

Internal Rationality and Asset Prices

Klaus Adam
Mannheim University

Albert Marcet
Institut d'Anàlisi Econòmica CSIC

Juan Pablo Nicolini
Universidad Torcuato di Tella

October 13, 2008
- preliminary and incomplete -

Abstract

We show how standard learning rules can be interpreted as small departures from rationality in the context of an asset pricing model. We propose a distinction between ‘internal rationality’, as agents that maximize discounted expected utility under uncertainty given consistent beliefs about the future, and ‘external rationality’ as agents that know perfectly the true stochastic process for fundamentals (dividends) and market determined variables (asset prices). Naturally, this distinction is irrelevant with complete markets and homogeneous agents. However, we show that the required amount of information and computational ability for achieving external rationality is gigantic, once one allows for weak forms of heterogeneity and market incompleteness. We show how simple models of learning that satisfy internal rationality can be interpreted as small deviations from rationality.

JEL Class. No.: G12, G14, D83, D84

1 Motivation

An increasing number of papers employs models of learning about expectations to explain macroeconomic data and for policy analysis.¹ Such models imply a departure from the full rationality assumptions that are standard in macroeconomics nowadays, giving rise to the following related questions: to what extent are agents in models of learning actually behaving ‘irrationally’? Can we define a metric to measure how ‘small’ a deviation from rationality is? Are there different ‘degrees’ or different ‘aspects’ of rationality within these models? The

¹For example, Adam, Marcet and Nicolini (2008), Adam (2005), Chakraborty and Evans (2008), Cogley and Sargent (2008), Eusepi and Preston (2008), Marcet and Nicolini (2003), and Timmermann (1993, 1994, 1996) use models of learning to explain various observations; Evans and Honkapohja (2003a, 2003b, 2005), Molnar and Santoro (2007), Orphanides and Williams (2006) and Sargent (1999) employ models of learning for policy analysis.

objective of this paper is to develop a conceptual framework to address these questions.

We perform the analysis within the context of the simplest asset pricing model, close to the well known ‘tree’ model of Lucas (1978), in which agents are endowed with stocks that deliver an exogenously given sequence of dividends of the unique perishable good. In this model, agents’ decision problems depend on future dividends and future prices. From the point of view of the theorist, dividends are exogenous, and prices are endogenous - that is - market determined. On the other hand, agents are competitive, so they take price as *given*. Given agents’ belief about prices, agents make consumption and stock holding decisions. When markets are complete, all prices are assumed to be observable so in fact, agents do know the stochastic properties of prices - they know all the distributions! However, when markets are incomplete, agents do not observe all the state contingent prices. In making decisions, therefore, we must provide agents in the model with some beliefs regarding prices of securities that will be traded in the future and that cannot be traded now.

The model we analyze allows for incomplete markets and heterogeneity. Agents are endowed with some perception about prices which is not necessarily the correct one, but other than that agents are fully rational in that they use their information optimally, they know their utility function, and they act to maximize their utility. We show that, in order for the equilibrium to be the rational expectations one, we must provide the agents in the model the same information the theorist has. In particular, agents must know all relevant details of every other agent in the economy (like taste shocks or credit constraints), they must understand how markets function, they need to know which markets are closed and they must be able to compute very high-dimensional general equilibrium problems.

We therefore propose a distinction between ‘internal’ and ‘external’ rationality. Internal rationality is satisfied if agents maximize utility under uncertainty, given their constraints and given a consistent set of probability beliefs about variables that are external (exogenous) to their decision problem. However, those probability beliefs may not coincide - during a transitional period - with the true distribution of those variables as they emerge in equilibrium. We show how the microfoundations of an agent’s decision problem can be changed in a natural way to incorporate such imperfect market knowledge. Specifically, we argue that the probability space over which agents form expectations should include also market-determined variables such as prices. This departs from the standard formulation in dynamic stochastic models, where the probability space of agents’ beliefs is reduced - from the outset - to contain only states of nature of exogenous variables. In the case of the stock pricing model we analyze, the standard assumption would be to assume that dividends span the probability space over which agents formulate their beliefs and that stock prices are a function of dividends.² Instead, we assume that agents’ beliefs are formulated over

²This assumption is also made in the literature on ‘rational bubbles’, e.g., Santos and Woodford (1997).

the space of dividends and stock prices. To put it differently, the standard way to formulate agents' beliefs is to assume from the outset that agents perceive that the joint distribution of the history of prices and dividends up to period t has a singularity, while we relax this imposition. Allowing agents to attach a non-degenerate distribution between prices and dividends is a natural and (potentially) small departure from rational expectations.

While we relax rationality about external aspects of agents' environment, we maintain all remaining rationality assumptions imposed in a Bayesian REE, namely those pertaining to the 'internal' aspects of agents' decision making: we consider agents who know their infinite horizon objective function, know their budget constraint and other constraints, hold complete and well-defined beliefs over the infinite sequence of payoff-relevant events (possibly giving rise to Bayesian learning as a way to formulate the relevant conditional expectations), and that maximize utility under uncertainty given this knowledge. All agents are thus maximizing utility conditional on their beliefs and we refer to this aspect of rationality as 'internal rationality'. Naturally, the equilibrium satisfies "external" rationality, when, in addition, the belief system of agents is constrained to be the one the model delivers in equilibrium for all periods.

It is in order to effectively distinguish between knowledge about the own decision problem (assumed perfect) and knowledge about market outcomes and fundamentals (allowed to be imperfect), that we allow for heterogeneous agents and incomplete markets. In this setup an investor's knowledge of the own decision problem will not imply knowledge of the market outcomes, unlike in the standard cases of a representative agent or complete markets, where these two aspects of knowledge cannot be disentangled.

To illustrate the previous point, consider the case of risk neutral agents with discount factor δ , and let \mathcal{P}^i denote the complete set of beliefs of agent i over possible future paths of dividend *and* price realizations. With incomplete markets and heterogeneity, some agents may choose not to participate in the market in a given period. Letting m_t be the marginal investor in period t , the equilibrium asset price will satisfy this marginal investor's optimality condition

$$P_t = \delta E_t^{\mathcal{P}^{m_t}} (P_{t+1} + D_{t+1}) \quad (1)$$

where $E_t^{\mathcal{P}^{m_t}}$ denotes the expectations of the marginal investor at time t . But as there is heterogeneity and markets are not complete, unless agents are endowed with a lot of additional information about the market (external rationality), internal rationality *will not imply* that future price expectations are given by

$$E_t^{\mathcal{P}^{m_t}} P_{t+1} = E_t^{\mathcal{P}^{m_t}} \sum_{j=1}^{\infty} \delta^j D_{t+1+j} \quad (2)$$

As a result, agents' price expectations will not be restricted by their beliefs about the dividend process. The underlying reason for this finding is that under incomplete markets the identity of the marginal investor changes from period to period, so that one cannot appeal to the law of iterated expectations

to transit from the first order condition (1) to the discounted sum expression (2). Therefore, equilibrium prices will generally deviate from the discounted sum of dividends, even if all agents are internally rational. In this way fluctuations in investors' expectations about future prices can become an independent source of asset price fluctuations.

We argue that it is natural to study models where the belief system \mathcal{P}^i does not incorporate a singularity in the joint distribution of prices and dividends, by showing explicitly what additional information about market behavior the agents need to possess to derive a discounted sum of dividend expression for the asset price. We argue that standard formulations of Bayesian REE are in fact imposing such a singularity in the beliefs of investors. The existence of a singularity appears to be in stark contrast with what economists seem to know about the relation between current prices and the observed history of dividends: the empirical literature in asset prices has had a very hard time in detecting a stable mapping between dividends and prices in the data, even though many researchers have tried. For this reason we think it is interesting to consider models that do not impute such market knowledge (external rationality) to investors.

Allowing for beliefs that do not necessarily exhibit the singularity has a very interesting implication: the standard equation used to value assets in finance, where the price is equal to the present value of future dividends fails to be an equilibrium condition. On the other hand, 'internal' rationality implies that the price of an asset is the discounted expectation of next period's stock price and dividend. This is important since, as we show in Adam, Marcet, Nicolini (2008) (AMN), a totally standard model can easily match some basic empirical puzzles of asset prices that have been documented in the finance literature, as long as agents have beliefs regarding future dividends and prices and stock prices are given by (1). We show here that in the model of AMN agents are 'internally' rational. In addition, we show that a specific probability measure that is generated by Bayesian updating about the price and the dividend gives rise to the same learning equations employed by agents in AMN.

A standard way to relax the strong informational assumptions underlying rational expectations equilibria (REE) has been the concept of Bayesian REE. We point out two shortcomings of this approach.³ First, while these equilibria allow for imperfect information about fundamentals, they imply knowledge about market outcomes far beyond the one implied by internal rationality. This knowledge takes the form of the *singularity* in agents' joint distribution over market outcomes and fundamentals. Second, Bayesian REE models may give rise to very unstable model behavior. To show this we analyze the situation in which agents possess sufficient market knowledge to derive a discounted sum of dividend expression for the asset price and show that the asset price becomes extremely sensitive to fine details of the prior distribution about the dividend process incorporated in agents' beliefs.⁴ We interpret this as saying that in this

³Bray and Kreps (1987) have pointed out further shortcomings.

