilibrium is also bounded below by that of the recursive equilibrium. The main result is to prove that a recursive equilibrium generates a sequential market equilibrium (Theorem 1) and to propose a simple operator, which characterizes the recursive equilibrium price as its fixed point (operator T). This operator is shown to be a monotone contraction, which gives not only uniqueness but also provides a simple criterion for the existence of a speculative bubble in equilibrium. This essay is closed by two illustrative examples.

- Essay 2 takes advantage of the techniques developed in Essay 1 to provide in-depth analysis of a very special case of 2-state Markov dividend process. This essay focuses solely on the dynamics of the speculative bubble. The main result here is to provide conditions for the agents' beliefs and the true data generating process so that learning lasts forever and the bubble is bounded away from zero asymptotically with probability 1.

The two essays are self contained, with an exception of application of Theorem 1 in Essay 2, which is clearly marked.

The integral part of the work is two technical appendices, to which some proofs are relegated. Appendix A covers the proofs from Essay 1, while Appendix B contains those of Essay 2.

Chapter 2

Speculative Trade with Learning

2.1 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.

As discussed in the Introduction these kind of bubbles are believed to be driven by a speculative trade among agents who disagree about the future returns of the assets.

One explanation of bubbles which is in that spirit was proposed by Harrison and Kreps (1978). In their model the equilibrium price is above any agent's buy-and-hold forever valuation of the asset, hence there is a bubble according to the definition below (Harrison and Kreps call it a speculative premium rather than a bubble). This bubble is generated by a speculative trade 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 problem in this example is that it is unclear how robust it would be to adding learning. One could expect that if agents were not dogmatic any initial disagreement would eventually disappear. The difficulty with verifying this statement is that adding learning to the environment with heterogeneous beliefs makes solving for equilibrium much more difficult. Even though some work has been done in asset pricing with learning in the context of verifying the market selection hypothesis (cf. Blume and Easley (2006) Sandroni (2000), or more recently, Beker and Espino (2009)) that methodology only applies to frictionless environments.

As mentioned in the Introduction, to analyze the dynamics patterns of speculative trade I will be focusing on the behavior of speculative bubbles. They are defined as the excess of the market price of the asset over its fundamental value.

This essay develops tools to analyze the equilibrium prices and bubbles in environments with short selling constraints and learning, in which the dividend process follows a hidden Markov process.

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 adding an 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.

An important supporting result is the one provided by Proposition 4, which states that recursive equilibrium is not just any sequential equilibrium, but it actually is a one with the lowest possible prices. This result is important in the context of analyzing speculation, namely if we have speculation in recursive equilibrium, then it must be the case in any other equilibrium.

In the end of this essay I propose two illustrative examples. In one example I prove that if there are only two hidden states then there is no speculation in any recursive equilibrium, no matter what initial beliefs are. In another setup with three hidden states I show that, for some particular initial belief, the equilibrium price has to feature a bubble.

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.

The rest of this essay is organized as follows. 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 a speculative bubble.

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.

I will close this introduction with a brief overview of the models of speculative trade, which is driven by heterogeneous priors. This trend in the literature was started by the paper by Harrison and Kreps (1978). This paper I discussed in detail above and also throughout the essay 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 the fundamentals 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 this essay 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.

2.2 The Model

2.2.1 Economy

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

There is one 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. Each agent starts with some initial holding of the asset, $\bar{\gamma}_0^i$, such that $\bar{\gamma}_0^1 + \bar{\gamma}_0^2 = 1$.

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 $(\Omega, \mathcal{F})$ 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, $(a_t)_{t=0,1,\ldots}$ be a stationary Markov process with the transition function $q : \mathcal{A} \to \Delta(\mathcal{A})$ assumed Borel-measurable, with $\Delta(A)$ 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 $\Phi_\cdot : \mathcal{A} \to \Delta(\mathcal{D})$ 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 discreet or Lebesgue). Denote this density by ϕ_a.

Having specified $\bar{a}_0 \in \mathcal{A}$, the family $\Phi \in \mathcal{B}(\mathcal{A}, \{\zeta \in \Delta(\mathcal{D}) | \zeta << \mu\})$, 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 $(\Omega, \mathcal{F})$ 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 $A_0 \in \mathcal{B}(\mathcal{A}), \ldots, A_t \in \mathcal{B}(\mathcal{A})$, $D_1 \in \mathcal{B}(\mathcal{D}), \ldots, D_t \in \mathcal{B}(\mathcal{D})$ we have:

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

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 $\Phi_\cdot$ or the transition function q. We will assume that $\Phi_\cdot$ and q can take values in some Borel sets of admissible values, $\mathbf{\Phi} \subseteq \mathcal{B}(\mathcal{A}, \{\zeta \in \Delta(\mathcal{D}) | \zeta << \mu\})$, and $\mathcal{Q} \subseteq \mathcal{B}(\mathcal{A}, \Delta(\mathcal{A}))$, respectively. Hence the agents formulate beliefs about the value of $(a_0, \Phi_\cdot, q)$.

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

$$\Pr^\pi(F) = \int_{\mathcal{A} \times \mathbf{\Phi} \times \mathcal{Q}} \Pr^{\bar{a}_0, \Phi, q}(F) \pi(d\bar{a}_0, d\Phi, dq) \tag{2.1}$$

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 $\pi_0^i \in \Delta(\mathcal{A} \times \mathbf{\Phi} \times \mathcal{Q})$.

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:

$$\begin{aligned} &\max E^{\pi_0^i} \sum \beta^t c_t^i \\ &\text{s.t. } c_t^i + p_t \gamma_{t+1}^i \le p_t \gamma_t^i + \gamma_t^i d_t \\ &\text{s.t. } c_t^i, \gamma_{t+1}^i \text{ — } \mathcal{F}_t^d\text{-measurable} \\ &\text{s.t. } \gamma_0^i = \bar{\gamma}_0^i, \gamma_t^i \ge 0 \end{aligned} \tag{2.2}$$

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

It is worth noting at this point, that it is an 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. In order to consider this equilibrium as a rational expectations equilibrium, the whole hierarchy of beliefs needs to be specified. Specifically, we can assume the common knowledge of beliefs.

Another technical issue associated with the definition above is that there might be a set of future contingencies, which is believed by an agent to be of zero probability, on which his behavior is inconsistent with Bayesian learning. In this paper I will be always assuming to use the version of conditional probability which is consistent with Bayes' rule, so any resulting equilibria will not have this problem.

2.2.2 Speculative trade

In this setup, it is the most natural way to define speculative trade as the situation in which the equilibrium price exceeds the fundamental valuation of the asset for the agent is an actual holder. 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.

It is a natual definition because agents are risk neutral and are sharing the discount factor, hence agents do not need the asset to smooth consumption or for insurance. If they decide to purchase the asset for the price which is higher than their subjective fundamental value it is because they use it as a betting device against the market which they perceive as not pricing the asset properly in the future. They understand the market does not price the asset properly because there are some other traders with fundamentally wrong beliefs.

This definition however is not the most practical one because to check if there is a speculative trade one would need to analyze the behavior of all the traders. A simpler statistic which can measure the strength of speculative trade is a speculative bubble.

Definition. The fundamental value of the asset for agent i at time t, given history $d^t \equiv (d_1, \ldots, d_t)$ is:

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

Definition. Speculative bubble is a current excess of the equilibrium prices over the market fundamental value,

$$s(d^t) \equiv p(d^t) - \max_i V_t^i(d^t)$$

Note that whenever $s > 0$ we do have a speculative trade, but not necessarily the other way round. Also note that the magnitude of the bubble is measuring how much the agents value the opportunity to bet against each other provided by the asset, and not necessary the volume of the speculative trade. For the remaining of this work I will be focusing on the speculative trade which is associated with a positive bubble, and i will often abuse the terminology by using speculative trade and a bubble interchangeably.

Example. Now let us see how the example of Harrison and Kreps fits into this notation.

In their model 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 $\mathcal{A} = \mathcal{D} = \{0, 1\}$, $\mathbf{\Phi} = \{\Phi\}$ 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 $\mathcal{A} \times Q$, 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:

$$V^1 = \begin{bmatrix} 4/3 \\ 11/9 \end{bmatrix} = \begin{bmatrix} 1.33 \\ 1.22 \end{bmatrix} \qquad V^2 = \begin{bmatrix} 16/11 \\ 21/11 \end{bmatrix} = \begin{bmatrix} 1.45 \\ 1.91 \end{bmatrix} \tag{2.3}$$

2.3 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 belief update operator, λ, which is crucial for applying the stochastic dynamic programming techniques. It maps the previous period beliefs into the current one, taking into account the current realization of the dividend d.

Definition 1. Given $\pi \in \Delta\big(\mathcal{A} \times \mathbf{\Phi} \times \mathcal{Q}\big)$, and $d \in \mathcal{D}$ define a measure $\lambda^d(\pi) \in \Delta\big(\mathcal{A} \times \mathbf{\Phi} \times \mathcal{Q}\big)$, by

$$\lambda^d(\pi)(A \times \Phi \times Q) \equiv \Pr^\pi(a_1 \in A \wedge \phi \in \Phi \wedge q \in Q | d_1 = d)$$

for each measurable $A \subseteq \mathcal{A}, \Phi \subseteq \mathbf{\Phi}, Q \subseteq \mathcal{Q}$, and some particular version of the conditional probability. Without loss of generality, throughout this paper I will be always using the version of the conditional probability which is given by Bayes' rule.

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

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

$$\Gamma(\gamma,\pi^1,\pi^2) \in \underset{\gamma'\geq 0}{argmax}\left\{(\gamma-\gamma')p(\pi^1,\pi^2)+\beta E^{\pi^1}\left(V(\gamma',\lambda^{d_1}(\pi^1),\lambda^{d_1}(\pi^2))+\gamma' d_1\right)\right\}$$

and for each γ,π^1,π^2,d we have:

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

It should be noted that in this case the symmetry reflects the fact that only one value function is used for both agents. When it is used for agent 1 it reads $V(\gamma,\pi^1,\pi^2)$ and when for agent 2 it becomes $V(\gamma,\pi^2,\pi^1)$. Also note that whenever d_1 appears under the expectation associated with measure π, (e.i. $E^{\pi}f(d_1)$ with f being any measurable real function) it refers to the first period dividend distributed according to the probability distribution Pr^{π} defined in (2.1).

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

$$\pi_1^i \equiv \lambda^{d_1}(\pi_0^i)$$
$$\pi_{t+1}^i \equiv \lambda^{d_{t+1}}(\pi_t^i)$$

Theorem 1. *Given some initial beliefs,* $\pi_0^1,\pi_0^2\in\Delta\big(\mathcal{A}\times\mathbf{\Phi}\times\mathcal{Q}\big)^2$*, if* $V:[0,1]\times\Delta\big(\mathcal{A}\times\mathbf{\Phi}\times\mathcal{Q}\big)^2\to\mathbb{R}_+$*,* Γ*,* $p:\Delta\big(\mathcal{A}\times\mathbf{\Phi}\times\mathcal{Q}\big)^2\to\mathbb{R}$ *constitute a recursive equilibrium then the processes:* $p_t^*\equiv p(\pi_t^1,\pi_t^2)$*,* $\gamma_{t+1}^{*i}\equiv\Gamma(\pi_t^i,\pi_t^{-i})$*,* $c_t^{*i}\equiv p_t(\gamma_t^{i*}-\gamma_{t+1}^{i*})+\gamma_t^{i*}d_t$ *constitute a sequential market equilibrium, given beliefs* π_0^1,π_0^2 *if the following transversality condition holds:*

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

for $i=1,2$*.*

The proof will follow from the following lemma:

Lemma 1. *For each* $s\in\mathbb{N}$*,* $d\in\mathcal{D}$*,* $A_0,\ldots,A_s\subseteq\mathcal{A}$*,* $\Phi\subseteq\mathbf{\Phi}$*,* $Q\subseteq\mathcal{Q}$*, such that* $A_0,\ldots,A_s,\Phi,Q$ *are measurable subsets, we have almost surely:*

$$\Pr^{\pi}(a_1\in A_0\wedge\ldots\wedge a_{s+1}\in A_s\wedge\phi\in\Phi\wedge q\in Q|d_1)$$
$$=\Pr^{\lambda^{d_1}(\pi)}(a_0\in A_0\wedge\ldots\wedge a_s\in A_s\wedge\phi\in\Phi\wedge q\in Q)$$

This lemma states that all the future distribution of the relevant processes at the next period is, from the perspective of player i, completely described by the updated measure $\pi_1^i = \lambda^{d_1}(\pi_0^i)$. This argument extends by induction to any future period: all the information about agent i's subjective future distributions of all the processes is completely encoded in π_t^i. This legitimates the introduction of the recursive equilibrium in this environment.

The proof of Theorem 1 as well as that of Lemma 1 are relegated to the appendix as they are purely technical.

