

Internal Rationality, Imperfect Market Knowledge and Asset Prices¹

Klaus Adam
Mannheim University and CEPR

Albert Marcet
London School of Economics, CEP and CEPR

November 2010

¹Thanks go to Andy Abel, Margaret Bray, Gaetano Gaballo, Katharina Greulich, Seppo Honkapohja, Erzo Luttmer, Ramon Marimon, Kris Nimark, Galo Nuno, Athanasios Orphanides, Bruce Preston, Iván Werning, Mike Woodford and, specially, to Juan Pablo Nicolini for interesting comments and suggestions. Comments from the associate editor and three referees were extremely useful. We gratefully acknowledge financial support from Fondation Banque de France. Marcet also acknowledges support from CIRIT (Generalitat de Catalunya), Plan Nacional project ECO2008-04785/ECON (Ministry of Science and Education, Spain). Author contacts: Klaus Adam, adam@uni-mannheim.de; Albert Marcet a.marcet@lse.ac.uk.

Abstract

We present a decision theoretic framework in which agents are learning about market behavior and that provides microfoundations for models of adaptive learning. Agents are ‘internally rational’, i.e., maximize discounted expected utility under uncertainty given dynamically consistent subjective beliefs about the future, but agents may not be ‘externally rational’, i.e., may not know the true stochastic process for payoff relevant variables beyond their control. This includes future market outcomes and fundamentals. We apply this approach to a simple asset pricing model and show that the equilibrium stock price is then determined by investors’ expectations of the price and dividend in the next period, rather than by expectations of the discounted sum of dividends. As a result, learning about price behavior affects market outcomes, while learning about the discounted sum of dividends is irrelevant for equilibrium prices. Stock prices equal the discounted sum of dividends only after making very strong assumptions about agents’ market knowledge.

Keywords: learning, internal rationality, consumption based asset pricing

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

1 Motivation

The rational expectations hypothesis (REH) places enormous demands on agents' knowledge about how the market works. For most models it implies that agents know exactly what market outcome will be associated with any possible contingency that could arise in the future.¹ This appears utterly unrealistic given that state contingent markets that could provide agents with such detailed information often fail to exist.

The objective of this paper is to present a rigorous decision-theoretic setup that allows to relax these strong informational assumptions about how the market works and that is useful for modeling learning about market behavior by agents. As we show, relaxing these informational assumptions can have important implications for model behavior.

The basic idea is to separate the standard rationality requirements embedded in the REH into an 'internal' and an 'external' rationality component. *Internal rationality* requires that agents make fully optimal decisions given a well defined system of subjective probability beliefs about payoff relevant variables that are beyond their control or 'external', including prices. *External rationality* postulates that agents' subjective probability belief equals the objective probability density of external variables as they emerge in equilibrium.

We propose to relax the external rationality assumption but to fully maintain internal rationality in a model with well specified microfoundations. This reflects the basic conviction that internal rationality is a good starting point for analyzing social interactions. As we show, however, internal rationality is *not* sufficient to achieve external rationality. Specifically, internally rational agents can not simply derive the equilibrium distribution of market prices through a deductive reasoning process. The REH is thus *not* a consequence of optimal behavior at the individual level. Instead, to achieve external rationality one typically needs to endow internally rational agents with a lot of additional information about the market.

While we propose to relax external rationality, we suggest at the same time to consider *small* deviations from the external rationality assumption that is embedded in the REH. Specifically, we consider agents who entertain subjective beliefs that are not exactly equal to the objective density of external variables but that will be close to the beliefs that an agent would entertain under the REH. This amounts to relaxing the 'prior beliefs' that agents are assumed to entertain under the REH and to study the economic implications of such a relaxation.

Doing so requires changing the microfoundations of our standard

¹Exceptions are models with private information, see section 7.

models. Specifically, it requires enlarging the probability space over which agents condition their choices, and including all payoff-relevant *external* variables, i.e., all variables that agents take as given. This includes (competitive) market prices. This departs from the standard formulation in the literature where agents' probability space is reduced from the outset to contain only *exogenous* (or 'fundamental') variables with prices being excluded from the probability space. In the standard formulation this is possible because prices are assumed to be a function of exogenous fundamentals, and the equilibrium pricing functions are *assumed* to be known to agents.² The standard procedure thus imposes a singularity in the joint density over market prices and fundamentals, with the singularity representing agents' exact knowledge about how prices are linked to fundamentals. It also implies that market outcomes carry only redundant information, so that agents do not need to condition on prices to behave optimally.

Assuming the existence of a singularity in agents' joint beliefs about prices and fundamentals, however, appears to be in stark contrast with what academic economists seem to know about the relation between prices and the observed history of fundamentals in the real world. This manifests itself in the fact that empirical economists often fail to agree on a dominant explanation for market price behavior and entertain competing models and explanations. In contrast to this, agents in RE models have reached an agreement on the correct model for the market price in period zero already. The existing uncertainty by expert economists suggests, however, to endow agents in our models with similar uncertainty about how prices link with fundamentals.

We do so by allowing agents to entertain a non-degenerate joint density over future prices and dividends, so that optimizing agents naturally need to condition decisions also on price realizations. Even though this is a potentially small departure from RE beliefs, we show that the model outcome can be quite different.

The literature on adaptive learning previously studied models in which agents learn about how to forecast future market outcomes. This literature, however, makes a number of ad-hoc assumptions about agents' behavior and learning mechanisms.³ As a result, the microfoundations of adaptive learning models have not been carefully laid out, and it is unclear to what extent agents in these models take rational decisions given the information they are assumed to possess.⁴ This generates

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

³We discuss these in detail in section 2 below.

⁴For example, the adaptive learning literature appeals to anticipated utility max-

controversy, specially in applications of models of learning to empirical work or for policy analysis, as is the case in an increasing number of contributions.⁵

Our approach can be used to provide microfoundations to models of adaptive learning. Similar to Muth (1961), who showed how adaptive expectations can be compatible with the REH, we demonstrate how ordinary least squares learning - a widely assumed learning rule in the adaptive learning literature - arises as the optimal way to update conditional expectations from a complete and dynamically consistent set of probability beliefs within a specific model.

To illustrate our approach for relaxing external rationality, we present a simple asset pricing model with risk-neutral investors. We include heterogeneous agents and standard forms of market incompleteness to insure that there exists a distinction between the agent's own decision problem, which we assume to be perfectly known, and market behavior, which we assume to be known only imperfectly.

We first show - perhaps surprisingly - that the equilibrium stock price is then determined by a *one-step ahead* asset pricing equation. More precisely, the equilibrium stock price equals the marginal investor's discounted expected sum of the total stock payoff (price plus dividend) *in the next period*. This differs from models with perfect market knowledge, where the equilibrium price equals the discounted sum of future dividends. Our one-step ahead equilibrium pricing equation implies different market outcomes because the marginal agent's expectations of tomorrow's price need not be related to the agent's expectations about future dividends. 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 - as long as the agent expects to be able to sell the stock at a higher price tomorrow.⁶ The agent may reasonably expect to be able to do so if she holds the expectation that the marginal agent tomorrow will hold more optimistic price and dividend expectations.

With imperfect market knowledge, beliefs about future prices thus become a crucial element for determining today's stock price. As a

imization in the sense of Kreps (1998), which is well-known to be *not* dynamically consistent.