⁴Related but different findings of this kind have shown up in Geweke (2001), Weitzman

approach agents' prior beliefs matter much more for stock prices than economic factors. This issue did not show up in most earlier papers using discounted dividends and learning, as they deviated from full Bayesian behavior.

Thus, as a side benefit, this paper provides microfoundations for the asset pricing model with learning presented in Adam, Marcet, Nicolini (2008) who show that the empirically observed volatility of stock prices arises naturally in standard asset pricing models as soon as *i*) agents decide whether to purchase or sell stocks by comparing today's price to the discounted expected stock payoff next period, as implied by equation (1); *ii*) agents do not have fully rational price expectations, instead learn to form price expectations using past observed price data. This paper shows that the assumptions in AMN are fully consistent with individual utility maximization of an infinite horizon objective. It also shows how the expectations updating implied by the learning model studied in AMN can be derived from internally rational agents, who compute the relevant conditional expectations using a complete and entirely standard probability measure over prices and dividends. In this formulation it becomes clear how agents' beliefs in AMN represent a small deviation from full (external) rationality: while agents' beliefs over the dividend and price process do contain a singularity as in a REE, agents understand that dividends and prices are strongly related. It is in this sense, that we argue the deviations from rationality in AMN to be small.

This paper is related to a number of papers in the learning literature that attempt to construct a full set of beliefs over long-horizons. This approach was first used in Marcet and Sargent (1989) (example e. in section 4) and has recently been extended by Preston (2005) and Eusepi and Preston (2008) in applications to monetary models. The main difference to this literature is that the present paper considers probability beliefs that take the standard form of a probability measure over a stochastic process, therefore are dynamically consistent over time. Instead, the papers mentioned use the anticipated utility framework of Kreps (1998) and construct each period a new set of probability beliefs, but one that is almost surely inconsistent with the complete set of beliefs held in the previous period. By comparison we consider agents who hold a consistent set of beliefs and the conditional expectations they compute are dynamically consistent.

The setup in this paper is also indirectly related to the literature on rational beliefs initiated by Mordecai Kurz (see chapter 1 in Kurz (1997) for an overview; chapter 9 for an application to asset pricing). In this literature agents also formulate beliefs about fundamentals and market outcomes, but unlike in the current paper, agents' probability distributions are assumed to be shifting in response to the realization of an extrinsic 'generating sequence'. This generating sequence causes agents' conditional expectations to differ from the true conditional expectation even asymptotically. In the present paper, belief revisions are either induced by endogenous market outcomes (price observations) or intrinsic outcomes (dividend realizations), and we consider models where beliefs asymptotically converge to the true distribution, so that asymptotically

(2007) and Pesaran, Pettenuzzo and Timmermann (2007).

deviations from full (external & internal) rationality will vanish. This has the advantage of being able to study the implications of ‘small’ deviations from full rationality.

The outline of the paper is as follows. Section 2 presents a simple incomplete markets model, derives investors’ optimality conditions, and considers agents that hold expectations about future prices and dividends. It shows how internal rationality fails to restrict price expectations as a function of dividend beliefs and how Bayesian updating of price and dividend expectations can give rise to the conditional expectations used in AMN (2008). Section 3 then derives the information about the market required by agents to be able to equate the price to the expected discounted sum of dividends (EDSD). This section also presents a formal result about the strong sensitivity of the EDSD to prior information about the dividend process. A conclusion briefly summarizes.

2 Rationality with imperfect market knowledge

We consider a simple asset pricing model with heterogeneous agents and incomplete markets. All agents are infinitely-lived and risk neutral, but agents differ because they may discount future payoffs differently and/or hold different (prior) beliefs. Markets are incomplete because of the existence of constraints that limit the amount of stocks investors’ can buy or sell. The model reduces to the setting studied in AMN in the limit with vanishing investor heterogeneity (identical discount factors and prior beliefs).

The presence of investor heterogeneity and market incompleteness allows us to distinguish between the investors’ knowledge of their own decision problem and their knowledge about market-determined variables, i.e., asset prices, which are also influenced by the discount factors and beliefs of other (possibly different) investors. This distinction is important to differentiate the implications of internal rationality from those implied by the full rationality assumptions present in (Bayesian) rational expectations equilibrium.

2.1 Basic model

The economy has $t = 0, 1, 2, \dots$ periods and is populated by I infinitely-lived risk-neutral investor types. There is a unit mass of investors of each type, all of them initially endowed with one unit of an infinitely lived stock. Agents of type $i \in \{1, \dots, I\}$ have a standard time-separable utility function

$$E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} (\delta^i)^t C_t^i \tag{3}$$

where C_t^i denotes consumption and δ^i a type-specific discount factor. The operator $E_0^{\mathcal{P}^i}$ denotes the agent’s expectations in some probability space $(\Omega, \mathcal{S}, \mathcal{P}^i)$, where Ω is the space of realizations, \mathcal{S} the corresponding σ -Algebra, and \mathcal{P}^i

a subjective probability measure over (Ω, \mathcal{S}) . The probability measure is allowed to be type-specific. We let Ω^t denote the set of histories from period zero up to period t , and ω^t an element of Ω^t . A routine application of probability rules will often imply that expectations conditional on Ω^t may be generated by some Bayesian updating scheme. Further details of the probability space will be specified at a later stage.

Investors of type i choose consumption and stock holdings (C_t^i, S_t^i) for all t where

$$(C_t^i, S_t^i) : \Omega^t \rightarrow R^2 \quad (4)$$

As usual, period t choices are contingent on all information available up to time t . Expected utility (3) can thus be written as

$$E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} (\delta^i)^t C_t^i = \int_{\Omega} \sum_{t=0}^{\infty} (\delta^i)^t C_t^i(\omega^t) d\mathcal{P}^i(\omega). \quad (5)$$

As usual, \mathcal{P}^i is given to the agents, incorporating the usual assumption that agents are small and can not influence the aggregate variables.

The agent faces the usual budget constraint

$$C_t^i + P_t S_t^i \leq (P_t + D_t) S_{t-1}^i \quad (6)$$

which has to hold for all t and all $\omega^t \in \Omega^t$ with P_t denoting the stock price, D_t the stock's dividend payments, and S_t^i the agent's stock holdings at the end of period t . Initial stockholdings $S_{-1}^i = 1/I$. Dividends evolve according to an exogenous stochastic process. Agents' perceptions about the evolution of dividends may or may not coincide with the objective probabilities. Prices $P_t : \Omega^t \rightarrow R_+$ are taken as given by the agent.

Besides the budget constraint, consumers face the following limit constraints on stock holdings:

$$S_t^i \geq 0 \quad (7)$$

$$S_t^i \leq \bar{S} \quad (8)$$

where $1 < \bar{S} < \infty$. Constraint (7) is a standard short-selling constraint and often used in the literature. The second constraint (8) is a simplified form of a leverage constraint capturing the fact that the consumer cannot buy arbitrarily large amounts of stocks. Constraint (8) simplifies our equilibrium calculation in the presence of risk neutral investors.

We consider agents who choose (4) in order to maximize utility (5) subject to the budget constraint (6), the limit constraints (7) and (8), taking as given the probability measure \mathcal{P}^i . Such agents are called internally rational. Throughout the paper we assume that \mathcal{P}^i assigns zero probability to negative prices and negative dividends and is such that

$$E^{\mathcal{P}^i} [P_{t+1} + D_{t+1} | \omega^t] < \infty \text{ for all } \omega, t \quad (9)$$

We also assume that a maximum of the investor's utility maximization problem exists.⁵

The setup just described may seem completely standard, but there exists a crucial difference from the standard formulation: we have *not* assumed that the space of outcomes Ω consists of histories of the state of nature (dividends) only. As a matter of fact, we will include also other elements and this will be crucial in our discussion.

2.2 Optimality conditions

The first order optimality conditions then require one of the following conditions to hold at all periods t and for almost all realizations in $\omega^t \in \Omega^t$:

$$P_t < \delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1}) \quad \text{and } S_t^i = \bar{S} \quad (10a)$$

$$P_t = \delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1}) \quad \text{and } S_t^i \in [0, \bar{S}] \quad (10b)$$

$$P_t > \delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1}) \quad \text{and } S_t^i = 0 \quad (10c)$$

where $E_t^{\mathcal{P}^i}$ denotes the expectation conditional on ω^t computed with the measure \mathcal{P}^i . Since the objective function is concave and the feasible set is convex these equations determine necessary and sufficient conditions for the agent's optimal investment decisions.

Importantly, the optimality conditions are of the one-step-ahead form, i.e., they involve today's price and the expected price and dividend tomorrow. Therefore, to take optimal decisions the agent only needs to know whether the observed realization ω^t implies that the expected stock return is higher, equal or lower than the inverse of the own discount factor. Since agents can trade stocks in any period, the one-step-ahead optimality conditions (10) deliver optimal investment choices even if stocks can be held for an arbitrary number of periods. Just to emphasize, it is not true that an internally rational agent has to compare today's price with the discounted sum of dividends in order to act optimally! As we show below, only in special cases are the above first order conditions equivalent with a discounted sum formulation that involves only agents' expectations of future dividends and knowledge of the agent's own utility function.

