DISCUSSION PAPER SERIES

No. 7498
CEPR/EACBN No. 45/2009

INTERNAL RATIONALITY AND ASSET PRICES

Klaus Adam and Albert Marcet

INTERNATIONAL MACROECONOMICS

€ABCN

Euro Area Business Cycle Network

WWW.EABCN.ORG

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP7498.asp

ISSN 0265-8003

INTERNAL RATIONALITY AND ASSET PRICES

Klaus Adam, University of Mannheim and CEPR
Albert Marcet, London School of Economics, CEP and CEPR

Discussion Paper No. 7498
CEPR/EABCN No.45/2009
October 2009

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 7498

October 2009

ABSTRACT

Internal Rationality and Asset Prices

We present a decision theoretic framework with agents that are learning about the behavior of market determined variables. Agents are 'internally rational', i.e., maximize discounted expected utility under uncertainty given consistent beliefs about the future, but may not be 'externally rational', i.e., may not know the true stochastic process for market determined variables (asset prices) and fundamentals (dividends). We apply this approach to a simple asset pricing model with heterogeneity and incomplete markets. We show how knowledge about dividends and optimal behavior alone fail to fully inform agents about equilibrium prices, so that learning about price behavior, as in Adam, Marcet and Nicolini (2008), is fully consistent with internal rationality. We also show that equilibrium prices depend on expectations of the discounted price and dividend in the next period only, rather than on the expected discounted sum of future dividends. Discounted sums emerge only after making very strong assumptions about agents' knowledge and prove extremely sensitive to the details about agents' prior beliefs about the dividend process.

JEL Classification: D83, D84, G12 and G14
Keywords: learning and rationality

Klaus Adam
Department of Economics
Mannheim University
Schloss
68131 Mannheim
GERMANY

Email: adam@uni-mannheim.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=151268

Albert Marcet
London School of Economics
Houghton Street
London
SW2A 2AE
UK

Email: a.marcet@lse.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=110735

Submitted 10 October 2009

The publication of this Paper is funded by the Euro Area Business Cycle Network (www.eabcn.org). This Network provides a forum for the better understanding of the euro area business cycle, linking academic researchers and researchers in central banks and other policy institutions involved in the empirical analysis of the euro area business cycle. Thanks go to Andy Abel, Gaetano Gaballo, Katharina Greulich, Seppo Honkapohja, Erzo Luttmer, Bruce Preston, Iván Werning and, specially, to Juan Pablo Nicolini for interesting comments and suggestions. Marcet acknowledges support from CIRIT (Generalitat de Catalunya), Plan Nacional project ECO2008-04785/ECON (Ministry of Science and Education, Spain) and the Wim Duisenberg fellowship from the European Central Bank.

1 Motivation

We argue that the literature on adaptive learning makes a number of ad-hoc assumptions about agents' behavior. These assumptions lead to a situation where it remains unclear whether or not agents in these models take rational decisions. It also generates controversy when models of adaptive learning are employed in empirical work or for policy analysis, as is the case in an increasing number of contributions.[1] The purpose of this paper is to study if and how these ad-hoc assumptions can be rationalized.

The source of the problem is the following: the learning literature takes as point of departure the first order optimality conditions that emerge under the rational expectations hypothesis (REH); it then replaces the rational expectations operator E appearing in these optimality conditions by an operator of *perceived* expectations $\widetilde{E}$. Agents are then assumed to constantly re-estimate the parameters involved in these perceived expectations in the light of new data. Yet, since first order conditions can be written in many equivalent ways under the REH, there are many ways how one could replace rational expectations by the subjective operator $\widetilde{E}$, and these can lead to rather different results.

As an example, consider the paper of Adam, Marcet and Nicolini (2009) - AMN henceforth. This paper shows that combining standard forms of adaptive learning with a standard asset pricing model allows to replicate a number of empirical stock price puzzles. AMN thereby assume that the stock is priced according to the one-step-ahead equation

$$P_t = \delta \widetilde{E}_t(P_{t+1} + D_{t+1}),$$

which relates the price today (P_t) to expectations of tomorrow's price $(P_{t+1})$and dividend (D_{t+1}) and that agents learn about the behavior of stock prices and dividends.[2] Previously, other papers have studied learning in stock markets but have assumed the stock price to be given by a present value equation

$$P_t = \widetilde{E}_t \sum_{j=1}^{\infty} \delta^j D_{t+j},$$

[1] To mention a few, see Adam, Marcet and Nicolini (2009), 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.

[2] The parameter $\delta \in (0, 1)$ denotes the discount factor, see section 2 for details.

which relates the current price to the discounted sum of dividends, thereby allowing for learning about dividend behavior only. These latter papers found that learning does not add much to explaining the puzzling behavior of stock prices. Which is the 'correct' way to model learning in stock markets?

The approach taken in AMN has been criticized along a number of dimensions: the one-step-ahead formula, it has been argued, implies myopic investment behavior and, therefore, a potentially large departure from rationality, while the discounted dividend formula implies that agents take into account the whole indefinite future; also, since agents in AMN change their expectations $\widetilde{E}$ about prices each period according to a simple updating rule, the resulting process for agents' expectations has been conjectured to be dynamically inconsistent; finally, it has been argued that optimizing agents should realize that dividend beliefs restrict price beliefs according to the present value equation, implying that there is just no role for learning about stock price behavior if agents are rational.[3]

The objective of this paper is to evaluate these arguments by considering strictly optimizing agents. As a first step towards this goal we propose to separate the standard rationality requirements of rational expectations (RE) models into an 'internal' and an 'external' rationality component. *Internal rationality* requires that agents make fully optimal decisions given a well defined system of beliefs. This implies that agents maximize utility under uncertainty, given their constraints and given a consistent set of probability beliefs about payoff relevant variables that are beyond their control or 'external'. *External rationality* postulates that agents know the objective probability distributions of external variables as they emerge in equilibrium. We propose to fully maintain internal rationality but argue that it should be of economic interest to study the implications of slightly relaxing the assumption of external rationality. This involves allowing for the possibility that agents' probability beliefs about external variables may not coincide - during a transitional period - with the true equilibrium distributions of these variables. Behind this proposal lies the basic idea that internal rationality is a good starting point for analyzing social interactions, but that achieving external rationality requires an amount of knowledge about the behavior of markets and fundamentals that not even expert economists have. Therefore, it should be of interest to allow for small deviations from the REH in economic models.