⁵For example, Adam, Marcet and Nicolini (2010), Adam (2005), Chakraborty and Evans (2008), Cogley and Sargent (2008), Eusepi and Preston (2010), Marcet and Nicolini (2003), and Timmermann (1993, 1996) use adaptive learning models to explain data; Evans and Honkapohja (2003a, 2003b, 2005), Molnar and Santoro (2007), Orphanides and Williams (2006) and Sargent (1999) employ such models for policy analysis.

⁶This is so because it is optimal to engage in speculative trading in the sense of Harrison and Kreps (1978).

result, revisions in price beliefs add to the volatility of stock prices. Moreover, if agents hold the view that prices differ from the discounted sum of dividends, then actual prices will do so, thereby supporting their initial view. Nevertheless, agents' beliefs will differ from the objective probability distribution of prices. Yet, as we show, agents can not derive the objective distribution of prices through a deductive reasoning process if they just know about their own dividend beliefs. This is possible only with additional information, for example, if the preferences and beliefs of all agents are common knowledge.

Intuitively, the stock price ceases to be a discounted sum of dividends because imperfect market knowledge (or alternatively lack of common knowledge of agents' preferences and beliefs) leads to a failure of the law of iterated expectations. Since the identity of the marginal agent that actually prices the stock is changing with time and because agents entertain heterogeneous beliefs, the equilibrium price is given by expectations evaluated under different probability measures each period. As a result, agents cannot iterate forward on the one-step ahead pricing equation.⁷

A standard way to relax the strong informational assumptions underlying RE has been the concept of Bayesian rational expectations equilibrium. This literature allows for imperfect information about the density of exogenous variables (fundamentals) but it maintains the assumption of perfect knowledge about the mapping from fundamentals to prices, thus assumes the existence of a singularity in agents' beliefs over prices and dividends. Bayesian RE equilibria thus deal with uncertainty about fundamentals (dividends) and market outcomes (prices) in a rather asymmetric way: while the process for fundamentals is imperfectly known, the contingent process for prices is assumed to be known perfectly.

Studying Bayesian RE equilibria, Bray and Kreps (1987) argued that it was unclear how much information agents need to possess about the market for a Bayesian RE equilibrium to emerge.⁸ Section 4 of this paper can be interpreted as addressing this issue. In the context of our asset

⁷This feature also emerged in Allen, Morris and Shin (2006), who study an asset pricing model with imperfect common knowledge. In their setting, the one-step ahead pricing equation emerges from the underlying two-period overlapping generations framework and differential information across generations is sustained by introducing a noise trader assumption. Both features together imply that one cannot easily iterate forward on the one-step ahead pricing equations. While Allen, Morris and Shin maintain RE in a model with private information, we depart from RE (by assuming imperfect market knowledge) but derive the one-step ahead pricing equation in a setting with infinitely lived investors. Preston (2005) also points out how imperfect market knowledge prevents the law of iterated expectations from giving a discounted sum formulation of a budget constraint.

⁸This point has been discussed more recently, for example, by Marimon (1997) and Sargent (2008).

pricing model, we show that a series of strong informational assumptions provide optimizing agent with sufficient information to map the process for dividends into a single price outcome. These assumptions endow the agent with a tremendous amount of additional knowledge about the market, over and above what can be derived from internal rationality alone. Roughly speaking, the Bayesian RE equilibrium emerges if all agents possess the same information as the theorist, i.e., agents need to know all details about all other agents in the economy, including other agents' probability beliefs, discount factors and so on, and all this needs to be common knowledge.

Considering agents whose beliefs about prices are different from the actual price distribution raises a number of issues. To many economists it may seem that the choice of agents' beliefs must be arbitrary. In section 5.5 we provide a discussion of this issue, we argue that there is no arbitrariness in applications where *i)* the beliefs of agents are near-rational, *ii)* the market outcome does not contradict agents' beliefs in an obvious way and *iii)* if the modeler's assumption about agents' beliefs is made in a reasonable way. Related to this, we end the paper by showing in section 6 that the REH does not prevent arbitrary outcomes: even when a Bayesian REE emerges, the asset pricing predictions prove extremely sensitive to fine details in agents' beliefs about the dividend process. Based on this we conclude that agents' prior beliefs may matter much more than other economic factors for the behavior of equilibrium stock prices in a Bayesian RE equilibrium. The pricing implications in Bayesian RE equilibrium models thus appear more arbitrary than previously recognized.

The outline of the paper is as follows. In section 2 we present a list of unresolved issues in the adaptive learning literature. In section 3 we introduce a simple stock pricing model with incomplete markets and heterogeneous agents, we show how to introduce internal rationality, derive investors' optimality conditions, and define a competitive equilibrium with internally rational agents. Section 4 compares our equilibrium concept to Bayesian RE equilibrium and shows how agents' market knowledge needs to be strengthened enormously in order for a discounted sum of dividends and the Bayesian RE equilibrium to arise. Section 5 presents a consistent set of beliefs where agents are uncertain about the mapping from dividends to prices. It shows how to entertain small deviations from REE beliefs and how least-squares learning equations then emerge from an optimal use of information in a specific case. Section 6 presents a formal result about the strong sensitivity of the discounted sum of dividends to prior information about the dividend process. Section 7 discusses some of the related literature. A conclusion summarizes.

2 Adaptive Learning Literature: Open Issues

The adaptive learning literature relaxes agents' knowledge about the behavior of market determined variables but also makes a number of ad-hoc assumptions on agents' behavior and learning mechanisms. These give rise to important questions regarding the microfoundations of adaptive learning models.

The source of the problem is as follows: the adaptive learning literature takes as point of departure the first order optimality conditions that emerge under the REH; it then replaces the rational expectations operator E appearing in these optimality conditions by an operator of *perceived* expectations \tilde{E} ; it then assumes that agents constantly re-estimate the parameters involved in these perceived expectations in light of new data using some stochastic approximation algorithm.

One element of arbitrariness arises because first order conditions under the REH can be written in many equivalent ways. One can then replace rational expectations by the subjective operator \tilde{E} in many different equations and, it turns out, depending on which version of the RE formulation is used one can end up with rather different outcomes under learning.

Adam, Marcet and Nicolini (2010), for example, consider an asset pricing model. They use a one-step-ahead asset pricing equation $P_t = \delta \tilde{E}_t(P_{t+1} + D_{t+1})$ and show that a number of empirical stock price puzzles can be explained if agents are learning about future price behavior. By contrast Timmermann (1996) and others set the stock price equal to expected discounted sum of dividends, i.e., uses $P_t = \tilde{E}_t \sum_{j=1}^{\infty} \delta^j D_{t+j}$, and studies learning about discounted dividends, finding a much more muted impact on stock prices from learning behavior. Which is the 'right' way to set up the learning model?

Likewise, Evans and Honkapohja (2003b) have formulated DSGE models under learning using one-step-ahead Euler equations while Preston (2005) showed that learning outcomes in a monetary model differ when using the budget constraint to obtain a discounted sum formulation of the optimality conditions. Again, which is the 'right' way to set up the learning model?

Another element of arbitrariness emerges because a large number of stochastic approximation algorithms are available to formulate estimates of the parameters that determine agents' perceptions \tilde{E} . The literature has used a range of stochastic approximation algorithms, e.g., ordinary least squares learning, constant gain learning, or switching gain algorithms. Which is the 'right' way to model the response of expectations to new data?