2.3 Beliefs

We now give some more structure to the non-standard part in our formulation, namely, the formulation of beliefs. We consider agents who view the process for $\{P_t, D_t\}$ as exogenous to their decision problem. More precisely, investors' probability measure \mathcal{P}^i is defined over the probability space

$$\Omega \equiv \prod_{t=0}^{\infty} R_+^2$$

⁵Appendix A shows that the existence of a maximum can be guaranteed by bounding the utility function. For notational simplicity we treat the case with unbounded utility in the main text and assume existence of a maximum.

which contains all possible realizations of $\{(P_t, D_t)\}_{t=0}^{\infty}$. Letting \mathcal{S} denote the sigma-algebra of all Borel subsets of Ω , we assume that type i 's beliefs are given by a well defined probability measure \mathcal{P}^i . We denote by Ω_P and Ω_D the sets of whole histories of P and D , respectively, so that $\Omega = (\Omega_P, \Omega_D)'$. Similarly, $\Omega^t = (\Omega_P^t, \Omega_D^t)'$.⁶

Throughout the paper we will make the statement that our agents have "a consistent set of beliefs". By this we simply mean that $(\Omega, \mathcal{S}, \mathcal{P}^i)$ is a proper probability space, that \mathcal{P}^i satisfies all the axioms that a probability measure on a stochastic process should satisfy.

With this setup ω^t is the history of all past prices and dividends, that is, investors take decisions conditional on the history of observed price and dividend realizations. It seems natural to assume that investors see prices as given and condition their decisions on realizations of prices. Moreover, \mathcal{P}^i gives proper joint probabilities to all possible values of prices and dividends in any set of dates, thus defines a complete set beliefs about all payoff-relevant events. These beliefs allow to evaluate first order conditions (10) at any possible history Ω^t .⁷

2.4 Standard belief formulation: a singularity

The setup for beliefs defined in the previous section may seem completely ordinary but it differs from standard dynamic economic modelling practice, which imposes additional restrictions on beliefs. Specifically, the standard belief specification in the above model would be to consider just Ω_D as the underlying state space over which probabilities are formed. The probabilities for the price process are then constructed by endowing agents with the knowledge that each realization $\omega_D^t \in \Omega_D^t$ is associated with a given level of the stock price P_t , which amounts to endowing agents with the knowledge of the function

$$P_t : \Omega_D^t \rightarrow R_+ \tag{11}$$

that maps dividend realization into prices. Agents then do not need to condition their actions on observed prices, since these carry redundant information. The standard belief specification can thus be interpreted as a special case of the formulation outlined in the previous section, namely one where \mathcal{P}^i is assumed to impose a degeneracy between pairs (ω_P^t, ω_D^t) . The more general belief formulation, outlined in the previous section, allows agents to be uncertain about the relation between prices and dividends.

Importantly, the standard singularity imposed on agents' beliefs is not a consequence of internal rationality. Instead, it is the result of an equilibrium consistency requirement imposed by the concept of (Bayesian) Rational Expectations Equilibrium. With rational expectations, such degenerate beliefs are

⁶Note that this setup is consistent with our previous general formulation of prices as $P_t : \Omega^t \rightarrow R_+$.

⁷This is related to the formulation in Anderson and Sonnenschein (1985).

consistent with the rational expectations equilibrium outcome, so no loss of generality is implied by imposing the degeneracy in \mathcal{P}^i from the outset.

Assuming that agents know about the existence of an exact equilibrium mapping determining prices as a function of past dividends seems to be very restrictive at a first sight. Indeed, when studying actual asset price and dividend data, such a relationship remains fairly elusive. One possible interpretation of our extended probability space is thus that it endows investors with the same doubts about the relationship between prices and dividends that appears to be present among economists who have been studying the behavior of actual stock prices for years.

2.5 Equilibrium

We now consider the process for equilibrium prices with internally rational agents. The objective of this section is to show that with internal rationality, agents' beliefs about dividends - even when combined with knowledge of the equilibrium asset pricing equation - do not impose restrictions on agents' price beliefs, thus also fails to imply that agents know the mapping (11).

As usual, equilibrium prices are defined as a stochastic process that clear all markets. Since agents do not necessarily hold rational expectations, we need to distinguish equilibrium prices from agents expectations about prices. We denote the equilibrium price as \mathbf{P}_t and will keep P_t inside the expectations of agents, since this is the variable that is perceived by agents. As a first approximation, we will assume that \bar{S} in constraint (8) is large enough such that it never binds. Thus, from the first order conditions (10), it is clear that in equilibrium the asset is held by the agent type with the most optimistic beliefs about the discounted expected price and dividend in the next period, i.e., equilibrium prices have to satisfy:

$$\mathbf{P}_t = \max_i \delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1}) \quad (12)$$

Since the expectations $E_t^{\mathcal{P}^i}$ are conditional on the realization \mathbf{P}_t , it is unclear - at this level of generality - whether this equation always has a unique solution. At this point, we proceed by assuming existence and uniqueness. The next section specifies \mathcal{P}^i in detail, which allows us to provide necessary and sufficient conditions for this to be the case.

The existence of a unique equilibrium price implies that it is indeed a function of the history ω_D^t of dividend realizations only, i.e., $\mathbf{P}_t : \Omega_D^t \rightarrow R_+$. Equilibrium prices will, of course, depend on the fundamentals of the problem, which in this case include all agents' beliefs \mathcal{P}^i in addition to the standard fundamentals such as utility function parameters, discount factors and dividend processes. It may thus appear that internal rationality implies it to be optimal (rational) for agents to hold degenerate beliefs of the kind imposed in standard rational expectations models. This statement is correct indeed, but fails to answer: which precise degeneracy agents should impose on their beliefs \mathcal{P}^i ? As we show

next, even if agents knew formula (12), they cannot derive the exact degeneracy just from knowledge of own utility functions, dividend beliefs and knowledge of the equilibrium asset pricing equation.

Let m_t denote the marginal agent pricing the asset in period t :⁸

$$m_t = \arg \max_i \delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1})$$

Since m_t is determined in equilibrium, we have $m_t : \Omega_D^t \rightarrow \{1, \dots, I\}$. The equilibrium price (12) can thus be written as

$$\mathbf{P}_t = \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (P_{t+1} + D_{t+1}) \quad (13)$$

Suppose it is common knowledge to agents that the equilibrium price satisfies equation (13) each period.⁹ Common knowledge thereby means that each agent knows that the asset is priced according to (13) each period, that each agent knows that other agents know this to be the case, knows that other agents know that others know it to be true, and so on to infinity.¹⁰ The question we are posing is: would common knowledge of equation (13) allow internally rational agents to impose restrictions on price beliefs as a function of their beliefs about dividends? The answer turns out to be ‘no’. An internally rational agent could rationally hold price beliefs that are either larger or smaller than the own expectations of the discounted sum of dividends.

Common knowledge of (13) implies that agents know that the equilibrium price must satisfy the recursion

$$\mathbf{P}_t = \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (\mathbf{P}_{t+1} + D_{t+1}) \quad (14)$$

This allows agents to iterate on this equation so as to obtain an expression for the equilibrium price in terms of expected dividends and expectations of some terminal price:

$$\begin{aligned} \mathbf{P}_t &= \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (D_{t+1}) \\ &+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} D_{t+2} \right) \\ &+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+3} \right) \right) \\ &+ \dots \\ &+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\dots \delta^{m_{t+T}} E_{t+T}^{\mathcal{P}^{m_{t+T}}} (\mathbf{P}_{t+T+1} + D_{t+T+1}) \right) \right) \end{aligned} \quad (15)$$

⁸If m_t is non-unique we can use a selection criterion from among all marginal agents. For example, we can assume m_t to be the marginal agent with the lowest index i .

⁹Internally rational agents do not need to have such knowledge to behave optimally conditional on their beliefs.

¹⁰See Aumann (1976) for a formal definition.

The terms on the second to the last line on the right hand side of the previous expression provide an alternative formulation for agents' price expectations. It shows that agents' price expectations are implied by the beliefs about which agents are going to be marginal in the future and by the beliefs about what beliefs future marginal agents will hold about dividends and some terminal price. Since agents are not marginal in each period and can rationally believe other marginal agents to hold rather different beliefs, own beliefs about dividends fail to restrict the beliefs agents could rationally hold about prices! This shows that dividend beliefs, knowledge of (13), and internal rationality fail to imply a *specific* singularity in agents' price and dividend beliefs \mathcal{P}^i . Indeed, despite internal rationality the equilibrium asset price could be different from the expectations of the discounted sum of dividends $E_t^{\mathcal{P}^i} \left(\sum_{j=1}^{\infty} (\delta^i)^j D_{t+j} \right)$ held by any agent i in the economy, as is the case in AMN (2008).

The previous derivation clarifies why, in our setup, it is unlikely that agents will take decisions by comparing the stock price to a discounted sum of dividends. The discounted sum of dividends is obtained in a standard way by applying the law of iterated expectations on the right side of equation (15). Yet, this can only be done when all the conditional expectations are with respect to the same probability measure, i.e., if m_t is constant through time. Since m_t is random in our model the law of iterated expectations can not be applied and the discounted sum of dividends does not emerge. This will occur whenever \mathcal{P}^i assigns positive probability to the event that the agent may not be marginal at some point in the future.