Perhaps not surprisingly, internal rationality implies external ratio-

[3]One of the critics in a seminar argued that 'agents [in AMN] can not add two plus two'.

nality about prices in a number of well-known - albeit highly restrictive - asset pricing models. For example, if markets are complete (in the sense that all time and state contingent markets are open at time zero), then agents can observe all prices and absence of arbitrage in equilibrium leaves no room for subjective price expectations. Likewise, when all agents are identical and this is commonly known, then equilibrium prices can be inferred from the agents' own decision problem.

Yet, in a setting with heterogeneous agents and standard forms of market incompleteness, the equilibrium stock price is determined by a one-step ahead asset pricing equation, where the price and dividend expectations are those held by the marginal investor. In such a setting an investor can not infer a discounted sum formulation of the asset price just from introspection and knowledge of the dividend process. Indeed, it can be optimal for the agent to pay a high price today - even if the agent expects the discounted sum of dividends to be low - provided the agent expects to be able to sell the stock at a higher price tomorrow. This is so because it is optimal to engage in speculative trading in the sense of Harrison and Kreps (1978).

In addition to showing that the asset is priced accordingto a one-step-ahead pricing equation, we demonstrate how the updating equations for the perceived expectations $\widetilde{E}$ used in AMN arise from a complete and dynamically consistent set of probability beliefs. Moroever, we show that these beliefs involve only small deviations form the beliefs entertained under the REH. We therefore show that the asset pricing implications in AMN are consistent with internal rationality and involve only small deviations from rational expectations beliefs.

More generally, this paper shows how to formulate decision problems when agents possess only imperfect market knowledge. Specifically, we argue that the probability space over which agents condition their choices and form expectations should include all *external* variables, including prices. This departs from the standard formulation in the literature where the probability space is reduced from the outset to contain only *exogenous* (or 'fundamental') variables, and prices are assumed to be a function of fundamentals. For example, in the stock pricing model the standard assumption is that dividends span the agents' probability space and that prices are a known function of dividends.[4] This imposes a singularity in the joint distribution of stock prices and dividends. Instead, we allow agents to entertain a non-degenerate joint distribution of prices and dividends. Even though this is a potentially small departure from rational expectations, we show that the one-step-ahead formula for

[4] This assumption is also made in the literature on 'rational bubbles', e.g., Santos and Woodford (1997).

stock prices arises under this probability space.

We then delineate conditions under which an optimizing agent can map the process for dividends into prices, so that a discounted sum equation emerges. It turns out that this is possible only if the agent is endowed with a tremendous amount of additional information about the market: the agent needs to possess the same information as the theorist, namely, the agent needs to know all details about other agents in the economy (like all other agents' probability beliefs, discount factors, etc.), and must possess the ability to compute solutions to very high-dimensional general equilibrium problems.

A standard way to relax the strong informational assumptions underlying rational expectations equilibria (REE) has been the concept of Bayesian REE. These equilibria allow for imperfect information about the distribution of exogenous variables (fundamentals). We argue, however, that Bayesian REE assumes much more knowledge about market outcomes than can be derived from internal rationality.[5] In particular, agents need to have the superior knowledge that *i)* there is a singularity in the joint distribution of market outcomes and fundamentals, *ii)* this singularity is consistent with market equilibrium and *iii)* this singularity solves a fixed point problem that only the authors of a paper will normally know how to compute. Notably, agents need to know *i), ii),* and *iii)* from the outset, they can not learn about these issues. Only under all these conditions does the equilibrium stock price equal the discounted sum implied by a Bayesian REE. Assuming that agents know this 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 in the real world: the empirical literature in asset prices has had a very hard time in detecting a stable mapping between dividends and prices in the data. For this reason we think it is interesting to consider models that do not impute such market knowledge (external rationality) to investors.

We also show that the results under Bayesian REE are extremely sensitive to fine details of the prior distribution about the dividend process incorporated in agents' beliefs.[6] Based on this we conclude that agents' prior beliefs matter much more than other economic factors for the equilibrium stock price in a Bayesian REE, so that the pricing implications in Baysian REE models appear highly ad-hoc.

Our work is related to a number of papers in the learning literature that attempt to construct a full set of beliefs over long-horizons, e.g.,

[5] See also the discussion in Bray and Kreps (1987).

[6] Related but different findings of this kind have shown up in Geweke (2001), Weitzman (2007) and Pesaran, Pettenuzzo and Timmermann (2007).

Preston (2005) and Eusepi and Preston (2008). The main difference is that our agents' beliefs take the form of a well defined probability measure over a stochastic process while these papers 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 set of beliefs held in the previous period. The setup in this paper is also indirectly related to the literature on rational beliefs initiated by Mordecai Kurz (1997), where agents' probability distributions are assumed to be shifting in response to the realization of an extrinsic 'generating sequence'. In the present paper belief revisions are triggered by intrinsic factors, namely market outcomes (price observations) and fundamentals (dividend realizations), and they are consistent with a well defined system of beliefs. The paper is also related to the model-consistent equilibrium concept of Anderson and Sonnenschein (1985) who assume that agents have a probability distribution over prices but impose a kind of rational expectations structure on the beliefs of agents.

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 (2009). Section 3 then derives the information about the market that is required to map dividend beliefs into stock prices. This section also presents a formal result about the strong sensitivity of discounted dividends 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 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 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 rational expectations assumptions often present in the literature.