Finally, while the perceptions \tilde{E} are constantly evolving over time, agents behave as if their current view will remain unchanged in the future, following the anticipated utility concept of Kreps (1998). It is unclear whether this way of decision making will lead to an admissible plan in the Bayesian sense, i.e., whether there exists at all a dynamically consistent subjective probability measure under which the agents' decisions resulting from this procedure are optimal.

Under the framework of this paper modeling choices are determined from rational behavior of agents and the microeconomic specifications of the agent's decision problem, including the agents' subjective beliefs about external variables. Surprisingly, it will turn out that some of the short-cuts of the adaptive learning literature are less ad-hoc than might initially appear, and we find that the one-step formulation of Adam Marcet and Nicolini (2010) under OLS is optimal in some specific models.

3 Internal Rationality with Imperfect Market Knowledge

This section introduces the concept of internal rationality, shows how to define agents' probability space and defines and characterizes the competitive equilibrium with internal rationality.

To illustrate our approach we study a risk-neutral asset pricing model with heterogeneous agents and incomplete markets. We choose such a model for its simplicity and because we obtain very different pricing implications from the standard case with perfect market knowledge.

Agents in our model differ in their discount factor and in their subjective beliefs. Markets are incomplete because of the existence of constraints that limit the amount of stocks investors can buy or sell and because contingent claim markets are unavailable. The presence of investor heterogeneity and market incompleteness allows us to distinguish between investors' knowledge of their own decision problem and their knowledge about market-determined variables, i.e., future asset prices, which are also influenced by the discount factors and beliefs of other (possibly different) investors.

3.1 Basic Asset Pricing 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 $1/I$ units of an infinitely lived stock. Agents of type $i \in \{1, \dots, I\}$ have a standard time-separable

utility function

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

where C_t^i denotes consumption at t 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}) . As usual, the probability measure \mathcal{P}^i is a model primitive and given to agents. It is allowed to be type-specific and, due to imperfect market knowledge, it may or may not coincide with objective probabilities. The stocks S_t^i owned by agents represent claims to an infinitely lived tree that yields each period D_t units of a perishable consumption good which are paid as dividend.

The non-standard part in our formulation is in the underlying probability space. 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 $\Omega_X = \prod_{t=0}^{\infty} R_+$ with $X \in \{P, D\}$. The probability space thus contains all possible sequences of prices and 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 (Ω, \mathcal{S}) . As usual we denote the set of all possible dividend histories up to period t by Ω_D^t and we let $D^t \in \Omega_D^t$ denote a typical dividend history. Using similar definitions for prices, the set of all histories up to period t is given by $\Omega^t = \Omega_P^t \times \Omega_D^t$ and its typical element is denoted by $\omega^t \in \Omega^t$.

With this setup rational investors will 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 stochastic variable that is beyond their control and since prices influence their budget constraint, investors want to condition their choices on the realization of prices, in addition to the realization of dividends.

Note that we have endowed agents with a dynamically consistent set of subjective beliefs, i.e., $(\Omega, \mathcal{S}, \mathcal{P}^i)$ is a proper probability space, \mathcal{P}^i satisfies all the standard probability axioms and gives proper joint probabilities to all possible values of prices and dividends in any set of dates. Moreover, although there is a time-invariant probability measure \mathcal{P}^i , our setup is general enough to allow for agents that are learning about the stochastic processes of prices and dividends. For example,

\mathcal{P}^i could arise from a view that agents entertain about the stochastic processes describing the evolution of prices and dividends and by some prior beliefs about unknown parameters of these processes. A particular example of this kind of subjective beliefs will be given in section 5.1.

Investors of type i choose consumption and stock holdings in period t , denoted (C_t^i, S_t^i) , contingent on the observed history $\omega^t = (P^t, D^t)$, i.e., investors choose a function

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

for all t . The expected utility (1) associated with any such contingent consumption choice can then be written as

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

The stock can be purchased and sold costlessly in a perfectly competitive spot market at ex-dividend price P_t . Agent i thereby faces the following flow budget constraint

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

which has to hold for all t and all $\omega^t \in \Omega^t$. Here ξ denotes a sufficiently large endowment of consumption goods, which is introduced for simplicity: it allows us to ignore non-negativity constraints on consumption.⁹

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

$$S_t^i \geq 0 \quad (5)$$

$$S_t^i \leq \bar{S} \quad (6)$$

where $1 < \bar{S} < \infty$. 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) helps to insure existence of a maximum in the presence of risk neutral investors.

We are now in a position to define internal rationality within the current setting:

Definition 1 (Internal Rationality) *Agent i is internally rational if she chooses the functions (2) to maximize expected utility (3) subject to the budget constraint (4), and the limit constraints (5) and (6), taking as given the probability measure \mathcal{P}^i .*

⁹No substantial result depends on the fact that the non-negativity constraint on consumption is not binding.

For more general settings, internal rationality requires that agents maximize their objective function taking into account all relevant constraints, that they condition their actions on the history of all observable external variables, and that they evaluate the probability of future external outcomes using a consistent set of subjective beliefs, which is given to them from the outset.

Within the context of the present model we assume that \mathcal{P}^i satisfies

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

and that a maximum of the investor's utility maximization problem exists.¹⁰

3.2 Optimality Conditions

Under internal rationality the space of outcomes Ω considered by agents includes all external variables, i.e., the histories of prices and the history of dividends. Agents can thus assign a consistent set of probabilities to all payoff relevant external events. Consequently, the first order optimality conditions are found in a standard way. In particular, one of the following conditions has to hold for all periods t and for almost all realizations in $\omega^t \in \Omega^t$:

$$P_t < \delta^i E_t^{\mathcal{P}^i}(P_{t+1} + D_{t+1}) \quad \text{and } S_t^i = \bar{S} \quad (8a)$$

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

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

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 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 without transaction costs, the one-step-ahead optimality conditions (8) deliver optimal investment choices, even if stocks can be held for an arbitrary number of periods.

¹⁰Appendix A.1 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.

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! Intuitively, our agents simply try to 'buy low and sell high' as much as the stock holding constraints allow them. This is the optimal strategy because it is optimal for agents to engage in speculative behavior in the sense of Harrison and Kreps (1978).

We show below that with imperfect market knowledge, an agent's expectations of the future price is not determined by the agent's dividend expectations and internal rationality. The first order conditions above, therefore, turn out to be equivalent to a discounted sum of dividend formulation only in very special cases.

3.3 Standard Belief Formulation: A Singularity

The setup for beliefs defined in the previous section differs from standard dynamic economic modeling practice, which imposes *additional restrictions* on beliefs. Specifically, the standard belief specification assumes that agents formulate probability beliefs only over the reduced state space Ω_D and that agents' choices are contingent on the history of dividends only. Agents are then endowed with the knowledge that each realization $D^t \in \Omega_D^t$ is associated with a given level of the stock price P_t , which amounts to endowing agents with knowledge of a function

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

The probabilities for the price process are then constructed from knowledge of this function and beliefs over Ω_D . Once this function is observed prices carry only redundant information, there is no need to condition choices on the history of prices and there is no loss in optimality by excluding prices from the state space. Clearly, knowledge of the function (9) represents knowledge regarding market outcomes: agents know exactly which market outcome is going to be associated with a particular history of fundamentals.

This 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 (P^t, D^t) . In contrast, our more general belief formulation outlined in section 3.1 allows agents to be uncertain about the relation between prices and dividends.