The previous discussion shows that agents' price expectations could be interpreted as a summary statistic of agents' beliefs about the 'deep' fundamentals δ^{m_t} and \mathcal{P}^{m_t} . A Bayesian modelling approach would postulate beliefs about these 'deep' fundamentals and update them in the light of new dividend realizations. This would of course imply restrictions on the behavior of price expectations and thus on prices as a function of the dividend realizations. Correctly updating the beliefs about the fundamentals δ^{m_t} and \mathcal{P}^{m_t} in the light of new dividend realizations requires, however, a tremendous amount of knowledge about the market. We come back to this issue in section 2.7.¹¹

2.6 Internal Rationality with Bayesian Learning

This section constructs a specific probability measure \mathcal{P} that is generated from Bayesian updating about a price and dividend process and shows how Bayesian updating gives rise to the learning equations employed by agents in AMN (2008).

¹¹The issues in this subsection are related to, but different from, the so-called 'infinite-regress problem' addressed by Townsend (1983). The point of the present paper is that agents can coordinate on Townsend's Bayesian REE solution only if they have a lot of external knowledge about other agents' characteristics, and that it is not an implication of internal rationality. Marcet and Sargent (1989) showed conditions guaranteeing that Townsend's equilibrium can be learnt in the long-run via least squares learning.

We also make precise the statement in what sense agents' prior beliefs in AMN (2008) deviate little from those imposed in the REE.

For analytical simplicity we consider a model in which all agents are identical in terms of beliefs and discount factors. This should be interpreted as the limiting case of a model where agents actually do have different but very similar discount factors and beliefs. We thus continue to consider agents that use the one-period optimality conditions (10) to decide on stockholdings and do not impose the singularity in price beliefs implied by a Bayesian REE. One interpretation of this setting is that agents are uncertain about whether or not other agents do actually hold similar beliefs.

The true process for dividends is

$$\log \frac{D_t}{D_{t-1}} = \log a + \log \varepsilon_t \quad (16)$$

with $\log \varepsilon_t \sim N(0, \sigma^2)$, D_{-1} given. With risk neutral investors with discount factor δ , the REE asset price is given by

$$P_t = \frac{\delta a E[e^{\log \varepsilon_t}]}{1 - \delta a E[e^{\log \varepsilon_t}]} D_t$$

so that the equilibrium process for the asset price follows

$$\log \frac{P_t}{P_{t-1}} = \log a + \log \varepsilon_t$$

Note that prices grow at the same rate as dividends and that the innovation in the price process is the same as in the dividend process. While it is well known that these aspects of the REE solution are empirically unappealing, the earlier discussion about the singularity in the joint distribution of prices and dividends suggests that they may be equally unappealing on theoretical grounds.

We consider agents whose perceived processes for prices and dividends encompasses the REE equilibrium, but do not impose some of the above-mentioned special features of the RE solution. More precisely, we assume that for a given value of the parameters $(\log \beta^P, \log \beta^D, \Sigma)$ agents perceptions satisfy

$$\begin{bmatrix} \log P_t/P_{t-1} \\ \log D_t/D_{t-1} \end{bmatrix} = \begin{bmatrix} \log \beta^P \\ \log \beta^D \end{bmatrix} + \begin{bmatrix} \log \varepsilon_t^P \\ \log \varepsilon_t^D \end{bmatrix} \quad (17)$$

given P_{-1}, D_{-1} , with

$$\begin{aligned} (\log \varepsilon_t^P, \log \varepsilon_t^D)' &\sim N(0, \Sigma) \\ \Sigma &= \begin{bmatrix} \sigma_P^2 & \sigma_{PD} \\ \sigma_{PD} & \sigma_D^2 \end{bmatrix} \end{aligned}$$

The previous specification allows prices and dividends to grow at different rates and innovations to prices and dividends to be only imperfectly correlated. Agents are uncertain about the mean growth rates of prices ($\log \beta^P$) and dividends ($\log \beta^D$) and about the covariance matrix of innovations (Σ). Agent's beliefs about these parameters are summarized by a distribution

$$(\log \beta^P, \log \beta^D, \Sigma) \sim f$$

We refer to f as the ‘prior’ distribution. Note that the prior together with the laws of motion (17) fully determine agents’ probability measure \mathcal{P} over future infinite sequences of price and dividends realizations.¹² We have thus specified the microfoundations of the model and of beliefs.

In what follows, we assume that f is of the Normal-Wishart conjugate form:

$$H \sim W(S_0, n_0) \tag{18a}$$

$$(\log \beta^P, \log \beta^D)' \Big| H \sim N \left((\log \beta_0^P, \log \beta_0^D)', (\nu_0 H)^{-1} \right) \tag{18b}$$

for given parameters $\log \beta_0^P, \log \beta_0^D, \nu_0, S_0$ and n_0 . The Wishart distribution W with precision matrix S_0^{-1} and $n_0 > 1$ degrees of freedom specifies agents’ marginal prior distribution about the inverse of the variance covariance matrix of innovations (Σ^{-1}), where n_0 scales the precision of prior beliefs. The normal distribution N specifies agents’ priors about the parameters $(\log \beta^P, \log \beta^D)$ conditional on the precision matrix H , where $(\log \beta_0^P, \log \beta_0^D)$ denotes the conditional prior mean and $\nu_0 > 0$ scales the precision of prior beliefs about $(\log \beta^P, \log \beta^D)$.

The previous specification nests rational expectations a special case. With RE agents know with certainty that $\beta^P = \beta^D = a$ and that $\sigma_P^2 = \sigma_D^2 = \sigma_{PD} = \sigma^2$. We call such a prior the ‘RE prior’. It requires centering beliefs at the RE outcome

$$\begin{aligned} (\beta_0^P, \beta_0^D) &= (a, a) \\ S_0 &= \sigma^2 \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \end{aligned}$$

and considering the limiting case of vanishing prior uncertainty:

$$\begin{aligned} n_0 &\rightarrow \infty \\ \nu_0 &\rightarrow \infty \end{aligned}$$

As we show below, the learning setup in AMN (2008) can be interpreted as a relaxation of this RE prior in the sense of allowing for non-vanishing prior uncertainty, i.e., for $n_0 < \infty$ and $\nu_0 < \infty$.

¹²Probabilities can be obtained mechanically as follows: for any Borel subset $s \subset \mathcal{S}$, determine the probability of s for any given value of $(\log \beta^P, \log \beta^D, \Sigma)$ using standard methods for Markov processes applied to equation (17), then integrate over these probabilities according to f_P .

Evaluating the first order conditions (10) based on \mathcal{P} requires computing the conditional price and dividend expectations, given information available up to period t . This is what we will deal with next.

Since the prior (18) is conjugate, agents' posterior beliefs at time t will have the same functional form as (18), but with location parameters $(\log \beta_t^P, \log \beta_t^D, \nu_t, S_t, n_t)$ instead of $(\log \beta_0^P, \log \beta_0^D, \nu_0, S_0, n_0)$. Defining the 'forecast error'

$$e_t = \begin{pmatrix} \log \frac{P_t}{P_{t-1}} - \log \beta_t^P \\ \log \frac{D_t}{D_{t-1}} - \log \beta_t^D \end{pmatrix}$$

it follows from chapter 9 in DeGroot (1970) that the location parameters evolve according to

$$\begin{pmatrix} \log \beta_{t+1}^P \\ \log \beta_{t+1}^D \end{pmatrix} = \begin{pmatrix} \log \beta_t^P \\ \log \beta_t^D \end{pmatrix} + \frac{1}{v_t + 1} e_t \quad (19a)$$

$$v_{t+1} = v_t + 1 \quad (19b)$$

$$S_{t+1}^{-1} = S_t^{-1} + \frac{n_t}{n_t + 1} e_t e_t' \quad (19c)$$

$$n_{t+1} = n_t + 1 \quad (19d)$$

The posterior mean for $(\log \beta^P, \log \beta^D)$ after integrating out uncertainty about the precision matrix and for given information up to period t is:¹³

$$E_t^{\mathcal{P}} [(\log \beta^P, \log \beta^D)] = (\log \beta_t^P, \log \beta_t^D)$$

Combining this result with (19a) shows that the dynamics of agents' one-step ahead conditional expectations of price and dividend growth evolve according to

$$E_{t+1}^{\mathcal{P}} \left[\begin{pmatrix} \log \beta_{t+1}^P \\ \log \beta_{t+1}^D \end{pmatrix} \right] = E_t^{\mathcal{P}} \left(\begin{pmatrix} \log \beta_t^P \\ \log \beta_t^D \end{pmatrix} \right) + \frac{1}{v_t + 1} e_t$$

$$v_{t+1} = v_t + 1$$

Except for a minor detail, this is the updating scheme for conditional expectations used in AMN (2008).¹⁴ We have thus shown how the pricing implications in AMN arise from a model populated with internally rational agents that hold a complete and consistent set of probability beliefs. Moreover, because the initial priors in AMN satisfy $\beta_0^P = \beta_0^D = a$, the difference with respect to the RE prior consists solely of allowing for non-vanishing prior uncertainty, i.e., $v_0 < \infty$.¹⁵

¹³This follows from the fact that the marginal posterior for price and dividend growth is t -distributed with $n_t - 1$ degrees of freedom, location vector $(\log \beta_t^P, \log \beta_t^D)'$ and precision matrix $v_t (n_t - 1) S_t$, see chapter 9 in DeGroot (1970).