2.1 Basic model

The economy has $t = 0, 1, 2, ...$ 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 $1/I$ units of an infinitely lived stock. Agents of type $i \in \{1, ..., I\}$ have a standard time-separable utility function

$$E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \left(\delta^i\right)^t C_t^i \tag{1}$$

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 $(\Omega, \mathcal{S})$. As usual, the probability measure $\mathcal{P}^i$ is a model primitive and given to agents. It is allowed to be type-specific and it may or may not coincide with objective probabilities. 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.

The non-standard part in our formulation is in the underlying probability space. Let P_t denote the stock price and D_t the stock's dividend payments. We consider agents who view the process for $\{P_t, D_t\}$ as external to their decision problem and the probability space over which they condition their choices is given by

$$\Omega \equiv \Omega_P \times \Omega_D$$

where $\Omega_P = \prod_{i=0}^{\infty} R_+$ ($\Omega_D = \prod_{i=0}^{\infty} R_+$) contains all possible sequences of prices (dividends). 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 $\mathcal{P}^i$ over $(\Omega, \mathcal{S})$. As usual we denote the set of all possible histories up to period t by $\Omega^t = \Omega_P^t \times \Omega_D^t$ and its typical element $\omega^t \in \Omega^t$.

With this setup investors can condition their decisions on the history of observed dividend and price realizations. This is a natural setup in a model of competitive behavior: since investors see prices as a shock that influences their budget constraint they want to make their future choices contingent on realizations of prices, in addition to dividends.

Throughout the paper we will make the statement that our agents have 'a consistent set of beliefs'. By this we simply emphasize that $(\Omega, \mathcal{S}, \mathcal{P}^i)$ is a proper probability space. $\mathcal{P}^i$ satisfies all the standard probability axioms and it gives proper joint probabilities to all possible values of prices and dividends in any set of dates.

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

$$\left(C_t^i, S_t^i\right) : \Omega^t \to R^2 \tag{2}$$

so that period t choices are contingent on all information available up to time t. Expected utility (1) can thus be written as

$$E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \left(\delta^i\right)^t \; C_t^i = \int_\Omega \sum_{t=0}^{\infty} \left(\delta^i\right)^t \; C_t^i(\omega^t) \; d\mathcal{P}^i(\omega). \tag{3}$$

The agent faces the following budget constraint

$$C_t^i + P_t \; S_t^i \leq (P_t + D_t) \, S_{t-1}^i + \xi \tag{4}$$

which has to hold for all t and all $\omega^t \in \Omega^t$. Here ξ denotes an endowment of consumption goods, this is introduced for simplicity, it allows us to ignore non-negativity constraints on consumption.

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

$$S_t^i \geq 0 \tag{5}$$

$$S_t^i \leq \overline{S} \tag{6}$$

where $1 < \overline{S} < \infty$. Constraint (5) is a standard short-selling constraint and often used in the literature. The second constraint (6) is a simplified form of a leverage constraint capturing the fact that the consumer cannot buy arbitrarily large amounts of stocks. Constraint (6) simplifies our equilibrium calculation in the presence of risk neutral investors.

We consider agents who choose (2) in order to maximize utility (3) subject to the budget constraint (4), the limit constraints (5) and (6), taking as given the probability measure $\mathcal{P}^i$. Such agents are called internally rational. Throughout the paper we assume that $\mathcal{P}^i$ satisfies

$$E^{\mathcal{P}^i}[P_{t+1} + D_{t+1} | \omega^t] < \infty \text{ for all } \omega, t \tag{7}$$

We also assume that a maximum of the investor's utility maximization problem exists.[7]

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

2.2 Optimality conditions

The only nonstandard element of the setup just described is that the space of outcomes Ω includes histories of prices, while in the standard formulation it would include only states of nature (dividends) only. This allows agents to evaluate first order conditions (8) at any possible history in Ω^t.

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

$$P_t < \delta^i \ E_t^{\mathcal{P}^i} \left(P_{t+1} + D_{t+1} \right) \qquad \text{and } S_t^i = \overline{S} \tag{8a}$$

$$P_t = \delta^i \ E_t^{\mathcal{P}^i} \left(P_{t+1} + D_{t+1} \right) \qquad \text{and } S_t^i \in \left[0, \overline{S}\right] \tag{8b}$$

$$P_t > \delta^i \ E_t^{\mathcal{P}^i} \left(P_{t+1} + D_{t+1} \right) \qquad \text{and } S_t^i = 0 \tag{8c}$$

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 (8) 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. Using popular language our agents simply try to 'buy low and sell high' as much as the stock holding constraints allow them to do. This is the optimal strategy under risk neutrality and without transaction costs.

2.3 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 knowledge of the function

$$P_t : \Omega_D^t \rightarrow R_+ \tag{9}$$

that maps dividend realization into equilibrium 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 Rational Expectations Equilibrium. With rational expectations, such degenerate beliefs are 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.

Yet, when departing the from rational expectations hypothesis, it appears to be very restrictive to assume that agents know about the existence of an exact equilibrium mapping determining prices as a function of past dividends. 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 allows to endow 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.

2.4 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 (9).

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 $\overline{S}$ in constraint (6) is large enough such that it never binds. Thus, from the first order conditions (8), 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:

$$\mathbf{P}_t = \max_{i \in I} \left[\delta^i \; E_t^{\mathcal{P}^i} \left(P_{t+1} + D_{t+1} \right) \right] \tag{10}$$

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 uniqueness assumption insures that the equilibrium price is indeed a function of the history ω_D^t of dividend realizations only, i.e., $\mathbf{P}_t : \Omega_D^t \to R_+$.[8] 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. Although this statement is indeed correct, it fails to answer the following simple question: which precise degeneracy should agents impose on their beliefs $\mathcal{P}^i$? As we show next, even if agents knew the equilibrium asset pricing equation (10), they still cannot derive the exact degeneracy just from knowledge of their own utility functions and their dividend beliefs. This suggests a natural interpretation of why agents beliefs might not contain a singularity, namely that agents are uncertain about its exact location.