2.3.1 Characterization of recursive pricing rule

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

$$p(\pi^1, \pi^2) = \max_{i=1,2} \beta E^{\pi_i}(p(\lambda^{d_1}(\pi^1), \lambda^{d_1}(\pi^2)) + d_1)$$

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

$$T^{(i)}p(\pi^1, \pi^2) \equiv \beta E^{\pi_i}\left\{d_1 + p\big(\lambda^{d_1}(\pi^1), \lambda^{d_1}(\pi^2)\big)\right\} \qquad i = 1, 2$$

$$Tp \equiv \max_{i=1,2} T^{(i)}p$$

With this notation, the equation for prices becomes:

$$p = Tp \tag{2.4}$$

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. 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):

$$\begin{aligned} T^{(i)}p(d) &= \beta E^{\pi_i}\left\{d_1 + p\big(d_1\big)\right\} \\ &= \begin{cases} \beta\left\{(1/2)p(0) + (1/2)(1+p(1))\right\} & \text{for } d=0, i=1 \\ \beta\left\{(2/3)p(0) + (1/3)(1+p(1))\right\} & \text{for } d=0, i=2 \\ \beta\left\{(2/3)p(0) + (1/3)(1+p(1))\right\} & \text{for } d=1, i=1 \\ \beta\left\{(1/4)p(0) + (3/4)(1+p(1))\right\} & \text{for } d=1, i=2 \end{cases} \end{aligned}$$

This allows us for almost immediate guess the solution to (2.4), 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 (2.3).

Let us investigate some properties of T.

Lemma 2. *Operators T,$T^{(i)}$ are all β-contractions w.r.t. the* sup*-norm.*

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

Lemma 3. *If p is continuous, then Tp is continuous.*

As a corollary we get:

Theorem 2. *If $\beta < 1$ then there is the unique bounded solution, p^* to the functional equation (2.4). Moreover, p^* is continuous and $p^* = \lim_{t\to\infty} T^t p$ (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. *If $Tp \geq p$ for some price system p, then $T^2p \geq Tp$. Hence also $p^* \geq p$*

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

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

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

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

□

As a corollary we get a useful fact.

Theorem 3. *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 $p^*(\pi^1, \pi^2) \geq Tp^F(\pi^1, \pi^2) > p^F(\pi^1, \pi^2)$ □

This theorem gives us an easy tool to check if we have a speculative bubble 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 (2.4) i.e. $p^* = p^f$ and the bubble is always zero. 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 $p^*(\pi^1, \pi^2) \geq Tp^f(\pi^1, \pi^2) > p^f0(\pi^1, \pi^2)$, which means that we have a positive bubble, and hence a speculative trade, in equilibrium.

Now, I will show that the price in the bounded recursive equilibrium provides a lower bound for the set of all sequential market equilibria. This justifies the use of applying recursive equilibrium in the contest of speculative trade. If there is a speculative bubble in the recursive equilibrium it is also positive in any sequential market equilibrium.

Proposition 4. *For any initial beliefs* π_0, *if* p *is a sequential equilibrium price system, and* p^* *is the bounded recursive equilibrium price, then for almost every history,* d^t, *we have* $p(d^t) \geq p^*(\pi_t(d^t))$.

Proof. If $(p_t(d^t))$ is an equilibrium price system, then it has to satisfy the first order conditions of (2.2), which, taking into account the fact that the market clearing condition must hold (i.e., $\gamma_t^i > 0$ for at least one agent), leads to:

$$p_t(d^t) = \max_{i=1,2} \beta E^{\pi_0^i}\big[p_{t+1}(d^{t+1}) + d_{t+1}|d^t\big]. \tag{2.5}$$

Define inductively a sequence of functions, $p_n : \Delta\big(\mathcal{A} \times \mathbf{\Phi} \times \mathcal{Q}\big)^2 \to \mathbb{R}$, by

$$\begin{aligned} &p_0^*(\pi) \equiv 0 \\ &p_{n+1}^*(\pi) \equiv \max_{i=1,2} \beta E^{\pi}\big[d_1 + p_n^*\big(\lambda^{d_1}(\pi)\big)\big]. \end{aligned}$$

Then by Proposition 2.7, we have $p^* = \lim_{n\to\infty} p_n^*$. I will show by induction that $p_t(d_t) \geq p_n^*(d_t, \pi_t^1(d_t), \pi_t^2(d^t))$, for each n, d^t.

For $n = 0$ this is obvious. Suppose that $p_t(d^t) \geq p_n^*(d^t, \pi_t(d^t))$ for some n and all t and d^t. Then, using (2.5) and Lemma 1 (recursively, t times) we get

$$\begin{aligned} p_t(d^t) &= \max_{i=1,2} \beta E^{\pi_0^i}\big[d_{t+1} + p_{t+1}\big(d^{t+1}\big)|d^t\big] \\ &\geq \max_{i=1,2} \beta E^{\pi_0^i}\big[d_{t+1} + p_n^*\big(\pi_{t+1}(d^{t+1})\big)|d^t\big] \\ &= \max_{i=1,2} \beta E^{\pi_t^i(d^t)}\big[d_1 + p_n^*\big(\lambda^{d_1}(\pi_t(d^t))\big)\big] \\ &= p_{n+1}^*(\pi_t(d^t)). \end{aligned}$$

□

2.4 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 then moving out of it. In that case I prove that no speculative trade (and hance no bubble) 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 show how Theorem 3 can be easily applied to show that there is a generic speculative pattern (for some open set of parameter values). In that example we will also see that the speculation may persist arbitrarily long even though the agents are learning.

2.4.1 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 consider the class of models in which: $\mathcal{A} = \{a_1, \ldots, a_n\}$, $\mathcal{D} = \{0, 1\}$, $\mathbf{\Phi} = \{\phi\}$ with $\phi_{a_j}(1) = 1 - \phi_{a_j}(0) \equiv \phi_j$, $\mathcal{Q} = \{q\}$, with $q = \begin{bmatrix} q_{11} & \cdots & q_{1n} \\ \vdots & \ddots & \vdots \\ q_{n1} & \cdots & q_{nn} \end{bmatrix}$ (Markov chain transition matrix).

In words this setup means that there is an underlying $\mathcal{A}$-valued Markov chain in the economy, $a_0, a_1, \ldots, a_t, \ldots$ 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 $\pi^i = [\pi^i_1, \ldots, \pi^i_n] \in \Delta(\mathcal{A}) \simeq \Delta^{n-1}$.

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 $V = [V_1, \ldots, V_n]'$ 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):

$$\lambda(d'|\pi^i) = \begin{cases} \left[\frac{\phi_1 \sum_{j=1}^n q_{j1}\pi_j^i}{\sum_{j'=1}^n \phi_{j'} \sum_{j=1}^n q_{jj'}\pi_j^i}, \dots, \frac{\phi_n \sum_{j=1}^n q_{jn}\pi_j^i}{\sum_{j'=1}^n \phi_{j'} \sum_{j=1}^n q_{jj'}\pi_j^i}\right] & \text{if } d' = 1 \\ \left[\frac{(1-\phi_1) \sum_{j=1}^n q_{j1}\pi_j^i}{\sum_{j'=1}^n (1-\phi_{j'}) \sum_{j=1}^n q_{jj'}\pi_j^i}, \dots, \frac{(1-\phi_n) \sum_{j=1}^n q_{jn}\pi_j^i}{\sum_{j'=1}^n (1-\phi_{j'}) \sum_{j=1}^n q_{jj'}\pi_j^i}\right] & \text{if } d' = 0 \end{cases} \tag{2.6}$$

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:

$$\Pr^{\pi^i}\{d' = 1\} = \pi^i q\phi \tag{2.7}$$

The formula for operator T becomes:

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

2.4.2 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

$$p^f(\pi^1, \pi^2) = \pi^1 V_h + (1 - \pi^1)V_l \tag{2.8}$$

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 occured). By (2.6) we have:

$$\lambda^d(\pi) = \begin{cases} \frac{\phi_h\left[(1-\epsilon^1)\pi + \epsilon^2(1-\pi)\right]}{\phi_h[(1-\epsilon^1)\pi+\epsilon^2(1-\pi)]+\phi_l[\epsilon^1\pi+(1-\epsilon^2)(1-\pi)]} & \text{for } d = 1 \\ \frac{\phi_l\left[(1-\epsilon^1)\pi + \epsilon^2(1-\pi)\right]}{\phi_l[(1-\epsilon^1)\pi+\epsilon^2(1-\pi)]+\phi_h[\epsilon^1\pi+(1-\epsilon^2)(1-\pi)]} & \text{for } d = 0 \end{cases} \tag{2.9}$$

A bit of algebra gives:

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

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 (2.7). Also I will need $\lambda^1(\pi) > \lambda^0(\pi)$, which is intuitively obvious and straightforward to check from (2.9).

This gives us:

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

where the second line comes from (2.8), 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. $\square$

2.4.3 Example with speculation

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

$$
q = \begin{bmatrix} q_{h_1h_1} & q_{h_1h_2} & q_{h_1l} \\ q_{h_2h_1} & q_{h_2h_2} & q_{h_2l} \\ q_{lh_1} & q_{lh_2} & q_{ll} \end{bmatrix} = \begin{bmatrix} 1-\epsilon^1 & 0 & \epsilon^1 \\ 0 & 1-\epsilon^2 & \epsilon^2 \\ 0 & 0 & 1 \end{bmatrix}
$$

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. *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 differers 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.

2.5 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.

Chapter 3

The Dynamics of a Speculative Bubble under Learning

3.1 Introduction

In many financial markets we observe periods of price behavior referred to as speculative bubbles. This term usually means that the price of an asset is significantly different from what is believed to be the asset's fundamental value. Furthermore, these events are often not isolated incidents involving a single asset, but rather entire regions or sectors of the economy. We observed this during the Japanese asset price bubble in the late 80's and during the "dot-com" bubble in the US in the late 90's. Other notable examples include the recent behavior of crude oil or real estate prices.

Harrison and Kreps (1972) provide one way to rationalize these events. In their paper, risk-neutral investors have heterogeneous prior beliefs about the persistence of the dividend process and face short-selling constraints. In some periods, the asset is not held by the agents for whom the expected discounted stream of dividends is higher, but by those who expect a higher dividend only in the following period and are hence hoping for profitable resale. This generates speculative trade and a price bubble, where the asset serves not only as a generator of dividends, but also as a period-to-period betting device. The bubble they obtain is not time-varying, in the sense that it depends on the history of dividends only through the current realization. This occurs because agents do not learn from the data. It seems that allowing for learning in this setting would quickly kill the bubble by removing heterogeneity from agents' posteriors. The goal of this paper is to address this objection by proposing a model in which rational traders fail to agree, despite learning from a common signal, resulting in a perpetual speculative bubble.

I develop a general equilibrium model of speculative bubbles in which agents differ in their prior beliefs. This model is based on Harrison and Kreps (1978), with an important addition, that I allow for learning. There is a market for a single asset that pays a stochastic dividend. The dividend follows a two-state Markov process, and can be either zero or one each period. There are two types of risk neutral agents, who face a short sale constraint, and both start with different prior beliefs concerning the transition matrix of the dividend process. The beliefs of both types are concentrated on two available matrices. This can be interpreted as a situation in which there are two theories or models being used by market practitioners with matrices representing the theories.

A bubble is defined as that portion of the equilibrium price over and above the market fundamental. The market fundamental is the maximum buy-and-hold-forever valuation of the asset. I define the fundamental value process independently of a given equilibrium. This statistic reflects how much agents value the asset in the absence of a market. Considering the fundamental value separately from the market can be justified by the assumed risk neutrality of the agents. This way the agents do not need the market for insurance purposes. The only reason for trading is speculation. This model naturally extends to utility functions with risk aversion, but in that case the fundamental value needs to account for the insurance motive.

I use a recursive methodology borrowed from Slawski (2008) to show that if agents are sufficiently impatient, there exists a linear lower bound on a current bubble magnitude as a function of current belief differentials. This result means that as long as belief differentials stay bounded away from zero so does the bubble, and applies regardless of the true data-generating process of the dividend. The main theorem provides non-generic conditions for the true data-generating process so that the differences in the posteriors persist forever on almost all sample paths. Under these conditions a bubble of high magnitude appears infinitely often with probability one. I provide the following generic complement to the main theorem: the data-generating process can be chosen so that, in expected terms, the belief difference and the high magnitude bubble can persist arbitrarily long.

For these results to apply, it is necessary that agents' beliefs put zero mass on the true data-generating process. Here, I take as a given that even though market traders might be testing for all existing models, it is never the case that one of the available models coincides exactly with reality.

To get an idea of how the lack of convergence in the posteriors is obtained, consider the following thought experiment: there are two agents who observe a coin being flipped. The agents know the experiment produces an iid binomially distributed sequence of outcomes ($d = (d_1, d_2, \ldots)$) which are in the set {heads,tails}. However, the agents do not know the probability of heads, $\theta \in [0, 1]$. They assume some prior beliefs about θ, which are concentrated on the set $\{0.4, 0.6\}$. Using Bayes' rule we get that the posteriors are equal to the priors each time the number of heads equals the number of tails in a current history. For each history, d^t, let $h_t(d^t)$ and $\tau_t(d^t)$ denote the number of heads and the number of tails, respectively. If the true data-generating process has $\theta = 0.5$, then it is clear that the process $h - \tau$ is a standard symmetric random walk and is thus persistent (i.e. $\Pr(\exists t > 0 \text{ s.th. } d_t = \tau_t) = 1$). Therefore agents' posterior beliefs fail to converge to each other on almost all sample paths.

One can see that the lack of convergence is achieved only for the true data-generating process with $\theta = 0.5$, and it is therefore not generic. However, using the weak law of large numbers we get the following robustness result: as θ approaches 0.5, the expected time that the process $h - \tau$ spends around the origin tends to infinity. This means that the initial disagreement can last arbitrarily long if θ is sufficiently close to 0.5. The generic result in this paper follows the above logic.