The standard formulation using degenerate beliefs is consistent with the rational expectations equilibrium outcome, so no loss of generality is implied by imposing the singularity in \mathcal{P}^i from the outset under the REH. But as we will show in sections 3.5 and 4 below, knowledge of this singularity is not a consequence of agents' ability to maximize their

utility or to behave rationally given their subjective beliefs. Instead, it is the result of a set of strong assumptions that imply that agents know from the outset how the market works. Indeed, a sufficient condition for agents to work out the equilibrium price function will be that agents know the market so well that they are able to map each potential future dividend sequence into a single value for the stock price. Given that such a relationship between dividends and prices remains fairly elusive to academic economists - these still entertain a range of alternative asset pricing models each of which implies a different function P_t - it seems equally reasonable to consider agents who are also not fully certain about the map linking dividends to prices. Imperfect knowledge about market behavior is thus naturally modeled by allowing agents to entertain beliefs about the joint process for prices and dividends that does not impose a singularity.

3.4 Internally Rational Expectations Equilibrium (IREE)

This section considers the process for competitive equilibrium prices with internally rational agents and defines an Internally Rational Expectations Equilibrium (IREE).

We propose a competitive equilibrium definition that is as close as possible to the standard formulation. The definition below is specific to our stock pricing model but is easily extended to more general setups. Let $(\Omega_D, \mathcal{S}_D, \mathcal{P}_D)$ be a probability space with Ω_D denoting the space of dividend histories and \mathcal{P}_D the ‘objective’ probability measure for dividends. Let $\omega_D \in \Omega_D$ denote a typical infinite history of dividends.

Definition 2 (IREE) *An Internally Rational Expectations Equilibrium (IREE) consists of a sequence of equilibrium price functions $\{\mathbf{P}_t\}_{t=0}^\infty$ where $\mathbf{P}_t : \Omega_D^t \rightarrow R_+$ for each t , contingent choices $\{C_t^i, S_t^i\}_{t=0}^\infty$ of the form (2) and probability beliefs \mathcal{P}^i for each agent i , such that*

- (1) *all agents $i = 1, \dots, I$ are internally rational, and*
- (2) *when agents evaluate $\{C_t^i, S_t^i\}$ at equilibrium prices, markets clear for all t and all $\omega^D \in \Omega_D$ almost surely in \mathcal{P}_D .*

Verbally, an IREE is a competitive equilibrium allowing for the possibility that agents’ subjective density about future prices and dividends is not necessarily equal to the objective density. Or equivalently, it is an equilibrium in which agents are internally rational but not necessarily externally rational.

Quite a few papers have previously studied Arrow-Debreu (AD) models in which agents’ subjective probability densities about fundamentals

may not coincide with the actual densities of the fundamentals.¹¹ It is important to note that an IREE is not a special case of this literature. The reason is that in the AD framework embodies two basic features: i) any physical good is treated as a different good if delivered in a different period or for a different realization; ii) agents observe the equilibrium prices for all goods. These two features together imply that the equilibrium price function (9) is known to agents. We consider cases where the singularity in beliefs is absent because there does not exist a full set of contingent claim markets.

We now determine the equilibrium price mappings \mathbf{P}_t in the above asset pricing model. Equilibrium prices will depend on standard micro-economic fundamentals such as utility functions, discount factors, and dividend beliefs, but also on agents' price beliefs given by the probability measures \mathcal{P}^i . Moreover, since agents do not necessarily hold rational price expectations, we need to distinguish between the stochastic process for equilibrium prices \mathbf{P}_t and agents' perceived price process P_t . The first order conditions (8) imply that the asset is held by the agent type with the most optimistic beliefs about the discounted expected price and dividend *in the next period*.¹² Equilibrium prices thus satisfy:¹³

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

The next section discusses, whether agents could deduce the equilibrium price function (10) from the information that is available to them.

3.5 Is Internal Rationality Sufficient to Derive a Singularity in Beliefs?

The equilibrium price function $\mathbf{P}_t : \Omega_D^t \rightarrow R_+$ emerging in an IREE is indeed a function of the history of dividends only. This implies that the objective density over prices and dividends features a singularity. In light of these observations it is natural to ask whether knowledge that this singularity exists would be sufficient to allow internally rational agents to compute the correct equilibrium price functions through a process

¹¹See Blume and Easley (2006) for a recent application.

¹²This emerges because we assume $\bar{S} > 1$ so that the constraint (6) never binds in equilibrium. Extensions to the case with $\bar{S} < 1$ are straightforward.

¹³Since expectations $E_t^{\mathcal{P}^i}$ are conditional on the realization P_t , the equilibrium price affects both sides of the expression above and, at this level of generality, it is unclear whether there always exists an equilibrium price P_t for any given dividend history D^t or whether it is unique, see also the discussion in Adam (2003). At this point, we proceed by simply assuming existence and uniqueness, we leave this issue for further research. See footnote 21 for a discussion of existence and uniqueness in the specific asset pricing model that we consider.

of deductive reasoning. Or equivalently, does internal rationality imply external rationality if agents know that dividends are the only source of fundamental disturbances?

The answer to both of these questions turns out to be ‘no’. As we show below, the problem is that knowledge of the existence of a degeneracy falls short of informing agents about its exact location. This holds true even if the equilibrium asset pricing equation (10) is common knowledge to all agents. This in turn provides a natural interpretation for why agents’ beliefs might *not* contain a singularity, even though in the model the objective density possesses a singularity: agents are simply uncertain about the correct model linking stock prices to the history of dividends, and they express this uncertainty using a non-degenerate system of beliefs \mathcal{P}^i over prices and dividends.

We now show that the singularity is not easily located. Let $m_t : \Omega_D^t \rightarrow \{1, \dots, I\}$ denote the marginal agent pricing the asset in period t in equilibrium:¹⁴

$$m_t = \arg \max_{i \in I} \left[\delta^i E_t^{\mathcal{P}^i} (P_{t+1} + D_{t+1}) \right] \quad (11)$$

Clearly the equilibrium price (10) can thus be written as

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

We now suppose that agents know that the equilibrium price satisfies equation (12) each period and that this is common knowledge.¹⁵ Doing so endows agents with a considerable amount of information about how the market prices the asset. Specifically, common knowledge implies that each agent knows that other agents know that the asset is priced according to (12) each period, that each agent knows that other agents know that others know it to be true, and so on to infinity.¹⁶ We can express this formally by saying that from the agents’ viewpoint the following equation holds

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

and that each agent has price and dividend beliefs \mathcal{P}^i that are consistent with this equation. The question we are posing is: would common

¹⁴If the argmax is non-unique we can use a selection criterion from among all marginal agents. For example, we can take m_t to be the marginal agent with the lowest index i .

¹⁵Internally rational agents do not need to have such knowledge to behave optimally conditional on their beliefs.

¹⁶See Aumann (1976) for a formal definition.

knowledge of equation (13) allow internally rational agents to impose restrictions on price beliefs as a function of their beliefs about dividends? Would it allow agents to determine a singularity?