Let m_t denote the marginal agent pricing the asset in period t:[9]

$$m_t = \arg\max_{i \in I} \left[\delta^i \; E_t^{\mathcal{P}^i} \left(P_{t+1} + D_{t+1} \right) \right]$$

Since m_t is determined in equilibrium, we have $m_t : \Omega_D^t \to \{1, \ldots, I\}$. The equilibrium price (10) can thus be written as

$$\mathbf{P}_t = \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(P_{t+1} + D_{t+1} \right) \tag{11}$$

[8] The mapping $\mathbf{P}_t$ will, of course, depend on all 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 and discount factors.

[9] If 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.

Suppose it is common knowledge to agents that the equilibrium price satisfies equation (11) each period.[10] Common knowledge thereby means that each agent knows that the asset is priced according to (11) each period, that each agents knows that other agents know this be the case, knows that other agents know that others know it to be true, and so on to infinity.[11] The question we are posing is: would common knowledge of equation (11) 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 still rationally hold price beliefs that are either larger or smaller than the own expectations of the discounted sum of dividends.

Common knowledge of (11) implies that agents know that the equilibrium price must satisfy the recursion

$$\mathbf{P}_t = \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(\mathbf{P}_{t+1} + D_{t+1}\right) \tag{12}$$

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:

$$\begin{aligned}
\mathbf{P}_t &= \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(D_{t+1}\right) \\
&+ \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m\,t+1}} D_{t+2}\right) \\
&+ \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m\,t+1}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m\,t+2}} D_{t+3}\right)\right) \\
&+ ... \\
&+ \delta^{m_t} \; E_t^{\mathcal{P}^{m\,t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m\,t+1}} \left(\ldots \delta^{m_{t+T}} E_{t+T}^{\mathcal{P}^{m\,t+T}} \left(\mathbf{P}_{t+T+1} + D_{t+T+1}\right)\right)\right)
\end{aligned} \tag{13}$$

The right hand side of the previous expression provides 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 their 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 (11), and internal rationality fail to imply a *specific* singularity in agents' price and

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

[11] See Aumann (1976) for a formal definition.

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 $E_t^{\mathcal{P}^i}\left(\sum_{j=1}^{\infty}\left(\delta^i\right)^j D_{t+j}\right)$ held by any agent i in the economy, as is the case in AMN (2008).

This explain why, in our setup, it is unlikely that agents will take decisions by comparing the stock price to a discounted sum of dividends. In the literature, the discounted sum of dividends is usually obtained by applying the law of iterated expectations on the right side of equation (13). 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 in (13) 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 2.6.[12]

Going back to the criticisms of adaptive learning (and in particular of AMN) discussed in the introduction, we have just shown, that beliefs about dividends and internal rationality do not imply a belief about prices. Moreover,.when agents hold reasonable amounts of uncertainty about the mapping between prices and dividends, the one-step-ahead pricing equation is not a consequence of myopic investment behavior, but follows from infinite horizon optimization by investors. Instead, agents' expectations about the discounted sum of dividends turns out to be irrelevant for determining their optimal investment choice.

[12] The 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.

2.5 Internal Rationality with Bayesian Learning

An important criticism of adaptive learning models is that it is unclear whether or not the learning equations used in this literature to describe the updating of expectations over time give rise to an expectations process that is dynamically consistent. In this paper we formulated our decision problem using a consistent system of beliefs $\mathcal{P}^i$ that is given to agents at the beginning of time. We now show that the learning schemes used in the adaptive learning literature can actually arise from such a consistent system of beliefs. Specifically, we construct a probability measure $\mathcal{P}$ for which optimal behavior gives rise to the ordinary least-squares learning equations employed by agents in AMN (2009). We also make precise the statement in what sense agents' prior beliefs $\mathcal{P}$ in AMN (2009) are a small deviation 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, it is trivial to include different discount factors. We thus continue to consider agents that use the one-period optimality conditions (8) to decide on stockholdings.

Let us first consider the equilibrium under rational expectations. The true process for dividends is

$$\log \frac{D_t}{D_{t-1}} = \log a + \log \varepsilon_t \tag{14}$$

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

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

so that the equilibrium process for the asset price follows

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

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 $\mathcal{P}$ encompasses the REE equilibrium, but does 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

$$\begin{bmatrix} \log P_t/P_{t-1} \\ \log D_t/D_{t-1} \end{bmatrix} = \begin{bmatrix} \log \beta^P \\ \log \beta^D \end{bmatrix} + \begin{bmatrix} \log \varepsilon_t^P \\ \log \varepsilon_t^D \end{bmatrix} \tag{15}$$

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

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

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 $\left(\log \beta^P\right)$ 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 (15) fully determine agents' probability measure $\mathcal{P}$ over infinite sequences of price and dividends realizations.[13] We have thus specified the microfoundations of the model and of beliefs.

Since the FOCs (8) use conditional expectations now we need to derive the evolution of the conditional distribution of prices and dividends. For this purpose it is convenient to assume that f is of the Normal-Wishart conjugate form

$$H \sim W(S_0, n_0) \tag{16a}$$

$$\left(\log \beta^P, \log \beta^D\right)' \Big| H \sim N\left(\left(\log \beta_0^P, \log \beta_0^D\right)', (\nu_0 H)^{-1}\right) \tag{16b}$$

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

[13] Probabilities for $\mathcal{P}$ can be obtained mechanically as follows: for any Borel subset $s \subset \mathcal{S}$, determine the probability of prices and dividends being in s for any given value of $(\log \beta^P, \log \beta^D, \Sigma)$ using standard methods for Markov processes applied to equation (15). Then integrate these probabilities over values of $(\log \beta^P, \log \beta^D, \Sigma)$ according to f.

agents' marginal prior distribution about the inverse of the variance covariance matrix of innovations $H \equiv \Sigma^{-1}$, where n_0 scales the precision of prior beliefs. The normal distribution N specifies agents' priors about the parameters $\left(\log\beta^P, \log\beta^D\right)$ conditional on the precision matrix H, where $\left(\log\beta_0^P, \log\beta_0^D\right)$ denotes the conditional prior mean and $\nu_0 > 0$ scales the precision of prior beliefs about $\left(\log\beta^P, \log\beta^D\right)$.