¹⁴The equation above requires using observed log growth rates to update expectations of the log of the growth rate, while AMN (2008) used the level of observed growth rates to update the expectations of the level of the growth rate.

¹⁵Note that in this setup the parameters (S_0, n_0) do not influence the evolution of the posterior means of $(\log \beta^P, \log \beta^D)$, thus need not necessarily be centered at the RE prior.

Before discussing in the next section how this setup deviates from a Bayesian REE, we wish to briefly address the issue of equilibrium existence. Defining expected price and dividend growth, respectively, as

$$\begin{aligned}\beta_t &= E_t^{\mathcal{P}} \left[e^{\log \beta_t^P} e^{\log \varepsilon_t^P} \right] \\ a_t &= E_t^{\mathcal{P}} \left[e^{\log \beta_t^D} e^{\log \varepsilon_t^D} \right]\end{aligned}$$

the first order condition of the marginal agent (14) with the beliefs \mathcal{P} delivers

$$\begin{aligned}\mathbf{P}_t &= \delta E_t^{\mathcal{P}} (P_{t+1} + D_{t+1}) \\ &= \delta \beta_t \mathbf{P}_t + \delta a_t D_t\end{aligned}$$

so that for $\beta_t < \delta^{-1}$ there exists a unique equilibrium price given by

$$\mathbf{P}_t = \frac{\delta a_t}{1 - \beta_t \delta} D_t \quad (20)$$

Existence thus requires that the posterior *mean* for expected price growth β_t remains below δ^{-1} . In AMN (2008) this condition was insured by imposing a projection facility on beliefs. While this facility constitutes a deviation from Bayesian learning, it operates only in a few periods and asset price dynamics - averaged across many periods - turned out not to be too sensitive to the precise value chosen as upper bound. Of course, in periods where beliefs are at or near the projection facility the equilibrium prices are influenced by details of how the projection facility is imposed, but again this happens in only in a few periods. One could formally incorporate a projection facility by choosing a prior f that imposes that the outcome $\beta^P \geq \delta^{-1}$ is actually less likely than suggested by the conjugate Normal-Wishart formulation. This would result in a modification of the above formulae and is likely to lead to Bayesian updating behavior very similar to that implied by the projection facility imposed in AMN. We do not pursue this line further at this point.

2.7 How this departs from Bayesian RE

The previous section considered agents that maximized utility given their best guess about the evolution of prices and dividends and that used standard rules of probability to update their beliefs about the laws of motion for P and D . For non-vanishing prior uncertainty ($n_t < \infty$ and $v_t < \infty$) this setup nevertheless departs from Bayesian rational expectations behavior. This is so because agents assume the economy to evolve according to equations (17), which imply that their perceived likelihood for prices conditional on past information is given by

$$l(\log P_t \mid \Omega^{t-1}; \log \beta^P, \sigma_P^2) \sim N(\log \beta_{t-1}^P + \log P_{t-1}, \sigma_P^2) \quad (21)$$

Agents thus believe that asset prices follow a random walk with unknown drift and unknown variance. The true likelihood, however, will differ from random walk behavior because prices depend on all past realized values of dividends

and not only on the last price, see the discussion following equation (12). While the random walk behavior implied by (21) will become true asymptotically - this is shown in AMN (2008) - lagged dividends do influence the likelihood along the transition path. Therefore, for $n_t < \infty$ and $v_t < \infty$ agents' price forecasts will not be externally rational from the start, but will become so only asymptotically.

Would it be easy for agents to detect that the stochastic process for prices underlying \mathcal{P} is at odds with the actual properties of the price data? If the model is close to replicating the actual process for prices and dividends in the data, then it appears rather challenging to be very certain about the exact likelihood, i.e., about the exact dependence of current prices on the past history of dividends (the location of the singularity). Indeed, given that the asset pricing literature has been struggling for decades to link prices to dividends, it seems of interest to consider agents who face a similar struggle and who are not one hundred percent sure about this link, thus not fully externally rational.

3 Special Cases with Discounted Dividends

We consider in this paper agents who are internally rational but not externally rational. We now show what kind of additional market knowledge agents need to possess to obtain a Bayesian Rational Expectations Equilibrium. We will show that sufficient knowledge about the market insures that agents value the asset according to the discounted sum of dividends, implying no independent role for agents' price beliefs.

We then go on to point out some limitations of the Bayesian REE concept. We demonstrate that the expected discounted sum of dividends proves to be extremely sensitive to fine details in the specification of agents' prior beliefs about the dividend process. Indeed, prior information about mean dividend growth is so important for equilibrium stock prices in each period that it appears that prior beliefs drive the asset price entirely and that there is little room for economic explanations of the asset price. We think of this as a rather unattractive feature of the Bayesian REE concept.

3.1 Deriving discounted dividend expressions

This section shows which assumptions - beyond internal rationality - are required to arrive at an asset pricing formula equating the asset price to the expected discounted sum of dividends.

The starting point of our analysis are the necessary and sufficient conditions for optimality implied by internal rationality, i.e., equations (10). To be able to substitute out the price expectations showing up in agents' first order conditions, one needs:

Assumption 1 It is common knowledge that equation (14) holds for all t and all $\omega_D \in \Omega_D$.

Assumption 1 provides agents with information about how the *market* prices the asset for all periods t and all states ω_D . This information allows agents to iterate on the equilibrium asset price (14) and to express it as a function of future dividends and some terminal price and, therefore, to obtain equation (15). Importantly, agents can not iterate on their *own* first order optimality conditions, as these do not hold with equality always. Therefore, they actually need to know the equilibrium relationship (14) which holds with equality in all periods and all contingencies, but also involves a variety of information about other agents.

The discounted sum expressions (15) still involves expectations about the terminal equilibrium price \mathbf{P}_{t+T} . To eliminate price expectations altogether, one thus needs to impose that all agents know that the equilibrium asset price satisfies a ‘no-rational-bubble’ requirement:

Assumption 2 It is common knowledge that

$$\lim_{T \rightarrow \infty} \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\dots \delta^{m_{t+T}} E_{t+T}^{\mathcal{P}^{m_{t+T}}} (\mathbf{P}_{t+T}) \right) \right) = 0$$

for all t and all $\omega_D \in \Omega_D$.

Assumption 2 again provides information about the market: all agents know that marginal agents expect future marginal agents to expect (and so on to infinity) that prices grow at a rate less than the corresponding discount factors. In the case with homogeneous expectations and discount factors this requirement reduces to the familiar condition

$$\lim_{T \rightarrow \infty} E_t^{\mathcal{P}} (\delta^T \mathbf{P}_{t+T}) = 0 \tag{22}$$

As in the general case with heterogeneous expectations, this more familiar ‘no-rational-bubble’ condition endows agents with knowledge of how the *market* prices the asset asymptotically, as equation (22) restricts the behavior of the equilibrium price \mathbf{P} .

Assumption 2 allows to take the limit $T \rightarrow \infty$ in equation (15) and to abstract from expectations about the terminal selling price. One thus obtains an expression for the asset price in terms of the expected discounted sum of marginal agents’ expectations of future marginal agents dividend expectations, etc.. Agents may, however, hold rather different views about who will be marginal in the future and what the expectations of such marginal agents are going to be. Therefore, agents might still not agree on what equilibrium price should be associated with any $\omega_D \in \Omega_D$. To put it differently, agents could still hold very different beliefs about the function $P_t : \Omega_D^t \rightarrow R$, i.e., fail to hold the kind of rational beliefs about the price process they are assumed to hold in a Bayesian REE.

Iterating forward on equation (15), shifting forward one period, and taking conditional expectations with respect to agent i 's probability measure \mathcal{P}^i conditional on period t , shows that agent i 's price expectations are given by

$$\begin{aligned}
E_t^{\mathcal{P}^i} \mathbf{P}_{t+1} &= E_t^{\mathcal{P}^i} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} (D_{t+2}) \right) \\
&+ E_t^{\mathcal{P}^i} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+3} \right) \right) \\
&+ E_t^{\mathcal{P}^i} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} \left(\delta^{m_{t+3}} E_{t+3}^{\mathcal{P}^{m_{t+3}}} D_{t+4} \right) \right) \right) \\
&+ \dots
\end{aligned} \tag{23}$$

The agent's price expectations are thus implied by their beliefs about the process $\{m_t\}$ and their beliefs about $\{\delta^j, \mathcal{P}^j\}_{j \neq i}$ where \mathcal{P}^j denotes agent type j 's probability measure over infinite histories of *dividend* realizations. In a Bayesian rational expectations equilibrium agents may have imperfect information about the process $\{m_t\}$ or about $\{\delta^j, \mathcal{P}^j\}_{j \neq i}$, but know that in equilibrium these fundamentals are functions of the history of dividend realizations, i.e.,

$$m_t : \Omega_D^t \rightarrow \{1, 2, \dots, I\} \tag{24}$$

$$\mathcal{P}^j : \Omega_D \rightarrow [0, 1] \tag{25}$$

$$\delta^j : \Omega_D \rightarrow [0, 1] \tag{26}$$

where δ^j takes on the same value for all $\omega_D \in \Omega_D$. Therefore, in a Bayesian REE agents can afford to form beliefs about the dividend process only, as equations (24)-(26) then provide them with the implied beliefs about which agent is marginal (m_t), the beliefs about other agent's beliefs (\mathcal{P}^j) and other agent's discount factors (δ^j). Combining these implied beliefs with equation (23) then determines agents' price expectations as a function of the history of dividend realizations.