This example illustrates the way the model should be interpreted. The fact that the true probability lies in-between the two existing models reflects the view that people's knowledge is not complete. Models simplify reality so much that none of them can be exactly right. This paper highlights that making such simplifying assumptions can potentially lead to perpetual disagreement, which is reflected in the existence of speculative bubbles. The primary reason that I cannot directly use this example is that this sort of disagreement about an iid process does not lead to speculation. As such, I must extend this technique to the Markov environment in which bubbles can emerge.

My paper naturally divides into two parts: Section 2 develops a general equilibrium model of speculative trading and provides a numerical link between the current difference between agents' posteriors and the size of the speculative bubble, while Section 3 focuses on showing how to generate the data so that agents's posteriors fail to converge, in the spirit of the coin flip example presented earlier. This section also delivers the main result.

I close the introduction with a brief literature review. I focus on those models in which bubbles are driven by heterogeneous beliefs and a short sales constraint. This approach was first taken by Harrison and Kreps (1978). The main objection to this model is that agents concentrate their priors on separate matrices, so they do not learn over time and the speculative bubble stays on the same level for every history. Morris (1996) extends the example of Harrison and Kreps by accounting for learning. He uses an iid dividend process and obtains a speculative bubble that vanishes over time at an exponential rate. Allen, Morris, and Postlewaite (1993) model bubbles by allowing for heterogeneous priors as well as private signals. Furthermore, there is a finite horizon, with an asset paying the dividend only once after the sequence of trading. This setup enables them to get a very strong notion of a bubble in which it is common knowledge that at some point the asset price is higher than the final dividend. The most recent contribution to this literature is the paper by Scheinkman and Xiong (2003). Their model can be treated as a fully developed continuous time version of Harrison and Kreps, in which (cumulative) dividend follows a diffusion process. Agents are getting two noisy signals and wrongfully interpreting one of them (each agent a different one) as carrying some information about the innovation part of the dividend process. In reality, neither of these processes is correlated with the innovation part. The dynamics of speculative trade are such that whenever the signal of an agent becomes more optimistic, the agent buy the asset. Even though this model predicts very interesting dynamics (high trading volumes and price volatility during bubbles), it still does not take into account the process of learning.

It is important to note that in the above models — including mine — having an infinite time horizon is not crucial to generate a bubble. The infinite time horizon is only used for analytical simplicity.

3.2 Speculative Equilibria

The main goal of this section is to provide a direct link between the current posteriors' asymmetry and the magnitude of the equilibrium bubble. I start with a detailed description of the market setup and the definition of sequential market equilibria. In Subsection 2.2, I introduce the notion of recursive equilibrium, which is a crucial tool in analyzing equilibria. I show that a recursive equilibrium generates a sequential market equilibrium with the lowest possible prices. I also introduce a functional operator, T, which characterizes the bounded recursive equilibria as its fixed point and serves as the main tool for analyzing them. Subsection 2.3 contains an illustrative benchmark case of no-learning, which can be explicitly solved. Subsection 2.4 contains the main result of this section, which is the linear lower bound for the speculative bubble.

3.2.1 Sequential Market Equilibria and Bubbles

I consider a market for a single stock with trading taking place at discrete time periods ($t = 0, 1, 2, \ldots$). The stock is in unit supply and the dividends are paid to the current holder at each period immediately prior to trading. Agents face a short sale constraint. I denote the random dividend process by $d = (d_0, d_1, d_2, \ldots)$. The dividends take values in the set $D \equiv \{0, 1\}$, and are observed up to the current period by all traders. The current owner has no control over the operation of the firm and the dividend process is the only (exogenous) source of uncertainty. The dividend process follows a two-state Markov chain with the transition matrix $Q = (q_{ij})_{0 \leq i,j \leq 1}$, which is not known to the traders.

There are two classes of investors, $i \in \{1, 2\}$. Each class is of measure one and is characterized by the initial beliefs about Q, $\pi_0^i \in \Delta\left(\{Q | Q \text{ is } 2 \times 2 \text{ probabilistic matrix}\}\right)$. Each π_0^i determines the probability distribution over the realizations of process d. This probability distribution is denoted by $\Pr^{\pi_0^i} \in \Delta(D^\infty)$. The expected value operator with respect to that measure is denoted by $E^{\pi_0^i}$.

For the purpose of this paper, I consider beliefs of a particular form. Specifically, I assume that the beliefs of both agents are concentrated on two transition matrices, $Q^1 = \left(q^1_{dd'}\right)_{0 \leq d,d' \leq 1}$ and $Q^2 = \left(q^2_{dd'}\right)_{0 \leq d,d' \leq 1}$ (i.e., $\pi_0^i \in \Delta\left(\{Q^1, Q^2\}\right)$). Belief π_0^i can be represented by one number, say $\pi_0^i(\{Q^1\}) \in [0, 1]$. In what follows I abuse the notation by writing π_0^i instead of $\pi_0^i(\{Q^1\})$.

One can think of Q^1 and Q^2 as the conceivable theories or models of the market. If both agents put positive probabilities on both, Q^1 and Q^2, then they have common support and I will refer to this situation as "learning." Agents may differ in the initial probabilities they assign to different theories, but they agree on what they consider possible.

The instances in which agents' beliefs are concentrated on single but different matrices (i.e., $\pi_0^i = \delta_{Q^i}$) are equivalent to the example of Harrison and Kreps, and I will refer to this situation as "no learning." In this case, none of the agents consider the theory used by the other type as possible, and therefore they do not update.

Both types' consumption sets are $C = \ell^\infty(\bigcup_t D^t) \equiv \bigcup_{\bar{c}}\{(c_t^i)_{t=0}^\infty | c_t^i : D^t \to [-\bar{c}, \bar{c}]\}$. They are risk neutral with utility functions $U^i : C \to \mathbb{R}$, given by

$$U^i(c^i) = E^{\pi_0^i} \sum_{t=0}^{\infty} \beta^t c_t^i,$$

for each $c^i \in C$. The assumed boundedness of the set C ensures the limit always exists.

A sequential market equilibrium consists of the following components:

$$\begin{aligned}
&c \in C^2 \\
&\gamma = ((\gamma_t^i)_{i=1}^2)_{t=1}^\infty, \gamma_t^i : D^{t-1} \to \mathbb{R}_+ \\
&p = (p_t)_{t=0}^\infty, p_t : D^t \to \mathbb{R}.
\end{aligned}$$

The first component is the consumption allocation, the second one is the allocation of the asset across agents, and the third one is the contingent price.

A triplet (c, γ, p) constitutes a sequential market equilibrium if for each $i = 1, 2$, c^i, and γ^i solves:

$$\begin{aligned}
&\max_{(c^i,\gamma^i)} U(c^i) \qquad (3.1) \\
&\text{s.t. } c_t^i(d^t) + p_t(d^t)\gamma_{t+1}^i(d^t) \leq p_t(d^t)\gamma_t^i(d^{t-1}) + \gamma_t^i(d^{t-1})d_t \qquad \text{for each } t \text{ and } d^t \\
&\text{s.t. } \gamma_0^i = \bar{\gamma}_0^i, \gamma_{t+1}^i(d^t) \geq 0
\end{aligned}$$

and the market-clearing conditions hold:

$$\begin{aligned}
c_t^1(d^t) + c_t^2(d^t) &= d_t, \text{ for each } t \text{ and } d^t \\
\gamma_t^1(d^t) + \gamma_t^2(d^t) &= 1, \text{ for each } t \text{ and } d^t.
\end{aligned}$$

I will now define the speculative bubble process for a given equilibrium. As mentioned before, the bubble process will relate the equilibrium price to the market fundamental. For each $i = 1, 2$ and for any t and d^t, define the process v^i by $v_t^i(d^t) \equiv E^{\pi_0^i}\left\{\sum_{\tau>t} \beta^t d_t | d^t\right\}$. This will be called the fundamental value of the asset for agent i, at time t, given the history d^t. It is the highest price agent i would be willing to pay for the asset if no future re-trade was possible. The market fundamental value of the asset at time t, given the history d^t, is the process v, defined by $v_t(d^t) \equiv \max_{i=1,2} v_t^i(d^t)$. This gives the market price of the asset if the current period was the last day of trade. Now, for a given equilibrium price process p, I define the speculative bubble process s as $s_t \equiv p_t - v_t$. This statistic measures how much agents value the possibility of future asset trading, purely as a betting device.

3.2.2 Recursive Equilibrium

In this subsection I will introduce an auxiliary concept of recursive equilibrium. The concept, which is a helpful tool in solving for equilibria, will also be used to see directly how the current difference in posterior beliefs translates into the size of the speculative bubble. Having the link between posteriors and the bubble is possible because in the stationary equilibrium, the price is a function of the current posterior belief profile.

To simplify the notation for the law-of-motion of the posterior beliefs, I need an additional piece of notation. Each history of length 2, $(d, d') \in D^2$, defines the function $\lambda^{dd'} : [0, 1] \to [0, 1]$ by

$$\lambda^{dd'}(\pi^i) \equiv \frac{\pi^i q^1_{d_{t-1}d'}}{\pi^i q^1_{d_{t-1}d'} + (1 - \pi^i) q^2_{dd'}}. \tag{3.2}$$

The above takes as an argument the previous beliefs and the type of transition that occurred in the last period, (d, d'). It gives as its value the new beliefs, updated using Bayes' rule. For notational convenience I denote the whole belief profile as $\pi = (\pi^1, \pi^2)$, and in what follows I will abuse notation by writing $\lambda^{dd'}(\pi)$, rather than $\left(\lambda^{dd'}(\pi^1), \lambda^{dd'}(\pi^2)\right)$.

A recursive equilibrium consists of:

$$W : \mathbb{R}_+ \times D \times [0,1]^2 \to \mathbb{R}^2$$
$$\gamma^* : \mathbb{R}_+ \times D \times [0,1]^2 \to \mathbb{R}^2_+$$
$$p^* : D \times [0,1]^2 \to \mathbb{R}.$$

The first component represents agents' value functions, and the second one their decision functions specifying agents' new asset holdings — as functions of the current asset holdings, the current dividend, and the current belief profile. The last component represents the price, which is assumed not to depend on the current asset holdings but only on the state variables.

A triplet (W, γ^*, p^*) is a recursive equilibrium if for each $i = 1, 2$, $\gamma = (\gamma^1, \gamma^2) \in \mathbb{R}^2_+$, $d \in D$, and $\pi \in [0,1]^2$, the following three conditions are satisfied.

1. Agent i's Bellman Equation holds

$$W^i(\gamma^i, d, \pi) = \max_{\gamma'^i \geq 0} \Big\{ \big[(\gamma^i - \gamma'^i) p^*(d, \pi) + \gamma^i d\big] + \sum_{d'} \big[\pi^i q^1_{dd'} + (1 - \pi^i) q^2_{dd'}\big] W\big(\gamma'^i, d', \lambda^{dd'}(\pi)\big) \Big\}. \tag{3.3}$$

2. $\gamma^{*i}(\gamma^i, d, \pi)$ is a solution to agent i's problem above.

3. Market clearing for the asset holds: if $\gamma^1 + \gamma^2 = 1$, then

$$\gamma^{*1}(\gamma^1, d, \pi) + \gamma^{*2}(\gamma^2, d, \pi) = 1.$$

The definitions of the fundamental values and the speculative bubble naturally extend to the recursive setup. The fundamental value of the asset for agent i, with beliefs π^i in state d, is $V^i(d, \pi^i) \equiv E^{\pi^i} \sum \beta^t d_t$. The market fundamental value is $V(d, \pi^1, \pi^2) \equiv \max_{i=1,2} V^i(d, \pi^i)$, and for a given recursive equilibrium price p^*, the speculative bubble is $S(d, \pi) \equiv p^*(d, \pi) - V(d, \pi)$.

Before introducing the tools to solve for a recursive equilibrium, I will provide a link between recursive equilibria and sequential market equilibria. If I specify the initial belief profile $\pi_0 = (\pi^1_0, \pi^2_0) \in [0,1]^2$, then any recursive equilibrium induces a sequential form price and allocation in a natural way. Namely, given a recursive equilibrium

(W, γ^*, p^*), define a sequential allocation and price (c, γ, p) together with the implied process of the posterior belief profiles π, recursively by:

$$\begin{aligned}
\pi_{t+1}(d^{t+1}) &= \lambda^{d_t d_{t+1}}(\pi_t(d^t)) \\
p_t(d^t) &= p^*(d_t, \pi_t(d^t)) \\
\gamma_{t+2}(d^{t+1}) &= \gamma^*(\gamma_{t+1}(d^t), d_{t+1}, \pi_{t+1}) \\
c_t(d^t) &= (\gamma_t - \gamma_{t+1})p_t + \gamma_t d_t.
\end{aligned}$$

In order to use recursive equilibrium techniques, we need to know that these sequences constitute a sequential market equilibrium. From Slawski (2008), we have the following:

Proposition 6. *Assume (W, γ^*, p^*) is a recursive equilibrium satisfying $W^1, W^2 \geq 0$. For any $\pi_0 = (\pi_0^1, \pi_0^2) \in [0,1]^2$, let (c, γ, p) be the sequential allocation and price induced by (W, γ^*, p^*) and π_0. If the following (transversality) condition is satisfied:*

$$\lim_{t \to \infty} \beta^t E^{\pi_0^i} W^i(\gamma_{t+1}^{*i}, d_t, \pi_t) = 0,$$

then (c, γ, p) is an equilibrium.

In particular, this means that any bounded recursive equilibrium generates a sequential market equilibrium for any initial belief profile.

Another important property of bounded recursive equilibria is that they generate sequential market equilibria with the lowest possible prices in the class of all sequential market equilibria.