The previous specification nests rational expectations a special case. Under 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'. This is a special case of the Normal-Wishart prior when

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

and considering the limiting case of vanishing prior uncertainty:

$$n_0 \to \infty$$
$$v_0 \to \infty$$

Therefore, for large n_0, v_0, the beliefs of agents are close to the beliefs of REE. More precisely, for any given subset of Ω^t, the probability assigned to this subset by $\mathcal{P}$ is arbitrarily close to the probability assigned by RE beliefs for sufficiently large n_0, v_0. In this sense, $\beta_0^P = \beta_0^D = a$ and large n_0, v_0 denote a small deviation from RE.

When $\mathcal{P}$ is given by (15) and (16) the conditional distribution of price and dividend can be derived using standard arguments. In particular, it is enough to find the distribution of β_P, β_D, Σ conditional on information up to t, since future innovations are uncorrelated with the past. In Bayesian language, we need to find the posterior distribution of β_P, β_D, Σ given information at time t. As is well know, the prior (16) is conjugate, so that the posterior distribution is also a Normal-Wishart of the form (16), but with location parameters $\left(\log\beta_t^P, \log\beta_t^D, \nu_t, S_t, n_t\right)$ instead of $\left(\log\beta_0^P, \log\beta_0^D, \nu_0, S_0, n_0\right)$. 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

$$\begin{pmatrix} \log \beta^P_{t+1} \\ \log \beta^D_{t+1} \end{pmatrix} = \begin{pmatrix} \log \beta^P_{t} \\ \log \beta^D_{t} \end{pmatrix} + \frac{1}{\upsilon_t + 1} e_t \tag{17a}$$

$$\upsilon_{t+1} = \upsilon_t + 1 \tag{17b}$$

$$S^{-1}_{t+1} = S^{-1}_t + \frac{\upsilon_t}{\upsilon_t + 1} e_t e'_t \tag{17c}$$

$$n_{t+1} = n_t + 1 \tag{17d}$$

In other words, the posterior distribution $\left(\log \beta^P, \log \beta^D\right)$ is

$$H \sim W(S_t, n_t) \tag{18a}$$

$$\left(\log \beta^P, \log \beta^D\right)' \Big| H, \omega^t \sim N\left(\left(\log \beta^P_t, \log \beta^D_t\right)', (\nu_t H)^{-1}\right) \tag{18b}$$

and the posterior mean for $\left(\log \beta^P, \log \beta^D\right)$ in period t is:[14]

$$E^{\mathcal{P}}_t \left[\left(\log \beta^P, \log \beta^D\right)\right] = \left(\log \beta^P_t, \log \beta^D_t\right) \tag{19}$$

Now recall that the first order condition of the marginal agent (11) with beliefs $\mathcal{P}$ delivers

$$\begin{aligned} \mathbf{P}_t &= \delta E^{\mathcal{P}}_t \left(P_{t+1} + D_{t+1}\right) \\ &= \delta E^{\mathcal{P}}_t \left(e^{\log \beta^P}\right) E\left(e^{\log \varepsilon^P_t}\right) \mathbf{P}_t + \delta E^{\mathcal{P}}_t \left(e^{\log \beta^D}\right) E\left(e^{\log \varepsilon^D_t}\right) D_t \end{aligned} \tag{20}$$

where the second equality follows from applying standard rules of conditional expectations using the form of the beliefs (15). The previous equation shows that one has to evaluate the expectation of exponential functions of $\log \beta^P$ and $\log \beta^D$, which can only be done numerically. We prefer, however, to obtain analytical results that involve the closed form expressions from the conjugate posterior beliefs. Therefore, we proceed by computing first order approximations to these expectations. Letting

[14] This follows from the fact that the marginal posterior for price and dividend growth is Student t-distributed with $n_t - 1$ degrees of freedom, location vector $\left(\log \beta^P_t, \log \beta^D_t\right)'$ and precision matrix $\upsilon_t \left(n_t - 1\right) S_t$, see chapter 9 in DeGroot (1970).

'$\approx$' denote an equality that is correct up to first order we have that

$$
\begin{aligned}
E^{\mathcal{P}}[e^{\log \beta^P}] &\approx e^{E^{\mathcal{P}} \log \beta^P} \\
&= e^{\frac{1}{t+v_0}\left(\sum_{j=1}^{t} \log \frac{P_j}{P_{j-1}}\right) + \frac{v_0}{t+v_0} \ln \beta_0^P} \\
&= \left(\left(\prod_{j=1}^{t} \frac{P_j}{P_{j-1}}\right) \left(\beta_0^P\right)^{v_0}\right)^{\frac{1}{t+v_0}} \\
&\approx 1 + \frac{1}{t+v_0} \sum_{j=1}^{t} \left(\frac{P_j}{P_{j-1}} - 1\right) + \frac{v_0}{t+v_0} \left(\beta_0^P - 1\right) \\
&= \frac{1}{t+v_0} \left(\sum_{j=1}^{t} \frac{P_j}{P_{j-1}}\right) + \frac{v_0}{t+v_0} \beta_0^P \\
&\equiv \widehat{\beta}_t^P
\end{aligned}
$$

where the first approximation has been taken with respect to $\log \beta^P$ at the point $E^{\mathcal{P}} \log \beta^P$ and the second approximation at $P_j/P_j = 1$ and $\beta_0^P = 1$. Expected price growth is thus given - up to a first order approximation - by a simple ordinary least squares (OLS) estimate involving observed price growth and prior beliefs β_0^P. Similar approximations for the other terms in (20) deliver

$$
\begin{aligned}
E_t^{\mathcal{P}}\left(e^{\log \beta^D}\right) &\simeq \widehat{\beta}_t^D \equiv \frac{1}{t+v_0} \sum_{i=0}^{t} \left(\frac{D_i}{D_{i-1}}\right) + \frac{v_0}{t+v_0} \beta_0^D \\
E\left(e^{\log \varepsilon_t^P}\right) &\simeq 1 \simeq E\left(e^{\log \varepsilon_t^D}\right)
\end{aligned}
$$