In a Bayesian REE the resulting price expectations have to be rational, i.e., objectively true given the dividend history. Therefore, it must be the case that the functions (24)-(26) used by agents to derive their beliefs about $\{m_t\}$, \mathcal{P}^j and δ^j are the ones that are objectively true in equilibrium and perfectly known to agents, i.e., we need:

Assumption 3 The function m_t , the discount factor δ^{m_t} and the probability distributions \mathcal{P}^{m_t} are common knowledge, for all t and all $\omega_D \in \Omega_D$.

The functions (24)-(26) incorporate a tremendous amount of knowledge about the market: for each possible dividend history they inform agents about which agent is marginal, the marginal agent's discount factor, and the marginal agent's belief system. This raises the important question of how agents could have possibly acquired such detailed knowledge about the working of the market already at period zero?

Assumptions 1-3 together imply that all agents are able to impose the following singularity on their joint beliefs about prices and dividends:

$$\begin{aligned}
\mathbf{P}_t &= \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (D_{t+1}) \\
&+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} D_{t+1} \right) \\
&+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+2} \right) \right) \\
&+ \dots
\end{aligned} \tag{27}$$

Moreover, this singularity proves to be correct in equilibrium.

The simplest and most common way in the literature to impose assumptions 1-3 in the literature is to consider the leading asset pricing example, i.e., the complete markets representative agent model with price beliefs that satisfy the no rational bubble requirement (22). The representative agent is marginal at all times and contingencies, so his own FOC holds with equality in all periods. Such an agent can simply iterate on the *own* first order optimality conditions and evaluate future expectations by applying the law of iterated expectations to own beliefs. Except for the issue of beliefs about limiting price growth, internal rationality thus implies *in this case* equality between equilibrium asset price and the discounted sum of dividends. The leading asset price example may thus erroneously suggest that this equality is the result of internally rational investment behavior on the side of agents, but as we have just shown this fails to be the case in more general settings with heterogeneous agents and incomplete markets. Agents that are not marginal in all times and all contingencies cannot iterate on their own first order conditions. Deriving an expression for the asset price in terms of discounted dividends then requires a tremendous amount of additional information about the market (Assumptions 1-3). Given that the equilibrium price does not even come close to revealing the underlying process for market fundamentals (m_t , δ^{m_t} and \mathcal{P}^{m_t}), it is hard to see how an agent could possibly be certain from the outset about how these fundamentals relate to the dividend process.¹⁶

3.2 Sensitivity of Bayesian RE asset prices

We now consider a representative agent model with risk-neutrality and complete markets. We impose a no-rational-bubble assumption on agents' price beliefs. As explained in the previous section, in this special setting individual rationality implies that prices equal the expected discounted sum of dividends. Specifically, letting δ denote the agent's discount factor and \mathcal{P} the agent's beliefs about the

¹⁶As we mentioned in footnote 11, this is related to, but different from, the issue of solving the "infinite regress" problem of Townsend (1983). Even though agents in a Bayesian REE have expectations that solve the infinite regress problem, this does not mean that internal rationality on the part of the agents is sufficient to derive which expectations would solve this problem.

dividend process (Ω_D), if it follows from equation (27) that the singularity $P_t : \Omega_D^t \rightarrow R$ in agents' joint beliefs about prices and dividends is given by

$$P_t = E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| \omega_D^t \right) \quad (28)$$

The goal of this section is to show that the expected discounted sum of dividends is extremely sensitive to the prior information about the dividend process incorporated in \mathcal{P} .¹⁷

The process for dividends evolves according to

$$D_t = aD_{t-1}\eta_t \quad (29)$$

where $\eta_t > 0$ is i.i.d. with $E[\eta_t] = 1$. We consider an agent that knows the distribution of innovations η_t but is uncertain about the true mean dividend growth rate $a > 0$. The agent's beliefs about a in period t are summarized by a posterior density $Post_t(\tilde{a})$ indicating the density assigned to $a = \tilde{a}$ given the history Ω_D^t . The prior information about dividends is given by $Post_{-1}(\cdot)$, i.e., by the beliefs prior to observing any dividend data. We assume that $Post_{-1}(\tilde{a}) = 0$ for $\tilde{a} < 0$, i.e., the agent assigns zero probability to dividends being negative.

For a given *known* dividend growth rate a the discounted sum expression (28) implies

$$P_t = \frac{\delta a}{1 - \delta a} D_t$$

The following proposition shows that once mean dividend growth is unknown, asset prices in a Bayesian RE model asset turn out to be extremely sensitive to slight changes in prior beliefs.¹⁸ The proof of the proposition can be found in appendix B.

Proposition 1 *Let B be the (possibly infinite) upper bound of the support of $Post_t(\cdot)$*

1. *If $B \geq \delta^{-1}$, then*

$$E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| \omega_D^t \right) = \infty$$

¹⁷The results regarding the sensitivity of the sum (28) is complementary to results derived in Geweke (2001), Pesaran, Pettenuzzo and Timmermann (2007), or Weitzman (2007), which require strictly positive risk aversion. Also, these authors consider the case of unknown dividend variance, while we consider the case of unknown mean dividend growth. Furthermore, our results apply to general prior functional forms, thus do not require conjugate prior specifications.

¹⁸This feature failed to show up in Pastor and Veronesi (2003). When studying the asset pricing implications of unknown earnings growth, these authors assumed the existence of a finite asset price at some terminal date $T < \infty$.

2. If $B < \delta^{-1}$ then

$$\begin{aligned} E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| \omega_D^t \right) &= D_t E_{Post_t} \left(\frac{\delta a}{1 - \delta a} \right) \\ &= D_t \int \frac{\delta \tilde{a}}{1 - \delta \tilde{a}} Post_t(\tilde{a}) d\tilde{a} < \infty \end{aligned}$$

3. Consider a family of posteriors $Post_t^k$ with upper bound for the support $B_k < \delta^{-1}$ that is converging to δ^{-1} as $k \rightarrow \infty$. Suppose $Post_t^k(B^k) > \gamma > 0$ and $Post_t^k(\cdot)$ is continuous from the left at B_k for all k , then

$$\lim_{k \rightarrow \infty} E^{\mathcal{P}^k} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| \omega_D^t \right) = \infty$$

The previous proposition shows that the asset price in Bayesian RE models depends almost exclusively on the specification of the upper bound of prior beliefs. If agents assign arbitrarily small but positive probability to dividend growth rates being larger than the inverse of the discount factor (B slightly larger than δ^{-1}) the asset price will be infinite. Moreover, if the distribution of the dividend growth innovations η has sufficiently large support, then agents will continue to assign positive probability to such events even after having observed an arbitrarily large amount of data (dividend growth realizations). The price will thus remain infinite forever.

The second part of the proposition proves that it is possible to obtain finite price levels by bounding the *support* of prior beliefs away from δ^{-1} . This insures that all posteriors are bounded in a similar way, so that the price remains finite forever. Yet, the final claim in the proposition illustrates that the asset price still remains dominated in *all* periods by the precise value chosen as upper bound. In particular, choosing a bound sufficiently close to δ^{-1} gives rise to arbitrarily high asset prices for *all* periods.

This allows us to conclude that in a Bayesian RE equilibrium the economics about stock prices do not matter nearly as much as the exact upper bound on prior beliefs.

The sensitivity of asset prices in a Bayesian RE has failed to show up in large part of the literature on Bayesian learning due to the use of a well-acknowledged shortcut.¹⁹ Instead of using the asset price equations in proposition 1, the following formula has typically been employed:

$$\mathbf{P}_t = D_t \frac{\delta E_{Post_t}(a)}{1 - \delta E_{Post_t}(a)} \quad (30)$$

¹⁹See, for example, Timmermann (1993, 1996), Brennan and Xia (2001), Cogley and Sargent (2008).

where

$$E_{Post_t}(a) = \int_0^\infty \tilde{a} Post_t(\tilde{a}) d\tilde{a}$$

In other words, standard practice has been to compute a discounted sum by assuming that dividends grow with certainty as suggested by the posterior mean of a . This differs notably from the correct expression which requires averaging over values of the discounted sum for all growth rates to which beliefs assign positive probability. This seemingly minor detail makes a huge difference for results: with equation (30) the issue of the prior *support* appears to be irrelevant, all that is required is that the prior *mean* is bounded away from δ^{-1} . The Bayesian literature introduced restrictions on belief updating (so-called projection facilities) that insured that this condition is satisfied each period and showed that the restrictions had to be applied only in a few periods. For the correctly computed stock price, however, the truncation of beliefs dominates the stock prices in all periods!

We now briefly discuss the issue of sensitivity of asset prices to prior beliefs in a setting with internally rational agents that cannot derive a discounted sum of dividend expression. Using the Bayesian learning setup from section 2.6, the pricing equation (20) shows that as long as the posterior *mean* about price growth belief β_t remains bounded away from δ^{-1} , prices remain well behaved. Importantly, with internal rationality the *support* of agents' beliefs about dividend or price growth is inessential for equilibrium asset prices so that asset prices do not depend in strong ways on aspects of beliefs about which economists generally possess little information.