Proposition 7. *For any initial beliefs π_0, if p is a sequential equilibrium price system, and p^* is the bounded recursive equilibrium price, then for every history, d^t, we have $p(d^t) \geq p^*(d_t, \pi_t(d^t))$.*

Note that this implies that the same equality holds for the bubble: $s_t(d^t) \geq S(d_t, \pi_t(d^t))$.

Proof. If $(p_t(d^t))$ is an equilibrium price system, then it has to satisfy the first order conditions of (3.1), which, taking into account the fact that the market clearing condition must hold (i.e., $\gamma_t^i > 0$ for at least one agent), leads to:

$$p_t(d^t) = \max_{i=1,2} \beta E^{\pi_0^i} \left[p_{t+1}(d^{t+1}) + d_{t+1} | d^t \right]. \tag{3.4}$$

Define inductively a sequence of functions, $p_n : \{0,1\} \times [0,1]^2 \to \mathbb{R}$, by

$$p_0(d,\pi) \equiv 0$$
$$p_{n+1}(d,\pi) \equiv \max_{i=1,2} \beta \sum_{d'} \left[\pi^i q^1_{dd'} + (1-\pi^i) q^2_{dd'}\right]\left[d' + p_n(d', \lambda^{dd'}(\pi))\right].$$

Then by Proposition 9, we have $p^* = \lim_{n\to\infty} p_n$. I will show by induction that $p_t(d_t) \geq p_n(d_t, \pi^1_t(d_t), \pi^2_t(d^t))$, for each n, d^t.

For $n = 0$ this is obvious. Suppose that $p_t(d^t) \geq p_n(d^t, \pi_t(d^t))$ for some n and all t and d^t. Then, using (3.2) and (3.4),

$$\begin{aligned} p_t(d^t) &= \max_{i=1,2} \beta E^{\pi^i_0}\left[d_{t+1} + p_{t+1}\big(d^{t+1}\big)|d^t\right] \\ &\geq \max_{i=1,2} \beta E^{\pi^i_0}\left[d_{t+1} + p_n\big(d_{t+1}, \pi_{t+1}(d^{t+1})\big)|d^t\right] \\ &= \max_{i=1,2} \beta \sum_{d_{t+1}\in\{0,1\}} \left[\pi^i q^1_{d_t d_{t+1}} + (1-\pi^i) q^2_{d_t d_{t+1}}\right]\left[d_{t+1} + p_n\left(d_{t+1}, \lambda^{d_t d_{t+1}}(\pi_t(d^t))\right)\right] \\ &= p_{n+1}(d_t, \pi_t(d^t)). \end{aligned}$$

□

In light of the last proposition, the significance of analyzing bounded recursive equilibria is that they provide the (highest) lower bound for the speculative bubble in sequential market equilibria. From now on I will focus solely on bounded recursive equilibria.

Next, I provide tools that allow me to solve for the bounded recursive equilibrium prices. These are then used to obtain the main results of this section. Consider the functional operator $T : B(D \times [0,1]^2) \to B(D \times [0,1]^2)$ defined by[1]

$$Tp(d,\pi) \equiv \max_{i=1,2} \beta \sum_{d'} \left[\pi^i q^1_{dd'} + (1-\pi^i) q^2_{dd'}\right]\left[d' + p^*(d', \lambda^{dd'}(\pi))\right]. \tag{3.5}$$

Proposition 8. *$p^* : D \times [0,1]^2 \to \mathbb{R}$ is the price of some bounded recursive equilibrium if and only if p^* is a fixed point of T.*

[1] For a given topological space X, $B(X)$ denotes the Banach space of real-valued bounded functions defined on X, and $C^0(X)$ denotes the Banach space of all real-valued bounded continuous functions defined on X

Proof. To show necessity take the first order conditions for (3.3) and use the envelope theorem to get $p^* = Tp^*$. For sufficiency, consider $W^i(\gamma, d, \pi) = \gamma^i(p^*(d, \pi) + d)$ and

$$\gamma^{*i}(\gamma, d, \pi) = \begin{cases} 0 & \text{if max in (3.5) is not achieved by } i \\ 1 & \text{if max in (3.5) is not achieved by } (-i) \\ 1/2 & \text{otherwise} \end{cases}$$

It is straightforward to verify this is a stationary equilibrium. □

Now I will briefly go over the properties of T.

Lemma 5. *T has the following properties:*

1. *It is a β-contraction with respect to the* sup*-norm on $B(D \times [0,1]^2)$*
2. *$T\big(C^0(D \times [0,1]^2)\big) \subseteq C^0(D \times [0,1]^2)$*
3. *It is monotone, i.e. for any $p, p' \in B\big(D \times [0,1]\big)$, if $p \geq p'$ then $Tp \geq p'$*

Proof. To get 1. apply Blackwell's sufficient conditions. 2. and 3. are obtained by a direct check. □

This lemma implies the following.

Proposition 9. *There exists a unique $p^* \in B(D \times [0,1]^2)$, such that $p^* = Tp^*$. Moreover, $p^* \in C^0(D \times [0,1]^2)$ and for any $p_0 \in B(D \times [0,1]^2)$, we have $p^* = \lim_{n\to\infty} T^n p_0$.*

The immediate consequence of monotonicity is a very useful lower/upper bound for prices, which will be the primary tool for proving the main result of this section.

Proposition 10. *If for some $p : D \times [0,1]^2 \to \mathbb{R}$ we have $Tp \geq p$ ($Tp \leq p$), then $p^* \geq Tp$ ($p^* \geq Tp$).*

3.2.3 No-Learning Benchmark

It is useful to go over the special case of no learning (i.e., $\pi^1 = 1$, $\pi^2 = 0$) in which I can explicitly solve for the fixed point of T. This is also an important benchmark for the general setup.

As there is no learning, the belief profile π stays the same for any history and the only relevant state variable is the dividend d. This allows me to omit beliefs from the argument of the recursive equilibrium objects throughout this subsection. For instance, $p(d)$ will stand for $p(d, 1, 0)$, and so on. First, I compute the fundamental values. I will use the notation $V_d^{Q^i}$ to denote the fundamental value corresponding to the transition matrix Q^i and the current dividend state, d. Standard recursive considerations lead to:

$$V^{Q^i} \equiv \begin{bmatrix} V_0^{Q^i} \\ V_1^{Q^i} \end{bmatrix} = \frac{\beta}{(1-\beta)(1+\beta(1-q_{00}^i - q_{11}^i))} \begin{bmatrix} q_{01}^i \\ q_{11}^i + \beta(1-q_{00}^i - q_{11}^i) \end{bmatrix}. \tag{3.6}$$

The operator T can be redefined as $T : \mathbb{R}^2 \to \mathbb{R}^2$ with:

$$Tp = T \begin{bmatrix} p(0) \\ p(1) \end{bmatrix} \equiv \begin{bmatrix} \max_{i=1,2} \beta \left(q_{00}^i p(0) + q_{01}^i (1 + p(1)) \right) \\ \max_{i=1,2} \beta \left(q_{10}^i p(0) + q_{11}^i (1 + p(1)) \right) \end{bmatrix},$$

where $p : D \to \mathbb{R}$ is a recursive price system, depending only on the dividend state, d. In this 2-dimensional case, the fixed point of T can be easily computed.

Proposition 11. *If $q_{00}^1 < q_{00}^2$ and $q_{11}^1 < q_{11}^2$, then the recursive equilibrium price system is given by*

$$p^* \equiv \begin{bmatrix} p^*(0) \\ p^*(1) \end{bmatrix} = \frac{\beta}{(1-\beta)(1+\beta(1-q_{00}^1 - q_{11}^2))} \begin{bmatrix} q_{01}^1 \\ q_{11}^2 + \beta(1-q_{00}^1 - q_{11}^2) \end{bmatrix},$$

and in those circumstances the speculative bubble is strictly positive. Analogous result holds for the symmetric case (i.e., $q_{00}^2 < q_{00}^1$ and $q_{11}^2 < q_{11}^1$). Otherwise, the recursive equilibrium price is equal to the market fundamental.

One should note that the pricing above is equal to the fundamental valuation of a hypothetical "hybrid agent" with beliefs concentrated on the transition matrix the first row of which is taken from Q^1 and the second from Q^2.

Proof. First consider the case $q_{00}^1 < q_{00}^2$ and $q_{11}^1 < q_{11}^2$. p^*, above, is the solution to

$$p^*(0) = q_{00}^1 p^*(0) + q_{01}^1 (1 + p^*(1))$$
$$p^*(1) = q_{10}^2 p^*(0) + q_{11}^2 (1 + p^*(1)).$$

The only thing which has to be shown is that the RHSs above constitute the operator T for p^*. This is equivalent to $p^*(0) \leq p^*(1) + 1$. Indeed, this can be easily verified.

To check that $S(d) > 0$ is a straightforward matter. The proof for the symmetric case $q^2_{00} < q^1_{00}$ and $q^2_{11} < q^1_{11}$ is the same.

If neither $q^1_{00} < q^2_{00}$ and $q^1_{11} < q^2_{11}$ nor $q^2_{00} < q^1_{00}$ and $q^2_{11} < q^1_{11}$, I apply the operator T to the market fundamental V to get $TV = V$. Consequently, $p^* = V$. □

I will now illustrate how this technique works on a specific numerical example. These numbers are taken from Harrison and Kreps (1978). Agents of respective groups have the following transition matrices:

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

with the discount factor $\beta = 0.75$. An agent of type one considers the dividend process more volatile than an agent of type two. Therefore, by Proposition 11 there must be a speculative bubble in the equilibrium. First, compute the fundamental value for each pair of beliefs and each state. Using (3.6), I obtain:

$$V_0^{Q^1} = 4/3 = 1.33 \qquad V_1^{Q^1} = 11/9 = 1.22$$
$$V_0^{Q^2} = 16/11 = 1.45 \qquad V_1^{Q^2} = 21/11 = 1.91$$

This reveals that in both states the asset is valued more by agent 2. The market fundamental is therefore $V = V^{Q^2}$.

Using Proposition 11, the recursive equilibrium price is $p_0^* = 24/13 = 1.85$, $p_1^* = 27/13 = 2.08$ with the bubble

$$S(0) = 1.85 - 1.45 = 0.40$$
$$S(1) = 2.08 - 1.91 = 0.17$$

Note that for any history the speculative bubble is on the same level (depending only on the current dividend).

In order to intuitively see why there is a positive speculative bubble, let us apply the operator T to the fundamental value V:

$$T(V) = T\left(\begin{bmatrix} 1.45 \\ 1.91 \end{bmatrix}\right) = 0.75 \begin{bmatrix} \frac{1}{2}1.45 + \frac{1}{2}(1 + 1.91) \\ \frac{1}{4}1.45 + \frac{3}{4}(1 + 1.91) \end{bmatrix} = \begin{bmatrix} 1.64 \\ 1.91 \end{bmatrix}.$$

This means that if the price next period was known to be the fundamental value then, in state $d = 0$, an agent of type one would find the asset worth 1.64, which is strictly more than the fundamental value of 1.45. This occurs because he would subjectively perceive a higher probability of transition to state 1 next period (than type 2), wherein not only is the fundamental price much higher but he also gains the dividend. Consequently, because of risk neutrality and infinite financial resources, he would show an infinite demand, which is not possible in equilibrium. Attempting to set 1.64 as the price in state $d = 0$, would also not work because now agents of type two in state $d = 1$ would show infinite demand. This corresponds to the second iteration of the operator T. Those considerations not only show why the market fundamental is not an equilibrium price but also illustrate why the equilibrium price has to be the limit of iterating the operator T.

3.2.4 Lower Bound for the Speculative Bubble

This subsection investigates under which conditions the recursive equilibrium bubble, S, is strictly positive for each $\pi^1 \neq \pi^2$. Additionally, it will be important to know if the magnitude of the bubble remains significant for a small belief asymmetry (i.e., $\pi^1 \approx \pi^2$).

In the previous subsection I characterized the situations in which there is a strictly positive bubble under the extreme belief differential (i.e., $\pi^1 = 1, \pi^2 = 0$). Specifically, by Proposition 11, if $\beta > 0$, $q_{00}^i < q_{00}^{-i}$ and $q_{11}^i < q_{11}^{-i}$ for some $i = 1, 2$, then $S(d, 1, 0) > 0$. By Proposition 9, p^* is continuous in beliefs and so is the bubble, S. This means that the speculative bubble is strictly positive in some neighborhood of $\pi = (0, 1)$. This, however, does not say how big the speculative bubble is when the beliefs are getting close to each other. It is not even clear if the bubble is strictly positive for all $\pi^1 \neq \pi^2$.

The following proposition provides conditions for the parameters of the model, ensuring that the bubble is not only strictly positive for any asymmetric beliefs ($\pi^1 \neq \pi^2$), but it is also bounded below by a strictly increasing linear function of $|\pi^1 - \pi^2|$. This means that any posterior belief asymmetry results in a speculative bubble of a significant magnitude and allows me to focus on the dynamics of the posteriors rather than the bubble dynamics, which I do in the next section.