Equation (20) thus implies - up to a linear approximation of beliefs - that

$$
\mathbf{P}_t = \delta \; \widehat{\beta}_t^P \; \mathbf{P}_t + \delta \; \widehat{\beta}_t^D \; D_t \tag{21}
$$

which is exactly the equation used in AMN with $\widehat{\beta}_t^P$ and $\widehat{\beta}_t^D$ denoting the least squares estimates.[15] We have thus shown how the pricing implications in AMN arise (approximately) from a model populated with internally rational agents that hold a complete and consistent set of probability beliefs. Moreover, setting initial prior means to $\beta_0^P = \beta_0^D = a$ and for very large n_0, v_0 this is arbitrarily close to the beliefs that give rise to the REE. It is not true, therefore, that using OLS each period to estimate the growth rate of prices amounts to changing beliefs each

[15] To avoid simultaneity, one final convenient approximation in the actual calculations of AMN is that $\widehat{\beta}_t^P$ and $\widehat{\beta}_t^D$ are computed disregarding the last observed growth rate.

period in an inconsistent manner. In fact this is what arises from behaving optimally under a consistent set of beliefs that we believe to be a natural one.

Before discussing in the next section how this setup deviates from a Bayesian REE, we wish to briefly address the issue of equilibrium existence. Clearly (21) implies that for that for $\widehat{\beta}_t^P < \delta^{-1}$ there exists a unique equilibrium price given by

$$\mathbf{P}_t = \frac{\delta \widehat{\beta}_t^D}{1 - \widehat{\beta}_t^P \delta} D_t \tag{22}$$

Existence thus requires that the (approximate) posterior *mean* for expected price growth $\widehat{\beta}_t^P$ remains below δ^{-1}. In AMN (2009) this condition was insured by imposing a projection facility on beliefs. This facility constitutes a deviation from Bayesian learning. The details of this facility do not affect substantially the results on second moments, since 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 alternative would be to choose a prior f that imposes that the outcome $\beta^P \geq \delta^{-1}$ occurs with zero probability. This would avoid the posterior ever being above δ^{-1} but it would result in a modification of the above formulae. It is likely to lead to Bayesian updating behavior very similar to that implied by the projection facility imposed in AMN. We have not pursued this line further at this point.

2.6 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 (15), which imply that their perceived likelihood for prices conditional on past information is given by

$$l(\log P_t \mid \Omega^{t-1}; \log \beta^P, \sigma_P^2) \sim N(\log \beta_{t-1}^P + \log P_{t-1}, \sigma_P^2) \tag{23}$$

Agents thus belief that asset prices follow 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 (10). While the random walk behavior implied by (23) will become true asymptotically - this is shown in AMN (2009) - 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.

Furthermore, the above model has a singularity, since the only fundamental shock is indeed dividends. We can express this singularity in the following way: letting l^L be the true likelihood under learning conditional on knowledge of the model and the actual value of all parameters, we have

$$l^L(P_t \mid \Omega_D^t) = \infty \qquad \text{for } P_t = \mathbf{P}_t \tag{24}$$

$$= 0 \qquad \text{otherwise} \tag{25}$$

while agents' perceived likelihood has a proper distribution.

The issue is: 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. Therefore, the results under Bayesian REE are very highly dependent on ad-hoc assumptions about extreme values of the prior distribution.

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 (8). 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 (12) 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 (12) and to express it as a function of future dividends and some terminal price and, therefore, to obtain equation (13). 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 (12) which holds with equality in all periods and all contingencies, but this equation also involves a variety of information about other agents.

The discounted sum expressions (13) also contains 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

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

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

$$\lim_{T\to\infty} E_t^{\mathcal{P}}\left(\delta^T\ \mathbf{P}_{t+T}\right) = 0 \tag{26}$$

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 (26) restricts the behavior of the equilibrium price $\mathbf{P}$.

Assumption 2 allows to take the limit $T \to \infty$ in equation (13) 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 $P_t : \Omega_D^t \to R$, 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 (13), 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

$$\begin{aligned} E_t^{\mathcal{P}^i}\mathbf{P}_{t+1} &= E_t^{\mathcal{P}^i}\left(\delta^{m_{t+1}}\ E_{t+1}^{\mathcal{P}^{m\,t+1}}\left(D_{t+2}\right)\right) \\ &\quad + E_t^{\mathcal{P}^i}\left(\delta^{m_{t+1}}\ E_{t+1}^{\mathcal{P}^{m\,t+1}}\left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m\,t+2}} D_{t+3}\right)\right) \\ &\quad + E_t^{\mathcal{P}^i}\left(\delta^{m_{t+1}}\ E_{t+1}^{\mathcal{P}^{m\,t+1}}\left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m\,t+2}}\left(\delta^{m_{t+3}} E_{t+3}^{\mathcal{P}^{m\,t+3}} D_{t+4}\right)\right)\right) \\ &\quad + \ldots \end{aligned} \tag{27}$$

The agent's price expectations are thus implied by their beliefs about the process $\{m_t\}$ and their beliefs about $\left\{\delta^j, \mathcal{P}^j\right\}_{j\neq i}$ 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 $\left\{\delta^j, \mathcal{P}^j\right\}_{j\neq i}$, but know that in equilibrium these fundamentals are functions of the history of dividend realizations, i.e.,

$$m_t : \Omega_D^t \to \{1, 2, \ldots, I\} \tag{28}$$
$$\mathcal{P}^j : \mathcal{S} \to [0, 1] \tag{29}$$
$$\delta^j : \Omega_D \to [0, 1] \tag{30}$$

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 (28)-(30) 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 (27) 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 (28)-(30) 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 (28)-(30) 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.