Common knowledge of equation (13) allows agents to iterate forward on this equation, say T times, to find

$$\begin{aligned}
P_t &= \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (D_{t+1}) \\
&+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} D_{t+2} \right) \\
&+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+3} \right) \right) + \dots \\
&+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\dots \delta^{m_{t+T}} E_{t+T}^{\mathcal{P}^{m_{t+T}}} (P_{t+T+1} + D_{t+T+1}) \right) \right)
\end{aligned} \tag{14}$$

The last three lines of the right-hand side of this equation provide an alternative expression for agents' discounted expectations of next period's price. They show that knowledge of (13) implies that agents' price expectations are given by their 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 future dividends and the terminal price. Since agent i is not marginal in all periods and since agent i can rationally believe other agents to hold rather different beliefs, own beliefs about dividends fail to restrict the beliefs agent i can entertain about prices. For example, agent i can believe the future discounted sum of dividends to be low but at the same time believe the future price to be high - all that is required is that the agent believes future marginal agents to be relatively more optimistic about future dividends and prices.

In the literature, the discounted sum of dividends is usually obtained by applying the law of iterated expectations on the right side of equation (14). This can be done whenever all conditional expectations are with respect to the same probability measure, e.g., if m_t is constant through time. In our model m_t is random whenever \mathcal{P}^i assigns positive probability to the event that the agent may not be marginal at some point in the future. If in addition the agent believes that other agents hold different (price and dividend) beliefs, then the law of iterated expectations can not be applied to (14).¹⁷ Price expectations then fail to be determined by agents' dividend expectations. This shows that own dividend beliefs, knowledge of (13), and internal rationality are not sufficient conditions for agents' beliefs P^i to contain a specific singularity where prices are equal to a discounted sum of dividends. The next section explores this issue further by actually providing sufficient additional conditions un-

¹⁷Allen, Morris and Shin (2006) and Preston (2005) make a similar point in different models.

der which internally rational agents would have to incorporate such a singularity.

4 Bayesian Rational Expectations Equilibrium

This section derives sufficient conditions that would allow internally rational agents to impose the ‘correct’ singularity in their subjective beliefs \mathcal{P}^i over prices and dividends, as it emerges in equilibrium. Specifically, we show that to be able to deduce the correct equilibrium pricing function \mathbf{P}_t , agents require a tremendous amount of information about the market. This confirms conjectures expressed previously by Bray and Kreps (1987) regarding the strong informational requirements underlying Bayesian REE models.

We start by providing a definition of a Bayesian REE. While our definition is stated in terms of our previous definition of an IREE, the resulting equilibrium notion nevertheless agrees with that provided in most of the literature.

Definition 3 (Bayesian REE) *A Bayesian Rational Expectations Equilibrium is an Internally Rational Expectations Equilibrium in which agents’ subjective beliefs \mathcal{P}^i are consistent with the equilibrium price function \mathbf{P}_t , i.e.*

$$\text{Prob}^{\mathcal{P}^i}(P_t = \mathbf{P}_t \mid D^t) = 1$$

for all t, ω, i .

Verbally, a Bayesian REE is an IREE in which all agents associate with each possible partial dividend history the correct equilibrium price. The term ‘Bayesian’ in this definition reflects the fact that agents’ knowledge about the dividend process is allowed to be imperfect. For example, agents may be uncertain about some of the parameters in the law of motion of dividends. When all agents know the true process for dividends, then the Bayesian REE simplifies further to a standard REE.

We now provide sufficient conditions on \mathcal{P}^i so that the IREE reduces to a Bayesian REE. As in the previous section, we start by endowing agents with knowledge of how the *market* prices the asset for all periods t and all states ω :

Assumption 1 It is common knowledge that equation (13) holds for all t and all $\omega \in \Omega$.

This allows agents to iterate on the equilibrium asset price equation (13) to obtain equation (14). Importantly, agents can not iterate on their *own* first order optimality conditions because these do not always hold with equality.

The discounted sum expression (14) still contains expectations about the terminal price 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

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

for all t and all $\omega \in \Omega$.

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:

$$\lim_{T \rightarrow \infty} E_t^{\mathcal{P}} (\delta^T P_{t+T}) = 0 \quad (15)$$

Note, that even this more familiar ‘no-rational-bubble’ condition endows agents with knowledge of how the *market* prices the asset asymptotically, it does not just arise from rational behavior.

Assumption 2 allows to take the limit $T \rightarrow \infty$ in equation (13) and to abstract from expectations about the terminal selling price to obtain:

$$\begin{aligned} \mathbf{P}_t &= \delta^{m_t} E_t^{\mathcal{P}^{m_t}} (D_{t+1}) \\ &+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} D_{t+1} \right) \\ &+ \delta^{m_t} E_t^{\mathcal{P}^{m_t}} \left(\delta^{m_{t+1}} E_{t+1}^{\mathcal{P}^{m_{t+1}}} \left(\delta^{m_{t+2}} E_{t+2}^{\mathcal{P}^{m_{t+2}}} D_{t+2} \right) \right) \\ &+ \dots \end{aligned} \quad (16)$$

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, and so on. Agents may, however, still entertain a range of views about who will be marginal in the future and what the dividend expectations of such marginal agents are going to be. Assumptions 1-2 are thus still not sufficient for rational agents to associate a single equilibrium price with each dividend history $D^t \in \Omega_D^t$.

For equation (16) to impose a singularity, agents have to believe in a given mapping $m_t : \Omega_D^t \rightarrow \{1, \dots, I\}$ and they must know the discount factor δ^i and the probability measure \mathcal{P}^i for all other agents i . Only then can a rational agent use equation (16) and the own beliefs about

the dividend process to evaluate the right side of (16), i.e., can associate a single price outcome with any dividend history D^t .

Furthermore, in a Bayesian REE the resulting price beliefs must be objectively true given the dividend history. This fails to be the case if agents employ an arbitrary mapping m_t . Therefore, agents must employ the mapping m_t that is objectively true in equilibrium! Letting $\mathbf{m}_t : \Omega_D^t \rightarrow \{1, 2, \dots, I\}$ denote this equilibrium mapping, we need

Assumption 3 The equilibrium functions \mathbf{m}_t for all t , the discount factors δ^i and the probability measures \mathcal{P}^i for all i are known to all agents.

Clearly, Assumption 3 incorporates a tremendous amount of knowledge about the market: agents need to know for each possible dividend history which agent is marginal, what is the marginal agent's discount factor, and the marginal agent's belief system. Only then can agents impose the correct singularity (16) on their joint beliefs about the behavior of prices and dividends.

The simplest and most common way in the literature to impose Assumptions 1-3 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 (15). If the representative agent knows that she is marginal at all times and contingencies, her first order condition holds with equality at all periods. She can then iterate on it and evaluate future expectations by applying the law of iterated expectations to own beliefs. In this specific case, internal rationality (plus assumption 2) implies equality between the 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 a natural outcome of internally rational investment behavior, but clearly this is not true more generally. Indeed, the commonly made assumption that the agent 'knows to be marginal at all times' is an indirect way to make Assumption 3, i.e., an assumption that appears rather unnatural when allowing for heterogeneity and incomplete markets. Since homogeneous agent models are best thought of as rough approximations to heterogeneous agent models, the assumption that agents know that they are marginal at each time should be as unappealing as making Assumption 3.

The fact that strong assumptions have to be imposed for the REH to emerge raises the important question of how agents could have possibly acquired such detailed knowledge about the working of the market? Given that equilibrium prices do not even come close to revealing the

underlying process for market microeconomic fundamentals (\mathbf{m}_t , δ^i and \mathcal{P}^i), it is hard to see how an agent could possibly be certain from the outset about the relation between dividends and prices. Given this, it seems to us worthwhile to pursue the concept of IREE, where agents do not know the equilibrium pricing function from the outset.