Proposition 12. *If $q_{00}^1 < q_{00}^2$, $q_{11}^1 < q_{11}^2$ and one of the following holds:*

1. $V^{Q^1} < V^{Q^2}$ *and* $\beta \leq \frac{q^1_{00}}{q^2_{00}}$

2. $V^{Q^1} > V^{Q^2}$ *and* $\beta \leq \frac{q^1_{11}}{q^2_{11}}$

3. $V_0^{Q^1} \geq V_0^{Q^2}$ *and* $V_1^{Q^2} \geq V_1^{Q^1}$,

then there exist constants $B_0 > 0$ *and* $B_1 > 0$ *such that for each* $d \in \{0, 1\}$,

$$S(d, \pi^1, \pi^2) \geq B_d|\pi^1 - \pi^2|.$$

The proposition comes as a direct corollary to the more general Lemma A1, which provides much bigger range of β for which the conclusion still holds. However, due to a complex statement, Lemma A1 is presented in the Appendix. The idea of the proof is to consider a pricing function, which is linear in beliefs in the following form:

$$p^A(d, \pi^1, \pi^2) \equiv (1 - \pi^1)V_d^{Q^2} + \pi^2 V_d^{Q^1} + (\pi^1 - \pi^2)A_d,$$

where $A = (A_0, A_1)$. Using the identity

$$\begin{bmatrix} V^i(0, \pi^i) \\ V^i(1, \pi^i) \end{bmatrix} = \pi^i V^{Q^1} + (1 - \pi^i)V^{Q^2},$$

it can be seen that for $A_d = \max\{V_d^{Q^1}, V_d^{Q^2}\}$, the function p^A is equal to the market fundamental, and any higher values of A_d will generate a bubble proportional to the belief differential. For each value of the parameters, one can explicitly compute the highest values of the constants A_0 and A_1 so that $Tp^A(d, \pi^1, \pi^2) \geq p^A(d, \pi^1, \pi^2)$ for any $\pi \in [0, 1]^2$ and $d \in D$. Using the monotonicity of T (Proposition 10), for such constants the equilibrium price p^* satisfies $p^* \geq p^A$. In order to obtain Proposition 12, I need to make sure the constants A_d, obtained in this way, satisfy $A_d > \max\{V_d^{Q^1}, V_d^{Q^2}\}$. Then I just put $B_d \equiv A_d - \max\{V_d^{Q^1}, V_d^{Q^2}\}$. The following proposition specifies the highest values of Bs this technique allows us to achieve.

Proposition 13. B_0 *and* B_1 *in Proposition 12 can be taken as follows:*

In case 1: $B_0/R_2 = B_1 = \dfrac{V_1^{Q^2} - V_1^{Q^1} - R_1(V_0^{Q^2} - V_0^{Q^2})}{R_1R_2 - 1}$

In case 2: $B_0 = B_1/R_1 = \dfrac{V_0^{Q^1} - V_0^{Q^2} - R_2(V_1^{Q^1} - V_1^{Q^2})}{R_1R_2 - 1}$

In case 3: $B_0 = \dfrac{V_0^{Q^1} - V_0^{Q^2} - R_2(V_1^{Q^1} - V_1^{Q^2})}{R_1R_2 - 1}$ *and* $B_1 = \dfrac{V_1^{Q^2} - V_1^{Q^1} - R_1(V_0^{Q^2} - V_0^{Q^2})}{R_1R_2 - 1}$, *where the constants R_1, R_2 are given by*

$$R_1 = \begin{cases} \frac{\left(1-\sqrt{\beta q_{00}^1 q_{00}^2}\right)^2}{\beta q_{01}^1 q_{01}^2} & \textit{if } q_{00}^1 \leq \beta q_{00}^2 \\ \frac{1-\beta q_{00}^2}{\beta q_{01}^2} & \textit{otherwise} \end{cases}$$

$$R_2 = \begin{cases} \frac{\left(1-\sqrt{\beta q_{11}^1 q_{11}^2}\right)^2}{\beta q_{10}^1 q_{10}^2} & \textit{if } q_{11}^1 \leq \beta q_{11}^2 \\ \frac{1-\beta q_{11}^2}{\beta q_{10}^2} & \textit{otherwise} \end{cases}$$

This proposition also follows directly from Lemma A1.

Now I will show how the values for Bs compare with the exact solution of the numerical example from the previous subsection. Since I allow for learning, the matrices Q^1 and Q^2 constitute the support of agents' beliefs. Without loss of generality I can assume their prior beliefs satisfy $\pi_0^1 > \pi_0^2$. In particular this assumption implies that $\pi_t^1(d^t) > \pi_t^2(d^t)$ for any t and any history d^t.

Since $V^{Q^2} > V^{Q^1}$ and $\beta = 3/4 = \frac{1/2}{2/3} = \frac{q_{00}^1}{q_{00}^2}$, the assumption of Proposition 12 is satisfied. Using the formulas from Proposiiton 13, I get $B_0 = .28$ and $B_1 = 0.11$. The linear lower bound for the equilibrium price bubble is then

$$S(0, \pi^1, \pi^2) \geq 0.28 \cdot |\pi^1 - \pi^2|$$
$$S(0, \pi^1, \pi^2) \geq 0.11 \cdot |\pi^1 - \pi^2|.$$

This compares to the exact values computed in the previous section for the no-learning environment (i.e., $\pi^1 = 1, \pi^2 = 0$). We had $S(0,1,0) = 0.40$ and $S(1,1,0) = 0.17$, suggesting that the linear lower bound captures a decent amount of speculation.

I close this section with a brief comment on high values of β, for which the assumption of Proposition 12 is not satisfied. Lemma A1 provides a much bigger range of βs but still does not work for $\beta \approx 1$. The question is: what is happening if β is high? Is it that the bubble cannot be bounded away from zero by a linear function of beliefs, or is it just a deficiency of the specific technique? Furthermore, is it true that for high values of β the bubble becomes insignificant for some small or intermediate values of $|\pi^1 - \pi^2|$? These questions are not addressed in this paper but some numerical experiments suggest that for $0 \ll \pi^1 \approx \pi^2 \ll 1$, and high values of β the bubble becomes very small, relative to its highest value under no learning. Moreover, these numerical results do not even elucidate that for some small values of $|\pi^1 - \pi^2|$, the bubble is strictly positive.

3.3 The Dynamics of the Speculative Bubble under the True Transition Matrix

In this section I analyze the dynamics of posterior beliefs for a given true data-generating process. According to Proposition 12, this translates into the dynamics of the speculative bubble in the circumstances specified therein. In Section 3.1, I state the main result of the paper, which gives the condition for the true data-generating process to achieve a perpetual disagreement with probability one. I then discuss the idea of the proof. Subsection 3.2 provides some further results concerning the posterior dynamics, most notably the generic supplement to the main theorem.

3.3.1 The Main Result

I will now apply the technique associated with the coin flip example from the introduction into the Markov environment of my model.

Let $Q = (q_{dd'})_{0 \le d, d' \le 1}$ be the transition matrix of the true data generating process.

Proposition 14. *If $Q = (q_{dd'})$ satisfies*

$$\frac{q_{00}}{q_{01}} \log\left(\frac{q_{00}^2}{q_{00}^1}\right) + \log\left(\frac{q_{01}^2}{q_{01}^1}\right) + \log\left(\frac{q_{10}^2}{q_{10}^1}\right) + \frac{q_{11}}{q_{10}} \log\left(\frac{q_{11}^2}{q_{11}^1}\right) = 0, \tag{3.7}$$

then for any $\alpha < 1$ the process of the posterior belief profile π satisfies $|\pi_t^1 - \pi_t^2| > \alpha|\pi_0^1 - \pi_0^2|$ infinitely often with Q-probability one.

This proposition, together with Propositions 7 and 12, leads to the following theorem as an immediate corollary:

Theorem 15 (Main Theorem)**.** *Suppose the assumptions of Proposition 12 are satisfied and B_0 and B_1 are as in Proposition 13. Let s be a sequential market equilibrium bubble process and $\alpha < 1$. If the true data-generating process transition matrix Q satisfies (3.7), then $s_t > \alpha \cdot \max\{B_0, B_1\} \cdot |\pi_0^1 - \pi_0^2|$ infinitely often with Q-probability one:*

$$\Pr^Q(s_t > \alpha \cdot \max\{B_0, B_1\} \cdot |\pi_0^1 - \pi_0^2| \text{ infinitely often}) = 1.$$

This theorem says that if Q^1, Q^2, and β satisfy the assumptions of Proposition 12, then the condition for the true data generating matrix Q to generate the bubble of high magnitude infinitely often with probability one is given by (3.7). This condition is not generic. It defines a 1-dimensional manifold in the set of 2×2 probabilistic matrices, which is of dimension 2. Hence, the Lebesgue measure of the set of Q's satisfying (3.7) is zero. In the next subsection I will provide a generic complement for the main theorem.

Throughout the rest of this subsection I discuss the derivation of Proposition 14, leaving some technical details for the Appendix. I also introduce notation that is needed later.

Using Bayes' rule, upon observing the history d^t, the current beliefs at time t are given by $\pi_t^i(d^t) = \left(1 + \frac{1-\pi_0^i}{\pi_0^i} L_t(d^t)\right)^{-1}$, where $L_t(d^t)$ denotes the current likelihood ratio of the history d^t for matrices Q^2 vs Q^1 (i.e., $L_t(d^t) \equiv \frac{\Pr^{Q^2}\{d^t\}}{\Pr^{Q^1}\{d^t\}}$). The current belief difference becomes:

$$\pi_t^1 - \pi_t^2 = (\pi_0^1 - \pi_0^2) \cdot (L_t^{-1}\pi_0^1 + 1 - \pi_0^1)^{-1} \cdot (L_t(1 - \pi_0^2) + \pi_0^2)^{-1},$$

and can be fully analyzed in terms of the process $L = (L_t)_t$. Note that the original beliefs are re-achieved whenever $L_t = 1$. The reason for this is that in those periods the current history does not favor Q^1 or Q^2. Similarly, each time, $L_t > 1$ it means that the current history favors the theory Q^2 vs. the theory Q^1 and both agents move towards theory Q^2 in their current posteriors. If $L_t < 1$, the situation is reversed.

Without loss of generality, I assume that the initial beliefs satisfy $1 - \pi_0^1 = \pi_0^2 = \epsilon < 1/2$. This assumption is purely for convenience. Otherwise, it could happen that for some future histories the belief difference is bigger than the original one. In that situation, the initial difference would not be a natural benchmark and the conclusion of Lemma 14 would be true even for some $\alpha > 1$.[2] We have the following:

$$|\pi_t^1 - \pi_t^2| = \pi_t^1 - \pi_t^2 = (1 - 2\epsilon) \cdot (L_t^{-1}(1 - \epsilon) + \epsilon)^{-1} \cdot (L_t(1 - \epsilon) + \epsilon)^{-1}$$

We can algebraically verify that for any $0 < \alpha < 1$, $\pi_t^1 - \pi_t^2 > \alpha \cdot (\pi_0^1 - \pi_0^2)$ if and only if

$$L_t \in I_\alpha \equiv \left(1 - \frac{2}{\sqrt{1 + \frac{4\alpha\epsilon(1-\epsilon)}{1-\alpha}} + 1}, 1 + \frac{2}{\sqrt{1 + \frac{4\alpha\epsilon(1-\epsilon)}{1-\alpha}} - 1}\right).$$

Clearly, $\pi_t^1 - \pi_t^2 = \pi_0^1 - \pi_0^2$ if and only if $L_t = 1$. It is therefore consistent to define $L_0 \equiv 1$. In order to analyze the dynamics of speculation for some given Q, it is enough

[2] If $1 - \pi_0^1 \neq \pi_0^2$, then for any $\delta > 0$ π_0^1, π_0^2 can be shown to be posterior beliefs of some priors $\tilde{\pi}_0^1, \tilde{\pi}_0^2$ such that $|1 - \tilde{\pi}_0^1 - \tilde{\pi}_0^2| < \delta$.

to see how much time the process L_t spends in the interval I_α, for $\alpha > 0$ under that particular Q.

I will now analyze the dynamics of the process L. It is easy to see that for a given history d^t, $L_t(d^t)$ is given by

$$L_t(d^t) = \prod_{\tau=0}^{t-1} \left(\frac{q^1_{d_\tau d_{\tau+1}}}{q^2_{d_\tau d_{\tau+1}}} \right).$$

For algebraic convenience, I will consider the process L in log terms:

$$l_t \equiv \log L_t = \sum_{\tau=0}^{t-1} \log \left(\frac{q^1_{d_\tau d_{\tau+1}}}{q^2_{d_\tau d_{\tau+1}}} \right). \tag{3.8}$$

Defining a process Δ by $\Delta_t \equiv (d_{t-1}, d_t) \in D^2$, $l = (l_t)_t$ can be expressed as $l_t = \sum_{\tau=0}^{t-1} f(\Delta_\tau)$, where function $f : D^2 \to \mathbb{R}$ is defined via (3.8).