3.3 When agents believe in discounted dividends

We now consider the price growth expectations implied by the discounted sum of dividend formula and show that these price growth expectations can take a rather counter-intuitive form. Internally rational agents can impose more appealing prior distributions on price growth and we show how less tight prior beliefs about price growth will allow these agents to outperform Bayesian agents with 'reasonable' subjective prior distributions for dividend growth.

Consider an economy with a representative Bayesian agent evaluating the asset according to the expected discounted sum of dividends. Dividends evolve according to (29) and the agent is uncertain about the true value of dividend growth. The agent's initial belief about mean dividend growth is captured by their prior belief distribution $Post_{-1}(\bullet)$ which has an upper bound on the support given by $B < \delta^{-1}$. From proposition 1 follows that equilibrium asset prices in this economy are given by

$$\mathbf{P}_t = D_t E_{Post_t} \left(\frac{\delta a}{1 - \delta a} \right) < \infty \quad (31)$$

Suppose that $Post_{-1}(\bullet)$ is centered around the true value a - so that agents have an ‘unbiased’ prior estimate of dividend growth, but that agents are uncertain about the true value of dividend growth and assign some (arbitrarily small) probability mass to dividend growth rates different from the true value a . Due to the convexity of $\delta a/(1 - \delta a)$ it follows from equation (31) that the asset is ‘overvalued’ in the sense that

$$\frac{\mathbf{P}_0}{D_0} > \frac{\delta a}{1 - \delta a}$$

where the right-hand side denotes the fundamental price dividend ratio under perfect information about dividend growth. As is clear from proposition 1, the amount of ‘overvaluation’ depends largely on how close the upper bound B is located to δ^{-1} .

Over time, Bayesian agents will learn the truth so that asymptotically one obtains

$$\lim_{t \rightarrow \infty} \frac{\mathbf{P}_t}{D_t} = \frac{\delta a}{1 - \delta a}$$

This together with the initial overvaluation implies that on average along the convergence process prices will grow at a rate *below* that of dividends.

What does a Bayesian agent expect prices to behave like in this economy? Price growth expectations can be derived from the first order condition (14), which together with equation (31) delivers

$$E_{Post_t} \left[\frac{\mathbf{P}_{t+1}}{\mathbf{P}_t} \right] = \frac{1}{\delta} - \frac{E_{Post_t} [a]}{E_{Post_t} \left(\frac{\delta a}{1 - \delta a} \right)} \quad (32)$$

Using again the convexity of $\delta a/(1 - \delta a)$, one obtains

$$E_{Post_t} \left[\frac{\mathbf{P}_{t+1}}{\mathbf{P}_t} \right] > \frac{1}{\delta} - \frac{E_{Post_t} [a]}{\frac{\delta E_{Post_t} [a]}{1 - \delta E_{Post_t} [a]}} = E_{Post_t} [a]$$

A Bayesian agent thus expects prices to grow *faster* than dividends. Moreover, from proposition (1) follows that $E_{Post_t} \left(\frac{\delta a}{1 - \delta a} \right) \rightarrow \infty$ as $B \rightarrow \delta^{-1}$, so that equation (32) implies that a Bayesian agent can expect prices to grow at a rate very close to the inverse of the discount factor for *any* amount of time! Yet, along the convergence process the opposite is true: prices will grow at a rate less than dividends and paradoxically the shortfall of price growth compared to dividend growth will be larger the closer B is located to δ^{-1} . Therefore, even if an arbitrary amount of data has accumulated indicating that prices grow at a rate much below dividends - they may actually strongly fall - a Bayesian agent may still believe the opposite to be true. Clearly, this occurs because the price growth beliefs implied by dividend growth beliefs become more and more concentrated at values close to the inverse of the discount factor as $B \rightarrow \delta^{-1}$.²⁰

²⁰This follows from (32) which shows that mean beliefs approach δ^{-1} in combination with the fact that beliefs are bounded above by δ^{-1} .

Small degrees of uncertainty about the value of dividend growth may thus imply that Bayesian agents become very dogmatic about expected price growth.

Ex-post these prior expectations about price growth turn out to be incorrect as the asset turns out to be less attractive than initially thought. This fact is reflected in the decreasing price dividend ratio along the convergence process.

Now suppose one would add to this economy an infinitesimally small internally rational agent. How would this agent behave? An internally rational agent does not need to impose a tight prior on price growth just because there is a small probability of dividends growing at a rate close to δ^{-1} . In fact, uncertainty about the mapping from dividends to prices could induce this agent to impose a much less informative prior on price growth. This allows the internally rational agent to learn more quickly that the asset is less attractive than initially thought, just by observing the disappointing price growth realizations. As a result, such an agent would sell the asset to Bayesian agents at a time when the price is still high. An internally rational agent with a less tight prior on price growth would thus outperform the (subjective) Bayesian agents. This occurs even if both agent types share the same beliefs about dividend growth.

4 Conclusion

We show how to formulate a model with internally rational but not externally rational agents. This entails changing the probability space underlying agents' contingent choices and beliefs. We propose this as a natural departure from the rational expectations assumption while maintaining internal rationality of agents.

The proposed concept works in different ways as Bayesian Rational Expectations Equilibria. When agents are not marginal in all periods and all contingencies - an assumption that may appear to be a reasonable working hypothesis and that arises naturally in models with heterogeneous agents and incomplete markets - internal rationality is fully consistent with expectations about future prices that deviate from expectations about the future discounted sum of dividends. In equilibrium the asset is evaluated according to the marginal agents' expectation of the discounted price and dividend in the next period. Deriving a discounted sum of dividend formulation for asset prices can be achieved by assuming that agents hold rational expectations about asset prices, i.e., are also externally rational. A discounted sum formulation can also be achieved by endowing agents with a tremendous amount of additional information about how the market prices the asset. Even in cases where prices are given by the expected discounted sum of dividends this sum proves to be very sensitive to prior information about dividend growth rates. Asset prices are considerably less sensitive to prior information when they equal the discounted expected price and dividend in the next period.

We conclude that internal rationality is an interesting approach to model rational behavior in models of learning.

A Existence of a Maximum

Strictly speaking the first order conditions (10) may not have to hold in the setup of the main text. This is so for the following reasons: with arbitrary price beliefs, an agent may assign positive probability to prices growing at a rate larger than the inverse of the discount factor, allowing the consumer to achieve arbitrary high levels of utility. When a maximum does not exist, the first order conditions do not have to hold.

This appendix shows that with slightly modified utility functions a maximum always exists for the investor's maximization problem and how the analysis in the main text applies to this modified setup. Consider the following alternative family of utility functions that is indexed by \bar{C}

$$U_{\bar{C}}(C_t^i) = \begin{cases} \bar{C} & C_t^i \leq \bar{C} \\ \bar{C} + f(C_t^i - \bar{C}) & C_t^i > \bar{C} \end{cases}$$

where f is a strictly increasing, strictly concave, differentiable and bounded function satisfying $f(0) = 0$, $f'(0) = 1$ and $f(\cdot) \leq \bar{f}$. Marginal utility of consumption is equal to one for consumption levels below \bar{C} but lower for higher consumption levels. For $\bar{C} \rightarrow \infty$ this utility function converges pointwise to the linear utility function in the main text.

For a given history $\omega = (P_0, D_0, P_1, D_1, \dots)$ the utility generated by some contingent stock holding plan $S = \{S_0, S_1, \dots\}$ with $S_t : \Omega^t \rightarrow [0, \bar{S}]$ is

$$V(S, \omega) = \sum_{t=0}^{\infty} \delta^t U_{\bar{C}}(S_{t-1}(\omega_{t-1})(P_t + D_t) - S_t(\omega_t)P_t)$$

Since

$$V(S, \omega) \leq \frac{\bar{C} + \bar{f}}{1 - \delta} \text{ for all } S \text{ and all } \omega \in \Omega$$

and since \mathcal{P} assigns zero probability to negative values of P and D , this implies that expected utility is bounded. Since the action space S is compact, an expected utility maximizing plan does exist.

Next, we show that for any finite number of periods $T < \infty$, the first order conditions with this bounded utility function are given - with probability arbitrarily close to one - by a set of first order conditions that approximate the ones used in the main text with arbitrary precision. The probability converges to one and the approximation error disappears as $\bar{C} \rightarrow \infty$.

The optimum with bounded utility functions is characterized by the first order conditions

$$\begin{aligned} U'_{\bar{C}}(C_t^i)P_t &< \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i)(P_{t+1} + D_{t+1})] && \text{and } S_t^i = \bar{S} \\ U'_{\bar{C}}(C_t^i)P_t &= \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i)(P_{t+1} + D_{t+1})] && \text{and } S_t^i \in [0, \bar{S}] \\ U'_{\bar{C}}(C_t^i)P_t &> \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i)(P_{t+1} + D_{t+1})] && \text{and } S_t^i = 0 \end{aligned}$$

In any period t , the agent's actual consumption C_t^i in *EQUILIBRIUM* is bounded by the available dividends D_t . Thus, for any $T < \infty$ the probability that

$\{D_t \leq \bar{C}\}_{t=0}^T$ in equilibrium can be brought arbitrarily close to one by choosing \bar{C} sufficiently high. Therefore, with arbitrarily high probability the agent's first order conditions in $t = 1, \dots, T$ are given by

$$P_t < \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] \quad \text{and} \quad S_t^i = \bar{S} \quad (33)$$

$$P_t = \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] \quad \text{and} \quad S_t^i \in [0, \bar{S}] \quad (34)$$

$$P_t > \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] \quad \text{and} \quad S_t^i = 0 \quad (35)$$

Since agents' beliefs satisfy (9) and assign zero probability to negative dividends and prices, we have from Lebegue's Dominated Convergence Theorem

$$\lim_{\bar{C} \rightarrow \infty} E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] = E_t^{\mathcal{P}^i} [(P_{t+1} + D_{t+1})] \quad (36)$$

This implies that for $\bar{C} \rightarrow \infty$ the first order conditions (33)-(35) approximate with arbitrary precision the first order conditions (10) used in the main text.