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

$$\begin{aligned}
\mathbf{P}_t = {} & \delta^{m_t}\ E_t^{\mathcal{P}^{m_t}}\left(D_{t+1}\right) \\
& + \delta^{m_t}\ E_t^{\mathcal{P}^{m_t}}\left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} D_{t+1}\right) \\
& + \delta^{m_t}\ E_t^{\mathcal{P}^{m_t}}\left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}}\left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+2}\right)\right) \\
& + \ldots
\end{aligned} \tag{31}$$

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., a representative agent model with sequentially complete markets and price beliefs that satisfy the no rational bubble requirement (26). 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 the equality between the market clearing asset price and the expected discounted sum of dividends is the result of (internally) rational investment behavior on the side of agents, but - as we have just shown - this is not the case once one considers 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]

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?. The interest in Bayesian REE arises from a desire to give agents less information than in a full information REE. For example, in the stock price model, one may want to relax the assumption that agents know exactly the average growth rate of dividends a and instead assume that agents have a prior distribution about a. But standard practice is then to assume that agents have an exact knowledge of the singularity that maps dividends into prices which means that, if indeed agents are heterogeneous and have different beliefs about a, that they have a huge amount of knowledge about how the market functions. This may be a useful modelling strategy in some cases, but leads to a very asymmetric treatment of knowledge about external variables.

[16] As we mentioned in footnote 12, 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.

3.2 Sensitivity of Bayesian RE asset prices

We now consider a representative agent model with risk-neutrality but we depart from the previous setup in that we assume enough assumptions hold so that the standard formula where prices equal the expected discounted sum of dividends arises. Specifically, letting δ denote the agent's discount factor and $\mathcal{P}$ the agent's beliefs about the dividend process we now consider a case where

$$P_t = E^{\mathcal{P}}\left(\lim_{T\to\infty} \sum_{j=1}^{T} \delta^j D_{t+j} \middle| \omega_D^t \right) \tag{32}$$

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

$$D_t = aD_{t-1}\eta_t \tag{33}$$

where $\eta_t > 0$ is i.i.d. with $E[\eta_t] = 1$ and $a > 0$. We consider an agent that knows the distribution of innovations η_t but is uncertain about the true mean dividend growth rate a. 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 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 (32) 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. The proof of the proposition can be found in appendix B.

[17] The results regarding the sensitivity of the sum (32) 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.

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

1. *Assume:*

   i) $B > \delta^{-1}$, or

   ii) $B = \delta^{-1}$ and $Post_t(B) > 0$,

   then

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

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

$$E^{\mathcal{P}}\left(\lim_{T\to\infty}\sum_{j=1}^{T}\delta^j D_{t+j}\middle|\,\omega_D^t\right)=D_t\ E_{Post_t}\left(\frac{\delta a}{1-\delta a}\right)$$
$$=D_t\int\frac{\delta\widetilde{a}}{1-\delta\widetilde{a}}\ Post_t(\widetilde{a})\ d\widetilde{a}<\infty$$

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 $k \to \infty$. 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*

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

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 *any* finite number of 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.[18,19] Instead of using the asset price equations in proposition 1, the following formula has typically been employed:

$$\mathbf{P}_t = D_t \frac{\delta \ E_{Post_t}(a)}{1 - \delta \ E_{Post_t}(a)} \tag{34}$$

where

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

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 (34) 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. As shown in proposition 1, 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 2.5, the pricing equation (22) shows that as long as the posterior *mean* about price growth belief β_t remains bounded away from

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

[19]To our knowledge Pastor and Veronesi (2003) is the only paper that correctly takes into account the non-linearity issue. They assumed the existence of a finite asset price at some fixed terminal date $T < \infty$, insuring that the discounted sum remains bounded.

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 (33) 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 $Post_{-1}(\cdot)$ 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

$$\mathbf{P}_t = D_t \; E_{Post_t}\left(\frac{\delta a}{1-\delta a}\right) < \infty \tag{35}$$

Suppose that $Post_{-1}(\cdot)$ 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 (35) 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

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

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 (12), which together with equation (35) delivers

$$E_{Post_t}\left[\frac{\mathbf{P}_{t+1}}{\mathbf{P}_t}\right]=\frac{1}{\delta}-\frac{E_{Post_t}[a]}{E_{Post_t}\left(\frac{\delta a}{1-\delta a}\right)} \tag{36}$$

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 $E_{Post_t}\left(\frac{\delta a}{1-\delta a}\right)\to\infty$ as $B\to\delta^{-1}$, so that equation (36) 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 $B\to\delta^{-1}$.[20] 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

[20] This follows from (36) which shows that mean beliefs approach δ^{-1} in combination with the fact that beliefs are bounded above by δ^{-1}.

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 agents that fail to be externally rational because they possess limited knowledge about the market. 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 rational decision making by agents given their beliefs about the external environment.

The proposed concept works in different ways as Bayesian Rational Expectations Equilibria. When agents are not marginal in all periods and all contingencies - an reasonable assumption that naturally holds 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 then evaluated according to the marginal agents' expectation of the discounted price and dividend in the next period. Optimal behavior, therefore, focuses on price prediction.

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. In a setup with incomplete information and heterogeneous agents this amounts to endowing agents with a tremendous amount of additional information about other agents and about how the market prices the asset. In effect, Bayesian REE assume that agents know exactly how to map fundamentals into prices, therefore it assumes that agents know exactly how to locate a singularity in the joint distribution of prices and dividends. Thus, a discounted sum formulation emerges in the Bayesian REE only because the equilibrium concept imposes very asymmetric assumptions on what agents know about fundamentals compared to prices.

Whenever prices are given by the expected discounted sum of dividends, the price proves to be very sensitive to ad-hoc assumptions on the prior information about dividend growth rates for high values of these 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 behavior in models of learning. It provides guidance about what models of learning are fully consistent with optimizing behavior and, in doing so, it provides a rationale for asset pricing equations with very different implications than those emerging from the standard REE approach.

A Existence of a Maximum

Strictly speaking the first order conditions (8) may not have to hold in the setup of the main text. This is so because 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 $\overline{C}$

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

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

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

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

Since

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

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 $\overline{C} \to \infty$.