The process Δ is a 4-state Markov chain with the transition matrix

$$\bar{Q} \equiv \begin{bmatrix} q_{00} & q_{01} & 0 & 0 \\ 0 & 0 & q_{10} & q_{11} \\ q_{00} & q_{01} & 0 & 0 \\ 0 & 0 & q_{10} & q_{11} \end{bmatrix},$$

$\bar{Q}$ is ergodic if and only if Q is ergodic, which occurs if and only if $0 < q_{01}q_{10} < 1$. Hence, in such situations, $\bar{Q}$ has the unique ergodic distribution ν^Q. Simple algebra reveals the following:

$$\nu^Q = \frac{1}{\frac{q_{00}}{q_{01}} + 2 + \frac{q_{11}}{q_{10}}} \left[\frac{q_{00}}{q_{01}}, 1, 1, \frac{q_{11}}{q_{10}} \right].$$

Now, the intuition associated with the ergodic theorem would be as follows: The process l should be recurrent if the increment $f(\Delta_t)$ has the expected value of zero on the stationary (long run) distribution ν^Q. In other words, Q should be such that

$$E^{\nu^Q} f(d_t, d_{t+1}) \equiv \nu^Q_{00} f(0,0) + \nu^Q_{01} f(0,1) + \nu^Q_{10} f(1,0) + \nu^Q_{11} f(1,1) = 0.$$

This intuition is confirmed by the following lemma:

Lemma 6. *Let $I \subseteq \mathbb{R}$ be an interval that can be reached by the process l with positive probability, then $\Pr^Q(l_t \in I$ infinitely often$) = 1$ if and only if Q satisfies (3.7).*

This result, together with all the previous considerations, completes the proof of Proposition 14 and subsequently the main theorem: since the process l visits any interval infinitely often, so it does the interval $\log(I_\alpha)$. Therefore, the process L visits I_α infinitely often with probability one.

The idea with regard to the proof of Lemma 6 is to transform the process l to a standard independent random walk $\bar{l}$, with mean zero using the appropriate stopping times. The detailed proof can be found in the Appendix. In what follows I briefly explain the way the process l is transformed into the iid increments process $\bar{l}$. This will be of use in the final subsection, wherein I discuss some dynamic properties of the process of the posterior beliefs, and I will do it in terms of the process $\bar{l}$.

Define a sequence of optional stopping times $\tau = (\tau_n)_{n\in\mathbb{N}}$:

$$\tau_1 \equiv \min\{t > 0 | \Delta_t = (0,0)\}$$
$$\tau_{n+1} \equiv \min\{t > \tau_n | \Delta_t = (0,0)\},$$

The realization of this random sequence gives us the number of periods in which the process Δ hits the state $(0,0)$. Now, I define the process $\bar{l}$ by

$$\bar{l}_t = l_{\tau_t}.$$

The process $\bar{l}$ shows the values of the process l only in periods in which $\Delta = (0,0)$. The strong Markov property for the process Δ roughly says that at any given random period, the future and the past of the process are independent, and also that the distribution of the future is solely determined by the current state. Now, since for all periods τ_t, the current state is $(0,0)$, it necessarily follows that the distribution of the future is the same for all periods τ_t. This means the increments of the process $\bar{l}$ are iid. It can also be argued that under the condition (3.7) they have zero expected value, so the process $\bar{l}$ has no drift and by standard properties of iid random walks it is recurrent. In other words, any interval that can be visited is visited infinitely often with probability one.

3.3.2 Further Insight into Bubble Dynamics

In this subsection I present several further results concerning bubble dynamics. The first result is a generic complement of the main theorem, and the second is that bubbles of high magnitude must appear very infrequently on typical sample paths. Finally, I present the so-called arc-sine law, which leads to a rather surprising conclusion. Even if none of the agents' models is a priori favored by the data generating process, on a typical sample path one of the models will be doing better most of the time. The results in this section are expressed in terms of the process $\bar{l}$ defined in the end of the last subsection. These results automatically translate into the log-likelihood process l and hence the bubble process s, but this transition requires some notational complications caused by the fact that the process $\bar{l}$ samples the process l at random times.

The problem with the main result is that it is not generic. In order to have persistent speculation we need to assume that the data generating matrix, Q, satisfies condition (3.7). As mentioned before, the set of Qs, satisfying (3.7) is a 1-dimensional manifold in the set of 2×2-probabilistic matrices. I will now show that if Q is approaching this manifold, then the persistency of the speculative bubble is going to infinity.

Lemma 7. *Let*

$$\mu = \frac{q_{00}}{q_{01}} \log\left(\frac{q_{00}^2}{q_{00}^1}\right) + \log\left(\frac{q_{01}^2}{q_{01}^1}\right) + \log\left(\frac{q_{10}^2}{q_{10}^1}\right) + \frac{q_{11}}{q_{10}} \log\left(\frac{q_{11}^2}{q_{11}^1}\right),$$

then for any interval I we have $E^Q \sum_{t=1}^{\infty} 1_{\{\bar{l}_t \in I\}} \sim \frac{1}{\mu}$. In particular, $E^Q \sum_{t=1}^{\infty} 1_{\{\bar{l}_t \in I\}} \to \infty$ as $\mu \to 0$.

The proof, being an easy application of the weak law of large numbers, can be found in the Appendix. Taking $\alpha < 1$ and $I = I_\alpha$ in Lemma 7, we get the generic complement to the main result.

Theorem 16. *If $0 < \alpha < 1$ and μ as in Lemma 7, then $E^Q \sum_{t=1}^{\infty} 1_{\{s_t > \alpha \cdot |\pi_0^1 - \pi_0^2|\}} \to \infty$ as $\mu \to 0$.*

As stated above, the expected total number of future periods in which the bubble exceeds the level $\alpha \cdot |\pi_0^1 - \pi_0^2|$ goes to infinity as μ goes to zero.

I will now examine the question of how often, on average, high bubbles should arise in this model. Since the process $\bar{l}$ is a standard random walk, we have the following result (cf. Theorem 2 in XII.2 of Feller, 1966):

Proposition 17. *For any data generating process Q and any interval $I \subseteq \mathbb{R}$, such that $\bar{l}_0 \notin I$*

$$E^Q \inf\{t > 0 | \bar{l}_t \in I\} = \infty,$$

This means that the expected waiting time for the process $\bar{l}$ to reach any given level different than the initial one is infinity. After translating the above to the process s_t, one can conclude that the expected waiting time for the bubble to reach any given magnitude higher than the current one is infinity. This suggests that even though under the conditions of the main theorem the bubbles are infinitely persistent, on average the periods of high magnitude rarely prevails. In practice, this infinite expected waiting time shows up in simulations in the form of having many runs with almost no bubbles. On some runs, however, we still observe a relatively high frequency of bubbles of a significant magnitude.

Another interesting property of the dynamics of $\bar{l}$ is provided by the arc sine law, which is a standard result in the theory of the iid random walk, and it extends naturally to my setup. For this to hold, it is required that $E\Delta_0^2 < \infty$, which can be easily verified. Necessarily, Theorems 1a and 2 in XII.8 of Feller (1966) imply the following:

Proposition 18. *Let* $\Pi_n \equiv \#\{t > 0 | \bar{l}_t > 0\}$, *then*

$$\frac{\Pi_n}{n} \to^d Beta(1/2, 1/2),$$

This means that the distribution of the fraction of strictly positive terms in $\bar{l}$, which can be thought of as an approximation of the fraction of time when the theory Q^2 has higher likelihood than theory Q^1, is asymptotically the Beta(1/2,1/2)-distribution with density $\frac{1}{\pi\sqrt{\alpha(1-\alpha)}}$.[3] It is unbounded at the endpoints 0 and 1 and has its minimum at $\frac{1}{2}$. This means that for long samples we are more likely to see either theory Q^1 or Q^2 explaining the data better rather than seeing close to equal lengths of one theory dominating. Since the speculative bubble is at its highest whenever the data is not conclusive, this suggests that, even though the periods of high bubble magnitude appear infinitely often, they will be observed relatively rarely and will typically be surrounded by much longer periods of low bubble magnitude. What we should generally see is a long period where theory Q^1 for example, dominates and both agents assign most of the beliefs to that theory. Indeed, during this time the speculative bubble is very small. At some point, the data will start to favor theory Q^2 which is associated with the beliefs being moved towards Q^2. During that time of transition the bubble explodes

[3] The cdf of this distribution is given by the scaled arc-sine function, which justifies the name of the law.

because one agent is more reluctant than the other to accept the new theory. Once both types settle their beliefs close to Q^2, speculative bubble vanishes again. This pattern is continued forever with probability one.

Bibliography

Acemoglu, D., V. Chernozhukov, and M. Yildiz (2007). Learning and disagreement in an uncertain world. *Unpublished manuscript*.

Alen, F., S. Morris, and A. Postlewaite (1993). Finite bubbles with short sale constraints and asymmetric information. *Journal of Economic Theory 61*.

Aumann, R. (1976). Agreeing to disagree. *The Annals of Statistics 4*.

Baccheta, P. and E. van Wincoop (2006a). Can information heterogeneity explain the exchange rate determination puzzle. *The American Economic Review 96*(3).

Baccheta, P. and E. van Wincoop (2006b). Higher order expectations in asset pricing. *Working Paper*.

Beker, P. F. and E. Espino (2009). The dynamics of efficient asset trading with heterogeneous beliefs. *unpublished manuscript*.

Blackwell, D. and L. Dubins (1962). Merging of opinions with increasing information. *The Annals of Mathematical Statistics 33*(3).

Blume, L. and D. Easley (2006). If you're so smart, why aren't you rich? *Econometrica 74*(4).

Bossaerts, P. (1995). The econometrics of learning in financial markets. *Economic Theory 11*(1).

Feller, W. (1966). *An Introduction to Probability Theory and Its Applications*, Volume II. Jon Wiley & Sons, Inc.

Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American Economic Review 70*.

Harrison, J. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics 92*(2).

Hellwig, M. (1980). On aggregation of information in competitive markets. *Journal of Economic Theory 22*.

Kyle, A. S. (1985). Continuous auction and insider trading. *Econometrica 53*(6).

Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of Economic Theory 12*.

Morris, S. (1996). Speculative investor behavior and learning. *The Quarterly Journal of Economics 111*(4).

Sandroni, A. (2000). Do makets favor agents able to make accurate predictions? *Econometrica 68*(6).

Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of Political Economy 111*(6).

Stokey, N. L., R. L. Lucas, and E. C. Prescott (2004). *Recursive Methods in Economic Dynamics* (7 ed.).

Appendix A

Proof of Lemma 1. First note, that since we defined:

$$\lambda^d(\pi)(A \times \Phi \times Q) = \Pr^\pi(a_1 \in A \wedge \phi \in \Phi \wedge q \in Q | d_1 = d)$$
$$= \frac{\int_{\phi\in\Phi}\int_{q\in Q}\int_{a_1\in A}\int_{a_0\in\mathcal{A}} \pi(da_0,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}{\int_{\phi\in\mathbf{\Phi}}\int_{q\in\mathcal{Q}}\int_{a_0,a_1\in\mathcal{A}^2} \pi(da_t,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}$$

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

$$\lambda^d(\pi)(da_1,d\phi,dq) \text{"="} \frac{\int_{a_0\in\mathcal{A}} \pi(da_1,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}{\int_{\phi\in\mathbf{\Phi}}\int_{q\in\mathcal{Q}}\int_{a_0,a_1\in\mathcal{A}^2} \pi(da_1,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}$$

Using this we have:

$$\Pr^{\lambda^d(\pi)}(a_0 \in A_0 \ldots, a_s \in A_s, \phi \in \Phi, q \in Q) =$$
$$= \int_{\phi\in\Phi}\int_{q\in Q}\int_{(a_0,\ldots,a_s)\in A_0\times A_s} \lambda^d(\pi)(da_0,d\phi,dq)q(a_0,da_1)$$
$$= \int_{\phi\in\Phi}\int_{q\in Q}\int_{(a_1,\ldots,a_{s+1})\in A_0\times\ldots\times A_s} \frac{\int_{a_0\in\mathcal{A}} \pi(da_0,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}{\int_{\phi\in\mathbf{\Phi}}\int_{q\in\mathcal{Q}}\int_{a_0,a_1\in\mathcal{A}^2} \pi(da_1,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}$$
$$= \frac{\int_{\phi\in\Phi}\int_{q\in Q}\int_{a_0\in\mathcal{A}}\int_{(a_1,\ldots,a_{s+1})\in A_0\times\ldots\times A_s} \pi(da_0,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}{\int_{\phi\in\mathbf{\Phi}}\int_{q\in\mathcal{Q}}\int_{a_0,a_1\in\mathcal{A}^2} \pi(da_1,d\phi,dq)q(a_0,da_1)\phi_{a_1}(d_1)}$$
$$= \Pr^\pi(a_1 \in A_0, \ldots, a_{s+1} \in A_s, \phi \in \Phi, q \in Q | d_1 = d)$$

□

Proof of Theorem 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 $\gamma^1 \equiv (\gamma^1_1, \gamma^1_2(d^1), \gamma^1_3(d^2)\ldots) \geq 0$ (following Stokey et al. (2004) we call it a plan) to maximize:

$$u(\gamma, \gamma^1_0, \pi^1, \pi^2) \equiv \lim_{T\to\infty} u_T(\gamma, \gamma^1_0, \pi^1, \pi^2)$$

where $u_T(\gamma, \gamma^1_0, \pi^1, \pi^2) \equiv \mathrm{E}^{\pi^1_0} \sum_{t=0}^{T} \beta^t [p_t(\gamma^1_t - \gamma^1_{t+1}) + \gamma_t d_t]$ taken as given $\gamma^1_0 = 0$.

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

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

$$V^*(\gamma^1_0, \pi^1_0, \pi^2_0) = \sup_{\gamma\in\Gamma} u(\gamma, \pi^1_0, \pi^2_0) \tag{A.1}$$

for each $\gamma^1_0 \geq 0$.