5 Asset Pricing with Imperfect Market Knowledge

This section presents a specific example showing how one can slightly relax the strong market knowledge assumptions underlying a Bayesian REE. The example is of interest because it shows - perhaps surprisingly - that for some models the standard approach taken in the adaptive learning literature, as discussed in section 2, can be consistent with internal rationality. Specifically, we show that the asset pricing model in Adam, Marcet and Nicolini (2010), which uses a one-step-ahead pricing equation and replaces the expectations operator in this equation by a least squares learning algorithm, can be derived from a model with internally rational agents whose prior beliefs are close to the RE beliefs. This is important because the learning model explored in Adam, Marcet and Nicolini gives rise to equilibrium prices dynamics that quantitatively replicate a wide range of asset pricing facts within a very simple setup.

The model below abstracts from heterogeneity amongst agents and considers instead a model with homogenous agents. Heterogeneity was useful in the previous section to highlight that in realistic models a huge amount of market knowledge is required for agents to deduce market outcomes and for a Bayesian REE to arise, but heterogeneity is not crucial for the characterization in this section. All we require is that homogeneity amongst agents fails to be common knowledge, so that agents cannot deduce the market outcome from what they know.¹⁸

We start by determining the REE, then show how one can relax slightly the singularity in prior beliefs that agents are assumed to entertain in the REE. Finally, we show how Bayesian learning about the price process gives rise to the ordinary least squares (OLS) learning equations assumed in Adam, Marcet and Nicolini (2010).

¹⁸Alternatively, the homogeneous agent model below could be interpreted as an approximation to the solution of a heterogeneous agent model in which the degree of heterogeneity is vanishing but where vanishing heterogeneity fails to be common knowledge.

5.1 Perfect Knowledge Benchmark (REE)

We consider risk neutral agents who share the same beliefs and the same discount factor δ . The true process for dividends is assumed to follow

$$\log D_t/D_{t-1} = \log a + \log \varepsilon_t \quad (17)$$

with $a > 0$, $\log \varepsilon_t \sim iin(0, \sigma^2)$ and $D_{-1} > 0$ given. Log dividends thus grow at the rate $\log a$ on average and dividend growth innovations are unpredictable. When the dividend process (17) and homogeneity of agents is common knowledge, i.e., if agents know all relevant features of other agents in the market, then internal rationality implies that the market equilibrium is given by the REE outcome, i.e.,

$$P_t^{RE} = \frac{\delta a e^{\sigma^2/2}}{1 - \delta a e^{\sigma^2/2}} D_t$$

The RE equilibrium price process thus evolves according to

$$\log P_t^{RE}/P_{t-1}^{RE} = \log a + \log \varepsilon_t \quad (18)$$

so that prices grow at the same rate as dividends. The stochastic innovation in the price growth process is thereby the same as in the dividend growth process, illustrating the existence of a singularity in the joint evolution of prices and dividends. While it is well known that these aspects of the REE solution are empirically unappealing, our discussion in section 3 about market knowledge suggests that they may be equally unappealing on theoretical grounds. The next section relaxes agents' knowledge about the stochastic processes (17) and (18).

5.2 Imperfect Knowledge: Relaxing REE Priors

We now relax the assumption that homogeneity of agents is common knowledge. Instead we assume that agents have imperfect knowledge about other agents' preferences and beliefs. The discussion in sections 3 and 4 show that this allows us to consider internally rational agents that hold subjective beliefs \mathcal{P} which differ from the ones they are assumed to entertain in the REE. We also show how one can choose agents' beliefs to be arbitrarily close to REE beliefs.

Specifically, we assume that agents believe prices and dividends to evolve according to the following process

$$\begin{bmatrix} \log P_t/P_{t-1} \\ \log D_t/D_{t-1} \end{bmatrix} = \begin{bmatrix} \log \beta^P \\ \log \beta^D \end{bmatrix} + \begin{bmatrix} \log \varepsilon_t^P \\ \log \varepsilon_t^D \end{bmatrix} \quad (19)$$

for given (P_{-1}, D_{-1}) and with

$$(\log \varepsilon_t^P, \log \varepsilon_t^D)' \sim iiN(0, \Sigma)$$

$$\Sigma = \begin{bmatrix} \sigma_P^2 & \sigma_{PD} \\ \sigma_{PD} & \sigma_D^2 \end{bmatrix}$$

This specification allows for different growth rates of prices and dividends and for innovations to prices and dividends that are only imperfectly correlated. Unlike section 5.1 we now consider agents who 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). We capture agents' uncertainty at time zero by prior beliefs about these unknown parameters and summarize these by a probability density function

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

The prior beliefs f together with the laws of motion (19) fully determine agents' probability measure \mathcal{P} over infinite sequences of price and dividends realizations.¹⁹

The previous system of beliefs gives rise to the beliefs that agents entertain in the REE in the special case when the prior f assigns probability one to the outcome

$$\beta^P = \beta^D = a, \quad \Sigma = \sigma^2 \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

We call this the 'RE prior' and let \mathcal{P}^{RE} denote the associated probability measure over sequences of prices and dividends. The singularity in this measure shows up in the form of a singular covariance matrix Σ .

We now relax these RE priors slightly. The relaxation gives rise to an alternative probability measure \mathcal{P} without a singularity in the joint density over prices and dividends. We also explain what we mean by a 'small deviation from REE beliefs'.

In the interest of obtaining closed form solutions for the evolution of the posterior beliefs, we use a conjugate prior specification for f that is of the Normal-Wishart form. Specifically, we consider prior beliefs of the form

$$H \sim W(S_0, n_0) \tag{20a}$$

$$(\log \beta^P, \log \beta^D)' \Big| H = h \sim N \left((\log \beta_0^P, \log \beta_0^D)', (\nu_0 h)^{-1} \right) \tag{20b}$$

¹⁹Given this structure, the probabilities assigned by \mathcal{P} can be obtained as follows: for any Borel subset $s \subset \mathcal{S}$, determine the likelihood 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 (19). Then integrate these probabilities over values of $(\log \beta^P, \log \beta^D, \Sigma)$ according to f .

for given parameters $\log \beta_0^P, \log \beta_0^D, \nu_0, S_0$ and n_0 . The Wishart distribution W with variance-covariance matrix S_0 and $n_0 > 1$ degrees of freedom specifies agents' marginal prior 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 $(\log \beta^P, \log \beta^D)$ conditional on the precision matrix H being equal to 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 $(\log \beta^P, \log \beta^D)$.

We further restrict our attention to prior beliefs such that

$$(\beta_0^P, \beta_0^D) = (a, a) \quad (21)$$

$$S_0 = \sigma^2 \begin{pmatrix} 1 & 1 - \kappa \\ 1 - \kappa & 1 \end{pmatrix} \quad (22)$$

This implies that agents' beliefs are initially centered at the REE. Moreover, in the limiting case with vanishing prior uncertainty ($n_0, \nu_0 \rightarrow \infty$ with $n_0/\nu_0 \rightarrow 1$) and perfectly correlated innovations ($\kappa \rightarrow 0$) the resulting beliefs P converge in distribution to P^{REE} . In other words, we have weak convergence to RE beliefs. It is in this precise sense that for large (n_0, ν_0) , small κ and initial values given by equations (21) and (22), agents' beliefs P involve only a 'small deviation from REE beliefs'.