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

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

In any period t, the agent's actual consumption C_t^i *in EQULIBRIUM* is bounded by the available dividends D_t. Thus, for any $T < \infty$ the probability that $\left\{D_t \le \overline{C}\right\}_{t=0}^{T}$ in equilibrium can be brought arbitrarily close to one by choosing $\overline{C}$ sufficiently high. Therefore, with arbitrarily high probability the agent's first order conditions in $t = 1, ..., T$ are given by

$$P_t < \delta^i E_t^{\mathcal{P}^i}\left[U'_{\overline{C}}(C_{t+1}^i)\left(P_{t+1}+D_{t+1}\right)\right] \qquad \text{and } S_t^i = \overline{S} \tag{37}$$

$$P_t = \delta^i E_t^{\mathcal{P}^i}\left[U'_{\overline{C}}(C_{t+1}^i)\left(P_{t+1}+D_{t+1}\right)\right] \qquad \text{and } S_t^i \in \left[0, \overline{S}\right] \tag{38}$$

$$P_t > \delta^i E_t^{\mathcal{P}^i}\left[U'_{\overline{C}}(C_{t+1}^i)\left(P_{t+1}+D_{t+1}\right)\right] \qquad \text{and } S_t^i = 0 \tag{39}$$

Since agents' beliefs satisfy (7) and assign zero probability to negative dividends and prices, we have from Lebegue's Dominated Convergence Theorem

$$\lim_{\overline{C}\to\infty} E_t^{\mathcal{P}^i}\left[U'_{\overline{C}}(C_{t+1}^i)\left(P_{t+1}+D_{t+1}\right)\right] = E_t^{\mathcal{P}^i}\left[\left(P_{t+1}+D_{t+1}\right)\right] \tag{40}$$

This implies that for $\overline{C} \to \infty$ the first order conditions (37)-(39) approximate with arbitrary precision the first order conditions (8) 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 $j \ge 1$ 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

$$\sum_{j=1}^{T} \delta^j\ D_{t+j}(\omega_D) = \sum_{j=1}^{T} \delta^j a(\omega_D)^j \prod_{\tau=1}^{T} \eta_{t+\tau}(\omega_D)\ D_t(\omega_D^t) \tag{41}$$

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

$$\begin{aligned}
& E^{\mathcal{P}}\left(\lim_{T\to\infty}\sum_{j=1}^{T}\delta^j\ D_{t+j}(\omega_D)\middle|\,\omega_D^t\right) \\
&= \lim_{T\to\infty} E^{\mathcal{P}}\left(\sum_{j=1}^{T}\delta^j\ D_{t+j}(\omega_D)\middle|\,\omega_D^t\right) && (42)\\
&= \lim_{T\to\infty} E^{\mathcal{P}}\left(\sum_{j=1}^{T}\delta^j a(\omega_D)^j\ \prod_{\tau=1}^{j}\eta_{t+\tau}(\omega_D)\ D_t(\omega_D^t)\right) && (43)
\end{aligned}$$

$$\begin{aligned}
&= \lim_{T\to\infty} D_t(\omega_D^t)\int_0^\infty\left(\sum_{j=1}^{T}\delta^j\ (\widetilde{a})^j\right) Post_t(\widetilde{a})d\widetilde{a} && (44)\\
&\geq \lim_{T\to\infty} D_t(\omega_D^t)\int_{\delta^{-1}}^\infty\left(\sum_{j=1}^{T}\delta^j\ (\widetilde{a})^j\right) Post_t(\widetilde{a})d\widetilde{a}\\
&\geq \lim_{T\to\infty} D_t(\omega_D^t)\cdot T\cdot\int_{\delta^{-1}}^\infty Post_t(\widetilde{a})d\widetilde{a}\\
&= \infty
\end{aligned}$$

where the first equality uses Lebesgue's monotone convergence theorem, the second the expression (41), 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 $\delta\widetilde{a} > 1$ over the considered range of integration. The last equality uses $\int_{\delta^{-1}}^\infty Post_t(\widetilde{a})d\widetilde{a} > 0$.

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:

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

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 (41) 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

$$\begin{aligned}
E^{\mathcal{P}}\left(\lim_{n\to\infty}\sum_{j=1}^{T}\delta^j D_{t+j}\middle|\,\omega_D^t\right) &= \lim_{T\to\infty} E^{\mathcal{P}}\left(\sum_{j=1}^{T}\delta^j\ D_{t+j}(\omega_D)\middle|\,\omega_D^t\right)\\
&= \lim_{T\to\infty} D_t(\omega_D^t)\int_0^\infty\left(\sum_{j=1}^{T}\delta^j\,(\widetilde{a})^j\right)Post_t(\widetilde{a})d\widetilde{a}\\
&= D_t(\omega_D^t)\int_0^\infty\left(\lim_{T\to\infty}\sum_{j=1}^{T}\delta^j\,(\widetilde{a})^j\right)Post_t(\widetilde{a})d\widetilde{a}\\
&= D_t(\omega_D^t)E^{\mathcal{P}}\left(\frac{\delta a}{1-\delta a}\middle|\,\omega_D^t\right) \qquad (45)
\end{aligned}$$

The second equality follows from (42)-(44), the third from applying dominated convergence once more, and the last equality uses the definition of posterior. This proves the second part of the proposition.

Next, we prove the third claim. It follows from (45) that for any $B<\delta^{-1}$ we have

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

Marcet and Nicolini (2003) show that

$$\int_0^B\frac{\delta\widetilde{a}}{1-\delta\widetilde{a}}Post_t(\widetilde{a})\ d\widetilde{a}\to\infty\quad\text{as } B\to\delta^{-1}$$

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 (2009): "Learning and Stock Market Volatility," *Mannheim University Mimeo.*

Anderson, R., and H. Sonnenschein (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," Coumbia 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.

HARRISON, J. M., AND D. M. KREPS (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323–336.

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.

MARCET, A., AND J. P. NICOLINI (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93, 1476–1498.

MARCET, 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.