We will show that $V(\gamma^1_0, \pi^1_0, \pi^2_0) = V^*(\gamma^1_0, \pi^1_0, \pi^2_0)$ and that proof will imply that γ^{*1} attains the sup in (A.1). First we prove that

$$V(\gamma^1_0, \pi^1_0, \pi^2_0) \geq u(\gamma, \gamma^1_0, \pi^1_0, \pi^2_0) \tag{A.2}$$

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

$$V(\gamma^1_0, \pi^1_0, \pi^2_0) = u(\gamma^{*i}, \gamma^1_0, \pi^1_0, \pi^2_0) \tag{A.3}$$

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

$$\begin{aligned}
V(\gamma^i_0, \pi^1_0, \pi^2_0) &= \max_{\gamma'\geq 0} \Big\{ (\gamma^i_0 - \gamma')p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\pi^1_0}\big(V(\gamma', \lambda(d_1|\pi^1_0), \lambda(d_1|\pi^2_0)) + \gamma' d_1\big)\Big\} \\
&\geq (\gamma^i_0 - \gamma^1_1)p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\pi^1_0}\big(V(\gamma^1_1, \lambda(d_1|\pi^1_0), \lambda(d_1|\pi^2_0)) + \gamma^1_1 d_1\big) \\
&= (\gamma^i_0 - \gamma^1_1)p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\pi^1_0}\Bigg(\max_{\gamma'\geq 0}\Big\{(\gamma^1_1 - \gamma')p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\lambda(d^1|\pi^1_0)}\big(V(\gamma', \lambda(d_1|\pi^1_0), \lambda(d_1|\pi^2_0)) + \gamma' d_1\big)\Big\} + \gamma^1_1 d_1\Bigg) \\
&= (\gamma^i_0 - \gamma')p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\pi^1_0}\Bigg(\max_{\gamma'\geq 0}\Big\{(\gamma^1_1 - \gamma')p(\pi^1_0, \pi^2_0) + \beta\big(V(\gamma', \lambda(d_2|\pi^1_0), \lambda(d_2|\pi^2_0)) + \gamma' d_2\big)\Big\} + \gamma^1_1 d_1\Bigg) \\
&\geq (\gamma^i_0 - \gamma')p(\pi^1_0, \pi^2_0) + \beta \mathrm{E}^{\pi^1_0}\Big((\gamma^1_1 - \gamma^1_2)p(\pi^1_0, \pi^2_0) + \beta\big(V(\gamma^1_2, \lambda(d_2|\pi^1_0), \lambda(d_2|\pi^2_0)) + \gamma^1_2 d_2\big) + \gamma^1_1 d_1\Big) \\
&= u_1(\gamma^1, \gamma^1_0, \pi^1, \pi^2) + \beta^2 \mathrm{E}^{\pi^1_0}\big(V(\gamma^1_2, \pi_1(d_2|\pi^1_0), \lambda(d_2|\pi^2_0)) + \gamma^2_1 d_2\big)
\end{aligned}$$

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

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

for all T. Now, using the assumption that $V \geq 0$ we conclude that (A.2) holds. Having (A.2) 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 (A.3). 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. $\square$

Appendix B

Lemma B1. *Suppose* $q_{00}^1 < q_{00}^2$ *and* $q_{11}^1 < q_{11}^2$*. Define the following constants:*

$$R_1 = \begin{cases} \frac{\left(1-\sqrt{\beta q_{00}^1 q_{00}^2}\right)^2}{\beta q_{01}^1 q_{01}^2} & \text{if } q_{00}^1 \leq \beta q_{00}^2 \\ \frac{1-\beta q_{00}^2}{\beta q_{01}^2} & \text{otherwise} \end{cases}$$

$$R_2 = \begin{cases} \frac{\left(1-\sqrt{\beta q_{11}^1 q_{11}^2}\right)^2}{\beta q_{10}^1 q_{10}^2} & \text{if } q_{11}^1 \leq \beta q_{11}^2 \\ \frac{1-\beta q_{11}^2}{\beta q_{10}^2} & \text{otherwise} \end{cases}$$

If one of the following conditions hold:

1. $V^{Q^1} > V^{Q^2}$ *and* $V_0^{Q^1} - V_0^{Q^2} > R_2\left(V_1^{Q^1} - V_1^{Q^2}\right)$
2. $V^{Q^1} > V^{Q^2}$ *and* $V_1^{Q^2} - V_1^{Q^1} > R_1\left(V_0^{Q^2} - V_0^{Q^1}\right)$
3. $V_0^{Q^1} \geq V_0^{Q^2}$ *and* $V_1^{Q^1} \leq V_1^{Q^2}$

then the recursive equilibrium price bubble, S, satisfies:

$$S(d, \pi^1, \pi^2) \geq (A_d - \max_i\{V_d^{Q^i}\})|\pi^1 - \pi^2| \tag{B.1}$$

where the constants A_0, A_1 are given by

$$A_0 = \frac{R_1 R_2 V_0^{Q^1} - R_2(V_1^{Q^1} - V_1^{Q^2}) - V_0^{Q^2}}{R_1 R_2 - 1}$$
$$A_1 = \frac{R_1 R_2 V_1^{Q^2} - R_1(V_0^{Q^2} - V_0^{Q^1}) - V_1^{Q^1}}{R_1 R_2 - 1}$$

and satisfy

$$A_0 > \max_i\{V_0^{Q^i}\} \text{ and } A_1 > \max_i\{V_1^{Q^i}\} \tag{B.2}$$

Proof. First I prove (B.2). Using the formula for A_0 and A_1 we get:

$$\begin{array}{c} A_0 - V_0^1 = \frac{A_1 - V_1^1}{R_1} = \frac{R_2(V_1^{Q^2} - V_1^{Q^1}) - (V_0^{Q^2} - V_0^{Q^1})}{R_1 R_2 - 1} \\ \frac{A_0 - V_0^{Q^2}}{R_2} = A_1 - V_1^{Q^2} = \frac{R_1(V_0^{Q^1} - V_0^{Q^2}) - (V_1^{Q^1} - V_1^{Q^2})}{R_1 R_2 - 1} \end{array} \tag{B.3}$$

Simple algebra reveals $R_1, R_2 > 1$. Using (B.3) it is easy to verify, that any of the assumptions, 1 or 2, imply (B.2). To get that assumption 3 also implies (B.2) use the formula (3.6) to get

$$V_0^1 - V_0^2 = \frac{\beta(q_0^2 - q_0^1 - \beta(q_1^2 - q_1^1) + \beta(q_0^1 q_1^2 - q_0^2 q_1^1))}{(1-\beta)(1+\beta(1-q_0^1-q_1^1))(1+\beta(1-q_0^2-q_1^2))}$$
$$V_1^1 - V_1^2 = \frac{\beta(-(q_1^2 - q_1^1) + \beta(q_0^2 - q_0^1) + \beta(q_0^1 q_1^2 - q_0^2 q_1^1))}{(1-\beta)(1+\beta(1-q_0^1-q_1^1))(1+\beta(1-q_0^2-q_1^2))}$$

which implies that we always have: $V_0^{Q^1} - V_0^{Q^2} \geq V_1^{Q^1} - V_1^{Q^2}$. This means (B.3) implies (B.2) also under assumption 3.

Now I move on to prove (B.1). I will also need the following identities, which are easy to verify,

$$\begin{array}{c} R_1 = \frac{A_1 - V_1^{Q^1}}{A_0 - V_0^{Q^1}} \\ R_2 = \frac{A_0 - V_0^{Q^2}}{A_1 - V_1^{Q^2}} \end{array} \tag{B.4}$$

Without loss of generality we may assume $\pi^1 \geq \pi^2$ and in that case for any d, d' we have $\lambda^{dd'}(\pi^1) \geq \lambda^{dd'}(\pi^2)$. This means by induction, that whenever agents start with beliefs satisfying $\pi^1 > pi^2$ than no matter what history of dividends they observe, their updated beliefs will also satisfy this inequality, $\pi_t^1(d^t) \geq \pi_t^2(d^t)$.

Let

$$p^A(d, \pi^1, \pi^2) \equiv (1-\pi^1)V_d^{Q^2} + \pi^2 V_d^{Q^1} + (\pi^1 - \pi^2)A_d$$

Note that

$$p^A(d, \pi^1, \pi^2) - V(d, \pi^1, \pi^2) = (A_d - \max_i\{V^{Q^i}(0)\})|\pi^1 - \pi^2|$$

Hence by Proposition 10 in order to get (B.1) it is enough to show $Tp^A \geq p^A$. We will do it for each state, d, separately.

First take $d = 0$. We have

$$
\begin{aligned}
Tp^A(0,\pi^1,\pi^2) &\geq T^1p^A(0,\pi^1,\pi^2) && \text{(B.5)}\\
&\equiv \beta(\pi^1 q^1_{00} + (1-\pi^1)q^2_{00})\left[(1-\lambda^{00}(\pi^1))V_0^{Q^2} + \lambda^{00}(\pi^2)V_0^{Q^1} + (\lambda^{00}(\pi^1)-\lambda^{00}(\pi^2))A_0\right]\\
&+ \beta(\pi^1 q^1_{01} + (1-\pi^1)q^2_{01})\left[1 + (1-\lambda^{01}(\pi^1))V_1^{Q^2} + \lambda^{01}(\pi^2)V_1^{Q^1} + (\lambda^{01}(\pi^1)-\lambda^{01}(\pi^2))A_1\right]
\end{aligned}
$$

Using

$$
\begin{aligned}
\lambda^{00}(\pi) &= \frac{\pi q^1_{00}}{\pi q^1_{00} + (1-\pi)q^2_{00}} = 1 - \frac{(1-\pi)q^2_{00}}{\pi q^1_{00} + (1-\pi)q^2_{00}}\\
\lambda^{01}(\pi) &= \frac{\pi q^1_{01}}{\pi q^1_{01} + (1-\pi)q^2_{01}} = 1 - \frac{(1-\pi)q^2_{01}}{\pi q^1_{01} + (1-\pi)q^2_{01}}
\end{aligned}
$$

we can see that both p^A and T^1p^A are linear in π^1. Also note that $T^1p^A = p^A$ for $\pi^1 = \pi^2$. In order to conclude $Tp^A(0,\pi^1,\pi^2) \geq p^A(0,\pi^1,\pi^2)$ it is enough then to prove that $T^1p^A \geq p^A$ for $pi^1 = 1$. Incidentally, note that we made use of the notation T^1p^A which was introduced when we defined operator T in (3.5).

We have:

$$
\begin{aligned}
&T^1p^A(0,1,\pi^2) - p^A(0,1,\pi^2) =\\
&=\beta q^1_{00}\left[\left(1 - \frac{(1-\pi^2)q^2_{00}}{\pi^2 q^1_{00} + (1-\pi^2)q^2_{00}}\right)V_0^{Q^1} + \frac{(1-\pi^2)q^2_{00}}{\pi^2 q^1_{00} + (1-\pi^2)q^2_{00}} \cdot A_0\right]\\
&\quad + \beta q^1_{01}\left[1 + (1 - \frac{(1-\pi^2)q^2_{01}}{\pi^2 q^1_{01} + (1-\pi^2)q^2_{01}})V_1^{Q^1} + \frac{(1-\pi^2)q^2_{01}}{\pi^2 q^1_{01} + (1-\pi^2)q^2_{01}} \cdot A_1\right]\\
&\quad - (1-\pi^2)(A_0 - V_0^{Q^1}) - V_0^{Q^1}\\
&=\beta(1-\pi^2)\left[\frac{q^1_{00}q^2_{00}}{\pi^2 q^1_{00} + (1-\pi^2)q^2_{00}}(A_0 - V_0^{Q^1}) + \frac{q^1_{01}q^2_{01}}{\pi^2 q^1_{01} + (1-\pi^2)q^2_{01}}(A_1 - V_1^{Q^1}) - \beta^{-1}(A_0 - V_0^{Q^1})\right]
\end{aligned}
$$

To show $T^1p^A(0,1,\pi^2) \geq p^A(0,1,\pi^2)$ for all $\pi^2 \in [0,1]$ it is enough to prove:

$$
\frac{q^1_{00}q^2_{00}}{\pi^2 q^1_{00} + (1-\pi^2)q^2_{00}}(A_0 - V_0^{Q^1}) + \frac{q^1_{01}q^2_{01}}{\pi^2 q^1_{01} + (1-\pi^2)q^2_{01}}(A_1 - V_1^{Q^1}) - \beta^{-1}(A_0 - V_0^{Q^1}) \geq 0
$$

for all $\pi^2 \in [0,1]$. Thanks to (B.2), which I already proved, I can divide both sides by $A_0 - V_0^{Q^1}$. By (B.4), this gives

$$
\frac{q^1_{00}q^2_{00}}{\pi^2 q^1_{00} + (1-\pi^2)q^2_{00}} + \frac{q^1_{01}q^2_{01}}{\pi^2 q^1_{01} + (1-\pi^2)q^2_{01}}R_1 - \beta^{-1} \geq 0
$$

Denoting $\tilde{\pi} \equiv (1-\pi^2)$, and $q^i \equiv q^i_{00} = 1 - q^i_{01}$ this is equivalent to

$$\begin{aligned} &\tilde{\pi}^2 \cdot \beta^{-1}(q^2-q^1)^2 + \tilde{\pi}(q^2-q^1)[R_1(1-q^1)(1-q^2) - q^1q^2 + \beta^{-1}q^1 - \beta^{-1}(1-q^1)] \\ &+ q^1(1-q^1)[q^2 + (1-q^2)R_1 - \beta^{-1}] \geq 0 \end{aligned} \tag{B.6}$$

for all $\tilde{\pi} \in [0,1]$. We have to consider two cases, which determine R_1 in the proposition. The first is $q^1 > \beta q^2$, then $R_1 = \frac{1-\beta q^2_{00}}{\beta q^2_{01}}$. Plugging this back into (B.6) we can see that the constant term becomes zero. It is also easy to see that the linear term must me positive under our assumptions (remember we assumed $q^1_{00} < q^2_{00}$). This, together with the fact that the quadratic term coefficient is positive means that the quadratic expression must be nonnegative for $\tilde{\pi} \geq 0$. Hence we are done in this case.