The next section determines the equilibrium asset prices implied by the beliefs \mathcal{P} .

5.3 Internally Rational Expectations Equilibrium

When all agents' beliefs are given by \mathcal{P} , it follows from equations (13) and (19) that the equilibrium asset price is given by

$$\begin{aligned} \mathbf{P}_t &= \delta E_t^{\mathcal{P}} (P_{t+1} + D_{t+1}) \\ &= \delta E_t^{\mathcal{P}} \left(e^{\log \beta^P} e^{\log \varepsilon_{t+1}^P} \right) \mathbf{P}_t + \delta E_t^{\mathcal{P}} \left(e^{\log \beta^D} e^{\log \varepsilon_{t+1}^D} \right) D_t \end{aligned} \quad (23)$$

The equilibrium price thus depends on agents' conditional time t expectations of $e^{\log \beta^X} e^{\log \varepsilon_{t+1}^X}$ (for $X \in \{P, D\}$). The next section determines the evolution of these conditional expectations over time and shows that they can be described by ordinary least squares (OLS) learning rules.

5.4 Bayesian Updating and OLS Learning

We now show that for the beliefs specified in the previous sections, optimal belief updating in an IREE gives rise to the least squares learning equations that have been typically studied in the adaptive learning literature.

We start by determining the posterior beliefs for $\log \beta_P, \log \beta_D$ and Σ . We thereby use the fact that our prior specification (20) is conjugate, so that the posterior in period t is again of the Normal-Wishart form (20).²⁰ The posterior is given by

$$H \mid \omega^t \sim W(S_t, n_t) \quad (24a)$$

$$(\log \beta^P, \log \beta^D)' \mid H = h, \omega^t \sim N \left((\log \beta_t^P, \log \beta_t^D)', (\nu_t h)^{-1} \right) \quad (24b)$$

where the parameters $(\log \beta_t^P, \log \beta_t^D, \nu_t, S_t, n_t)$ evolve recursively as follows:

$$\begin{pmatrix} \log \beta_{t+1}^P \\ \log \beta_{t+1}^D \end{pmatrix} = \begin{pmatrix} \log \beta_t^P \\ \log \beta_t^D \end{pmatrix} + \frac{1}{\nu_t + 1} e_t \quad (25a)$$

$$\nu_{t+1} = \nu_t + 1 \quad (25b)$$

$$S_{t+1}^{-1} = S_t^{-1} + \frac{\nu_t}{\nu_t + 1} e_t e_t' \quad (25c)$$

$$n_{t+1} = n_t + 1 \quad (25d)$$

with e_t denoting the one-step-ahead ‘forecast error’:

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

Letting ‘ \approx ’ denote an approximation that is correct up to first order, appendix A.2 shows that the posterior beliefs imply that the conditional expectations appearing in the pricing equation (23) are given by

$$E_t^P \left(e^{\log \beta^P} e^{\sigma_P^2/2} \right) \approx \widehat{\beta}_t^P \quad (26)$$

with $\widehat{\beta}_t^P$ denoting the ordinary least squares (OLS) estimate of stock price growth given by

$$\widehat{\beta}_t^P \equiv \frac{1}{t + \nu_0} \sum_{j=1}^t \frac{P_j}{P_{j-1}} + \frac{\nu_0}{t + \nu_0} \beta_0^P$$

²⁰The subsequent result follows from Theorem 1, chapter 9.10 in DeGroot (1970). Our variables can be mapped into the ones employed in DeGroot’s theorem using: $h \rightarrow r$, $H \rightarrow R$, $n_t \rightarrow \alpha$, $S_t^{-1} \rightarrow \tau$, $(\log \beta_t^P, \log \beta_t^D)' \rightarrow \mu$, $\nu_t \rightarrow \nu$. Unlike DeGroot, we parameterize Normal and Wishart distributions using variance-covariance matrices rather than precision matrices, as this has become more common nowadays.

Note that the OLS estimator incorporates the prior β_0^P by treating it like ν_0 observations of stock price growth in the data. Similar approximations for dividend expectations yield

$$E_t^{\mathcal{P}} \left(e^{\log \beta^D} e^{\sigma_D^2/2} \right) \approx \widehat{\beta}_t^D \equiv \frac{1}{t + \nu_0} \sum_{i=0}^t \frac{D_i}{D_{i-1}} + \frac{\nu_0}{t + \nu_0} \beta_0^D$$

so that the pricing equation (23) implies - up to a first order approximation of conditional expectations - that

$$\mathbf{P}_t = \delta \widehat{\beta}_t^P \mathbf{P}_t + \delta \widehat{\beta}_t^D D_t \quad (27)$$

or, equivalently

$$\mathbf{P}_t = \frac{\delta \widehat{\beta}_t^D}{1 - \delta \widehat{\beta}_t^P} D_t \quad (28)$$

This is the equation studied by Adam, Marcet and Nicolini (2010). This equation clearly says that learning about price growth behavior influences equilibrium stock prices in a model with internally rational agents who hold a complete and consistent set of probability beliefs. Specifically, it is clear that higher expected price growth $\widehat{\beta}_t^P$ implies a higher price-dividend ratio and this is what generates the price dynamics studied in that paper.²¹

5.5 Specifying Agents' Beliefs

Our definitions of internal rationality and Internally Rational Expectations Equilibrium (IREE) take agents' beliefs \mathcal{P} as exogenously given. This allows to formulate simple microfoundations and to apply standard dynamic modeling tools. It also reflects the fact that - once one departs from the REH - agents' beliefs \mathcal{P} inevitably become part of the microfoundations of the model and thus a choice variable for the economic

²¹Equation (28) reveals that existence of an equilibrium price requires that the (approximate) posterior *mean* for expected price growth $\widehat{\beta}_t^P$ remains below δ^{-1} . In Adam, Marcet and Nicolini (2010) this condition is insured by imposing an ad-hoc continuous projection facility on beliefs, which bounded mean price growth expectations below δ^{-1} . We conjecture that this can be obtained from a consistent system of beliefs \mathcal{P} by truncating the upper tail of the prior density f for the unknown parameter $\log \beta^P$. Also, since the equilibrium price \mathbf{P}_t affects the left and right hand side of equation (28) there may actually exist multiple or no mutually consistent equilibrium price and belief pairs. To avoid this, the estimate $\widehat{\beta}_t^P$ in Adam, Marcet and Nicolini (2010) is computed using prices up to period $t - 1$ only. Adam and Marcet (2010) provide an information structure where such delayed updating arises from fully Bayesian updating behavior in a setting where agents nevertheless observe contemporaneous prices and dividends.

modeler. In contrast, when working under the REH, agents' beliefs \mathcal{P} are largely dictated by the remaining micro-foundations of the model, so that economic modelers typically do not have to specify agents' beliefs.²² While the REH simplifies model construction enormously, we have argued that this way of specifying beliefs may not be the most plausible one: it incorporates a singularity into agents' beliefs that contains so much information about market behavior that not even expert economists can plausibly claim to possess.

When departing from the REH, one inevitably has to specify agents' beliefs; and this endows the economic modeler with additional degrees of freedom. While these could be used in an unfruitful way, it is equally possible that economists will learn to employ these degrees of freedom productively. In this section we offer some preliminary comments on belief selection - most of the issues we address are still part of ongoing research.

To avoid arbitrary belief specifications, it appears desirable when specifying agents' belief systems to incorporate some concept of near-rationality with respect to the resulting equilibrium outcomes.