B Proof of Proposition 1

Proof. Fix t and ω_D^t . For any realization $\omega_D \in \Omega_D$ for which the first t elements are given by ω_D^t , the law of motion for dividends for all $j \geq 1$ is

$$D_{t+j}(\omega_D) = a(\omega_D)^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t(\omega_D^t)$$

so that the partial sum can be expressed as

$$\sum_{j=1}^T \delta^j D_{t+j}(\omega_D) = \sum_{j=1}^T \delta^j a(\omega_D)^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t(\omega_D^t) \quad (37)$$

The partial sums are positive and monotonically increasing in T . The following proves the claim made in part 1.) of the proposition:

$$\begin{aligned} & E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j}(\omega_D) \middle| \omega_D^t \right) \\ &= \lim_{T \rightarrow \infty} E^{\mathcal{P}} \left(\sum_{j=1}^T \delta^j D_{t+j}(\omega_D) \middle| \omega_D^t \right) \end{aligned} \quad (38)$$

$$= \lim_{T \rightarrow \infty} E^{\mathcal{P}} \left(\sum_{j=1}^T \delta^j a(\omega_D)^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t(\omega_D^t) \right) \quad (39)$$

$$\begin{aligned}
&= \lim_{T \rightarrow \infty} D_t(\omega_D^t) \int_0^\infty \left(\sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} & (40) \\
&\geq \lim_{T \rightarrow \infty} D_t(\omega_D^t) \int_{\delta^{-1}}^\infty \left(\sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} \\
&\geq \lim_{T \rightarrow \infty} D_t(\omega_D^t) \cdot T \cdot \int_{\delta^{-1}}^\infty Post_t(\tilde{a}) d\tilde{a} \\
&= \infty
\end{aligned}$$

where the first equality uses Lebesgue's monotone convergence theorem, the second the expression (37), and the third the independence of future η 's from ω_D^t . The first inequality uses the fact that dividends are positive and the second the assumption that $\delta \tilde{a} > 1$ over the considered range of integration. The last equality uses $\int_{\delta^{-1}}^\infty Post_t(\tilde{a}) d\tilde{a} > 0$.

We now prove the second part of the proposition. Define the function

$$\mathcal{F}(\omega_D) = \sum_{j=1}^{\infty} \delta^j B^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t(\omega_D^t)$$

By standard arguments, the infinite sum on the right side exists almost surely and is finite. Therefore, \mathcal{F} is well defined for almost all ω_D and is integrable:

$$E^{\mathcal{P}}(\mathcal{F} | \omega_D^t) = \frac{\delta B}{1 - \delta B} D_t(\omega_D^t) < \infty$$

Moreover, for all n and for given ω_D^t

$$\sum_{j=1}^T \delta^j D_{t+j}(\omega) \leq \mathcal{F}(\omega) \quad a.s.$$

Therefore, the partial sums (37) are bounded a.s. by the integrable function \mathcal{F} , so that we can apply Lebesgue's dominated convergence theorem to obtain the first equality in

$$\begin{aligned}
E^{\mathcal{P}} \left(\lim_{n \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| \omega_D^t \right) &= \lim_{T \rightarrow \infty} E^{\mathcal{P}} \left(\sum_{j=1}^T \delta^j D_{t+j}(\omega_D) \middle| \omega_D^t \right) \\
&= \lim_{T \rightarrow \infty} D_t(\omega_D^t) \int_0^\infty \left(\sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} \\
&= D_t(\omega_D^t) \int_0^\infty \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} \\
&= D_t(\omega_D^t) E^{\mathcal{P}} \left(\frac{\delta a}{1 - \delta a} \middle| \omega_D^t \right) & (41)
\end{aligned}$$

The second equality follows from (38)-(40), the third from applying dominated convergence once more, and the last equality uses the definition of posterior. This proves the first claim made in the second part of the proposition.

Next, we prove the second claim. It follows from (41) that for any $B < \delta^{-1}$ we have

$$\begin{aligned} E^{\mathcal{P}} \left(\sum_{j=1}^{\infty} \delta^j D_{t+j} \middle| \omega_D^t \right) &= D_t(\omega_D^t) E^{\mathcal{P}} \left(\frac{\delta a}{1 - \delta a} \middle| \omega_D^t \right) \\ &= D_t(\omega_D^t) \int_0^B \frac{\delta \tilde{a}}{1 - \delta \tilde{a}} Post_t(\tilde{a}) d\tilde{a} \end{aligned}$$

Marcet and Nicolini (2003) show that

$$\int_0^B \frac{\delta \tilde{a}}{1 - \delta \tilde{a}} Post_t(\tilde{a}) d\tilde{a} \rightarrow \infty \text{ as } B \rightarrow \delta^{-1}$$

by exploiting the fact that this integral behaves like the integral $\int_0^\theta \frac{1}{x} dx$, which is infinite for any $\theta > 0$. This completes the proof. ■

References

- ADAM, K. (2005): “Learning to Forecast and Cyclical Behavior of Output and Inflation,” *Macroeconomic Dynamics*, 9(1), 1–27.
- ADAM, K., A. MARCET, AND J. P. NICOLINI (2008): “Learning and Stock Market Volatility,” *ECB Working Paper No. 862*.
- ANDERSON, R., AND H. SONNENSCHNEIN (1985): “Rational Expectations Equilibrium with Econometric Models,” *Review of Economic Studies*, 52, 359–369.
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *The Annals of Statistics*, 4 (6), 1236–1239.
- BRAY, M., AND D. KREPS (1987): “Rational Learning and Rational Expectations,” in *Arrow and the Ascent of Modern Economic Theory*, ed. by G. Feiwel, pp. 597–625. New York University Press, New York.
- BRENNAN, M. J., AND Y. XIA (2001): “Stock Price Volatility and Equity Premium,” *Journal of Monetary Economics*, 47, 249–283.
- CHAKRABORTY, A., AND K. MITRA (2008): “Can Perpetual Learning Explain the Forward Premium Puzzle?,” *Journal of Monetary Economics (forthcoming)*.
- COGLEY, T., AND T. J. SARGENT (2008): “The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?,” *Journal of Monetary Economics*, 55, 454–476.

- DEGROOT, M. (1970): *Optimal Statistical Decisions*. McGraw-Hill, New York.
- EUSEPI, S., AND B. PRESTON (2008): “Expectations, Learning and Business Cycle Fluctuations,” Columbia University mimeo.
- EVANS, G., AND S. HONKAPOHJA (2003a): “Adaptive Learning and Monetary Policy Design,” *Journal of Money Credit and Banking*, 35, 1045–1072.
- (2003b): “Expectations and the Stability Problem for Optimal Monetary Policy,” *Review of Economic Studies*, 70, 807–824.
- (2005): “Policy Interaction, Expectations and the Liquidity Trap,” *Review of Economic Dynamics*, 8, 303–323.
- GEWEKE, J. (2001): “A Note on some Limitations of CRRA Utility,” *Economics Letters*, 71, 341–345.
- KREPS, D. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge University Press, Cambridge.
- KURZ, M. (ed.) (1997): *Endogenous Economic Fluctuations*. Springer, Berlin.
- MARCEY, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93, 1476–1498.
- MARCEY, A., AND T. J. SARGENT (1989): “Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information,” *Journal of Political Economy*, 97, 1306–22.
- MOLNAR, K., AND S. SANTORO (2007): “Optimal Monetary Policy When Agents are Learning,” *Norwegian School of Economics Working Paper No. 7*.
- ORPHANIDES, A., AND J. C. WILLIAMS (2006): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford. University of Chicago Press.
- PASTOR, L., AND P. VERONESI (2003): “Stock Valuation and Learning About Profitability,” *Journal of Finance*, 58(5), 1749–1789.
- PESARAN, H., D. PETTENUZZO, AND A. TIMMERMANN (2007): “Learning, Structural Instability, and Present Value Calculations,” *Econometric Reviews*, 26, 253–288.
- PRESTON, B. (2005): “Learning About Monetary Policy Rules When Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1(2), 81–126.
- SANTOS, M. S., AND M. WOODFORD (1997): “Rational Asset Pricing Bubbles,” *Econometrica*, 65, 19–57.

- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton Univ. Press, Princeton.
- TIMMERMANN, A. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1994): “Present Value Models with Feedback,” *Journal of Economic Dynamics and Control*, 18, 1093–1119.
- (1996): “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.
- TOWNSEND, R. M. (1983): “Forecasting the Forecasts of Others,” *Journal of Political Economy*, 91, 546–588.
- WEITZMAN, M. L. (2007): “Subjective Expectations and Asset-Return Puzzles,” *American Economic Review*, 97(4), 1102–1130.