Now assume $q^1 > \beta q^2$. Then $R_1 = \frac{(1-\sqrt{\beta q^1_{00} q^2_{00}})^2}{\beta q^1_{01} q^2_{01}}$. In this case it is a matter of doing a simple algebra to check that the determinant of the quadratic is zero, which, together with the fact that the quadratic term coefficient is zero implies that (B.6) holds for all $\tilde{\pi}$. So we are done with the case $d = 0$.

Now consider $d = 1$. We have

$$\begin{aligned} Tp^A(1,\pi^1,\pi^2) &\geq T^2p^A(1,\pi^1,\pi^2) \\ &\equiv \beta(\pi^2 q^1_{10} + (1-\pi^2)q^2_{10}) \left[(1-\lambda^{10}(\pi^1))V_0^{Q^2} + \lambda^{10}(\pi^2)V_0^{Q^1} + (\lambda^{10}(\pi^1) - \lambda^{10}(\pi^2))A_0\right] \\ &+ \beta(\pi^2 q^1_{11} + (1-\pi^2)q^2_{11}) \left[1 + (1-\lambda^{11}(\pi^1))V_1^{Q^2} + \lambda^{11}(\pi^2)V_1^{Q^1} + (\lambda^{11}(\pi^1) - \lambda^{11}(\pi^2))A_1\right] \end{aligned} \tag{B.7}$$

Using

$$\begin{aligned} \lambda^{10}(\pi) &= \frac{\pi q^1_{10}}{\pi q^1_{10} + (1-\pi)q^2_{10}} = 1 - \frac{(1-\pi)q^2_{10}}{\pi q^1_{10} + (1-\pi)q^2_{10}} \\ \lambda^{11}(\pi) &= \frac{\pi q^1_{11}}{\pi q^1_{11} + (1-\pi)q^2_{11}} = 1 - \frac{(1-\pi)q^2_{11}}{\pi q^1_{11} + (1-\pi)q^2_{11}} \end{aligned}$$

we can see that both $p^A(1,\pi^1,\pi^2)$ and $T^2p^A(1,\pi^1,\pi^2)$ are linear in π^2. Also note that $T^2p^A = p^A$ for $\pi^1 = \pi^2$. In order to conclude $Tp^A(0,\pi^1,\pi^2) \geq p^A(0,\pi^1,\pi^2)$ it is enough then to prove that $T^1p^A(1,\pi^1,\pi^2) \geq p^A(1,\pi^1,\pi^2)$ for $\pi^2 = 0$. Note that I introduced auxiliary notation T^2 ("2" refers to agent 2 not to iterating T twice).

We have:

$$
\begin{aligned}
&T^2p^A(1,\pi^1,0) - p^A(1,\pi^1,0) = \\
&=\beta q^2_{10}\left[\left(1-\frac{\pi^1 q^1_{10}}{\pi^1 q^1_{10}+(1-\pi^1)q^2_{10}}\right)V_0^{Q^2}+\frac{\pi^1 q^1_{10}}{\pi^1 q^1_{10}+(1-\pi^1)q^2_{10}}\cdot A_0\right] \\
&\quad+\beta q^2_{11}\left[1+(1-\frac{\pi^1 q^1_{11}}{\pi^1 q^1_{11}+(1-\pi^1)q^2_{11}})V_1^{Q^2}+\frac{\pi^1 q^1_{11}}{\pi^1 q^1_{11}+(1-\pi^1)q^2_{11}}\cdot A_1\right] \\
&\quad-\pi^1(A_1-V_1^{Q^2})-V_1^{Q^2} \\
&=\beta\pi^1\left[\frac{q^1_{10}q^2_{10}}{\pi^1 q^1_{10}+(1-\pi^1)q^2_{10}}(A_0-V_0^{Q^2})+\frac{q^1_{11}q^2_{11}}{\pi^1 q^1_{11}+(1-\pi^1)q^2_{11}}(A_1-V_1^{Q^2})-\beta^{-1}(A_1-V_1^{Q^2})\right]
\end{aligned}
$$

To show $T^2p^A(1,\pi^1,0)\geq p^A(1,\pi^1,0)$ for all $\pi^1\in[0,1]$ it is enough to prove:

$$
\frac{q^1_{10}q^2_{10}}{\pi^1 q^1_{10}+(1-\pi^1)q^2_{10}}(A_0-V_0^{Q^2})+\frac{q^1_{11}q^2_{11}}{\pi^1 q^1_{11}+(1-\pi^1)q^2_{11}}(A_1-V_1^{Q^2})-\beta^{-1}(A_1-V_1^{Q^2})\geq 0
$$

for all $\pi^1\in[0,1]$. The already proved inequality (B.2) allows me to divide both sides by $A_1-V_1^{Q^2}$. I get (using (B.4)) that it is enough to prove:

$$
\frac{q^1_{10}q^2_{10}}{\pi^1 q^1_{10}+(1-\pi^1)q^2_{10}}R_2+\frac{q^1_{11}q^2_{11}}{\pi^1 q^1_{11}+(1-\pi^1)q^2_{11}}-\beta^{-1}\geq 0
$$

Denoting $\bar{\pi}\equiv 1-\pi^1$, and $q^i\equiv q^i_{11}=1-q^i_{10}$ this is equivalent to

$$
\begin{aligned}
&\bar{\pi}^2\cdot\beta^{-1}(q^2-q^1)^2+\bar{\pi}(q^2-q^1)[R_2(1-q^1)(1-q^2)-q^1q^2+\beta^{-1}q^1-\beta^{-1}(1-q^1)] \\
&+q^1(1-q^1)[q^2+(1-q^2)R_2-\beta^{-1}]\geq 0 \qquad \text{(B.8)}
\end{aligned}
$$

for all $\bar{\pi}\in[0,1]$.

This looks the same as (B.6) and the proof goes exactly the same. We have to consider two cases, which determine R_2 in the proposition. The first is $q^1>\beta q^2$, then $R_2=\frac{1-\beta q^2_{11}}{\beta q^2_{10}}$. Plugging this back into (B.8) we can see that the constant term becomes zero. It is also easy to see that the linear term must me positive under our assumptions (remember we assumed $q^1_{11}<q^2_{11}$). This, together with the fact that the quadratic term coefficient is positive means that the quadratic expression must be nonnegative for $\bar{\pi}\geq 0$. Hence we are done in this case.

Now assume $q^1\geq\beta q^2$. Then $R_2=\frac{(1-\sqrt{\beta q^1_{11}q^2_{11}})^2}{\beta q^1_{10}q^2_{10}}$. In this case it is an easy calculation to check that the determinant of the quadratic is zero, which, together with the fact that the quadratic term coefficient is zero implies that (B.8) holds for all $\bar{\pi}$. So we are done with the case $d=1$. □

Proof of Lemma 6. Consider the Markov family of distributions associated with $\bar{Q}$, with $P^{\bar{Q}}_{\bar{\Delta}}$ denoting the probability distribution associated with the Markov chain Δ starting from $\Delta_0 = \bar{\Delta}$. Obviously the distribution of interest is the one with $\bar{\Delta} = (0, d_0)$ but we need the whole Markov family to use recursive techniques.

Define a sequence of optional stopping times:

$$\tau_1 \equiv \min\{t > 0 | \Delta_t = \Delta_0\}$$
$$\tau_{n+1} \equiv \min\{t > \tau_n | \Delta_t = \Delta_0\}$$

Ergodicity of the process Δ_t implies that each state is recurrent (cf. Billingsley (1986) Sec. 8). This mean the stopping times defined above are finite. Using Lemma 8.3 from Bilingsley (1986) we get that also their first moments are finite:

$$E\tau_n = \frac{1}{\nu_1^Q},$$

for each $n \in \mathbb{N}$. Now we consider a selection of process l, given by stopping times τ_n i.e. process $\bar{l}_n \equiv l_{\tau_n}$. Now we will show that this process has iid increments. It is known that any time-homogeneous Markov chain has the strong Markov property (cf. Paragraph 8.6.3 in Wentzel (1980)), which in the case of process Δ can be expressed as:

$$P^Q_{\bar{Delta}}(\theta_\tau^{-1} B | \mathcal{F}_\tau) = P^Q_{\Delta_\tau}(B) \qquad P_s\text{-almost surely on } \{\tau < \infty\}$$

for any $\bar{\Delta} \in D^2$, stopping time τ and event B measurable w.r.t the process Δ. $\theta_\tau : \{D^2\}^\infty \to \{D^2\}^\infty$ denotes the shift operator, i.e. $\theta_\tau(\Delta)_t \equiv \Delta_{t-\tau}$.

For any $m < n$, consider nth and mth increments of $\bar{l}$, $\Delta^{\bar{l}}_n \equiv \sum_{t=\tau_n}^{\tau_{n+1}-1} f(\Delta_t)$ and $\Delta^{\bar{l}}_m \equiv \sum_{t=\tau_m}^{\tau_{m+1}-1} f(\Delta_t)$, respectively. For any Borel set $A \subseteq \mathbb{R}$ we have: $\{\Delta^{\bar{l}}_n \in A\} = \theta^{-1}_{\tau_n}\{\Delta^{\bar{l}}_0 \in A\}$, and $\{\Delta^{\bar{l}}_m \in A\} = \theta^{-1}_{\tau_m}\{\Delta^{\bar{l}}_0 \in A\}$, hence the strong Markov property implies:

$$P^Q_{(0,d_0)}(\{\sum_{t=\tau_n}^{\tau_{n+1}-1} f(\Delta_t) \in A\} | \mathcal{F}_{\tau_n}) = P^Q_{(0,d_0)}(\{\sum_{t=\tau_m}^{\tau_{m+1}-1} f(\Delta_t) \in A\} | \mathcal{F}_{\tau_m}) = P^Q_{(0,d_0)}(\{\sum_{t=0}^{\tau_1-1} f(\Delta_t) \in A\})$$

The rightmost probability is unconditional hence it is a constant then also two other probabilities are constant. This implies that the nth increment is independent of $\mathcal{F}_{\tau_n}$ hence of all previous increments. Also the distribution is exactly the same as all the previous increments. By induction this lets us conclude that all the increments are iid.

This proves, that the process $\bar{l}$, which is some strictly increasing random selection from the process l is a random walk with iid increments. In order to prove it is recurrent we will use some standard results from the random walk theory. In order to apply those results we need to know that the increments have zero expected value. Let for each $D\bar{e}lta \in D^2$, $e_{\bar{\Delta}} = E^Q_{\bar{\Delta}}(\Delta^{\bar{l}}_0)$. The strong Markov property implies that e must satisfy the following recursive equation:

$$e = \begin{bmatrix} f(0,0) \\ f(0,1) \\ f(1,0) \\ f(1,1) \end{bmatrix} + \bar{Q} \begin{bmatrix} 0 \\ e_{(0,1)} \\ e_{(1,0)} \\ e_{(1,1)} \end{bmatrix}$$

Pre-multiplying both sides by the stationary distribution ν^Q we get:

$$\nu^Q e = \nu^Q \begin{bmatrix} f(0,0) \\ f(0,1) \\ f(1,0) \\ f(1,1) \end{bmatrix} + \nu^Q \begin{bmatrix} 0 \\ e_{(0,1)} \\ e_{(1,0)} \\ e_{(1,1)} \end{bmatrix} = \nu^Q \begin{bmatrix} 0 \\ e_{(0,1)} \\ e_{(1,0)} \\ e_{(1,1)} \end{bmatrix}$$

so $e_{(0,0)}$=0.

Now we can apply Theorem 3 and Theorem 4 from VI.10 of Feller (1966) to conclude that for any interval I the process $\bar{l}$ visits this interval infinitely often with probability 1, i.e. $P^Q(\bar{l}_n \in I \text{ infinitely often}) = 1$. □

Proof of Theorem 7. Let F denote the distribution of Δ_0. For any interval $I \subseteq \mathbb{R}$ denote

$$U(I) = E^Q \sum_{t=1}^{\infty} 1_I(\bar{l}_t) = \sum_{k=0}^{\infty} F^{k\bigstar}(\{I\})$$

Using the weak law of large numbers we have $P(|\bar{l}_n - \mu n| < \epsilon n) > \frac{1}{2}$ for $n > n_\epsilon$. It follows that $F^{k\bigstar}([-a,a]) > 1/2$ for $n_\epsilon < k < \frac{a}{\epsilon+\mu}$ so $a^{-1}U([-a,a]) > \frac{1}{2}(\frac{1}{\epsilon+\mu} - n_\epsilon a^{-1})$. It is a standard fact from the renewal theory (cf. Theorem 1 in VI.10 of Feller (1966)) that for any $a > 1$, $U([-a,a]) \leq (2a+1)U[-1,1])$. Using this we get

$$U([-1,1]) \geq \frac{1}{2a+1} U([-a,a]) \geq \frac{1}{3} a^{-1} U([-a,a]) > \frac{1}{6} (\frac{1}{\epsilon + \mu} - n_\epsilon a^{-1})$$

for any $a > 1$ and any $\epsilon > 0$, hence taking the limit $a \to \infty$, and $\epsilon \to 0$ we get

$$U([-1,1]) > \frac{1}{6\mu}$$

□