In the present paper we propose to consider agents whose beliefs are close (in distribution) to REE beliefs, as discussed in section 5.2. Studying such small deviations from REE beliefs appears of interest because anyone willing to assign to agents the REE beliefs should be equally willing to explore the consequences of endowing agents with small deviations from such beliefs. This is especially true if such small deviations are associated with a considerable improvement in the empirical performance of the model, as shown in Adam, Marcet and Nicolini (2010). In the present model, this feature also insures that agents' beliefs are close to the resulting equilibrium outcomes because the equilibrium outcomes are continuous with respect to agents' beliefs.

More generally, one might ask whether decision makers maintain a system of beliefs \mathcal{P} , as we assumed, if they observe equilibrium prices that have a distribution different from \mathcal{P} ? This can only be expected if the belief system \mathcal{P} is not too obviously wrong after observations of the actual equilibrium outcomes in the model are available. This is the case, if beliefs and actual outcomes are close (in distribution), or if agents' beliefs converge over time to such a situation, or if agents even learn over time to make RE predictions. In the latter case, agents would realize in the long run that their beliefs are correct and thus have no incentive to change their belief system \mathcal{P} . For the homogenous agent version of the asset pricing model considered in this paper, Adam, Marcet and Nicolini

²²Exceptions arise whenever the RE equilibrium fails to be unique, as is the case in many monetary models, for example.

(2010) show that least squares learning converges globally and almost surely to RE. Therefore, agents in the setup studied in this paper will not change the model on which they are basing their beliefs after observing the equilibrium outcomes for a large number of periods.

Belief systems \mathcal{P} are unlikely to be kept for long, however, if they assign zero likelihood to outcomes that agents can actually observe in equilibrium. The observation of zero probability (more precisely, zero density) events would tell agents that their belief systems are incompatible with the observed equilibrium path and it should induce reasonable agents to immediately adjust their beliefs. In the present paper, we avoided this problem by specifying beliefs that have full support over all outcomes, therefore, any combination of prices and dividends has positive density.²³

Future research might show that other features could also give rise to unattractive belief systems. Indeed, determining what constitutes acceptable belief specifications promises to be an exciting avenue for further research.²⁴

Macroeconomists sometimes appear reluctant to consider relaxations of the REH. This may partly be motivated by the fear that economists could then start choosing belief systems for their models that appear unattractive when compared to the equilibrium implications they generate. We propose to interpret such outcomes simply as bad modeling choices, something that can happen in all aspects of economic modeling, e.g., even when it comes to choosing utility functions. While economists could make unreasonable utility specifications, this does not imply that the task of specifying appropriate utility functions would not be useful: important progress has been made, for example, by showing which preference specifications are consistent with the Kaldor facts. Similar progress appears possible when it comes to determining appropriate belief systems.

Finding the right way to model expectations promises to be an exciting avenue for research. Discipline can be imposed, not only by re-

²³Some of the literature employing ‘near-rational’ or ‘robust’ belief systems, e.g., Woodford (2010) or Hansen and Sargent (2009), has avoided such inconsistencies by requiring that near-rational beliefs are absolutely continuous with respect to the equilibrium outcomes.

²⁴For example, we find it equally acceptable in our model to assume that agents use a constant gain learning algorithm instead of the OLS learning equations. This would be an IREE if the belief system is as specified in section 5.2 but has time-varying parameters ($\log \beta^P, \log \beta^D$) in the law of motion (19), with these parameters following independent unit root processes. As long as the variance of the innovation in the unit root processes also becomes arbitrarily small, the implied constant gain becomes small and we still have weak convergence in distribution to REE beliefs.

stricting agents to be near-rational, but also by using the many available sources of indirect data on expectations, including survey evidence and financial market data which is closely linked to investors' expectations.

6 Sensitivity of Bayesian REE Asset Prices

This section demonstrates that the asset price in a Bayesian REE is extremely sensitive to fine details in the specification of agents' prior beliefs about dividend growth. Indeed, details of agents' prior beliefs about dividend growth matter much more for asset prices than the microeconomic structure of the economy. Since economists will probably never find out about details on the prior, this represents a degree of freedom in Bayesian REE modeling that strongly influences the Bayesian REE asset price.

For simplicity, we consider again a homogeneous agent model in which all agents hold the same discount factor and dividend beliefs and where this is common knowledge. Each internally rational agent can then deduce the Bayesian REE asset price associated with any history of dividends. This price is given by

$$\mathbf{P}_t = E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \middle| D^t \right) \quad (29)$$

but turns out to be extremely sensitive with respect to the prior beliefs about the dividend process incorporated in \mathcal{P} . Specifically, as we show in proposition 2 below, the equilibrium price can be increased by any desired amount by simply reallocating an arbitrarily small amount of prior probability mass.

To illustrate this point we rewrite the dividend process as

$$D_t = aD_{t-1}\eta_t \quad (30)$$

where $\eta_t > 0$ is i.i.d. with $E[\eta_t] = 1$ and $a > 0$.²⁵ Agents' prior density about a is denoted by f and satisfies $f(\tilde{a}) = 0$ for all $\tilde{a} \leq 0$. The posterior density about a conditional on any observed history D^t is denoted by $Post_t$.

The following proposition provides a first result. It shows that unless the posterior beliefs about dividend growth are bounded by the inverse of the discount factor, equilibrium prices are infinite. The proof of the proposition can be found in appendix A.3.

²⁵The parameter a in the equation above is not exactly equal to the one employed in equation (17), but this is of no importance for the results that follow.

Proposition 1 Consider the Bayesian REE asset price (29). For any t and D^t :

1. If $Post_t(a \geq \delta^{-1}) > 0$, then

$$\mathbf{P}_t = \infty$$

2. Let B denote the upper bound of the support of $Post_t$. If $B < \delta^{-1}$ then

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

The previous proposition is closely related to results derived in Pesaran, Pettenuzzo and Timmermann (2007), but also differs because it does not rely on parametric forms for the prior beliefs.²⁶ This will prove useful below for showing that the sensitivity of prices to priors is a general phenomenon.

Proposition 1 above shows that the asset price is finite whenever the support of the posterior density is bounded below δ^{-1} . This can be insured, for example, by choosing prior beliefs f with an upper bound of the support $B < \delta^{-1}$. The posterior beliefs will then inherit this property. This in turn might suggest that Bayesian REE asset prices can not be arbitrarily high provided one imposes an upper bound $B < \delta^{-1}$ on the support of prior dividend growth beliefs. Yet, the following proposition, which is proven in appendix A.3, shows that this fails to be true:

Proposition 2 Consider a prior density f with upper bound $B < \delta^{-1}$ for its support. There exists a sequence of densities $\{f^k\}$ with upper bound $B^k < \delta^{-1}$ and $\int |f^k - f| \rightarrow 0$ such that the Bayesian REE price implied by f^k converges to infinity as $k \rightarrow \infty$.

Bounding the prior support is thus not a very robust solution to the problem for the high sensitivity of Bayesian REE asset prices: given any prior that implies bounded prices, there exists another prior that is arbitrarily close to it and that gives rise to prices that are arbitrarily large.

Due to a (well-acknowledged) shortcut this sensitivity of Bayesian REE asset prices sometimes failed to show up in some of the Bayesian

²⁶The result in proposition 1 also differs from the examples in Geweke (2001) and Weitzman (2007) where non-existence of expected utility does not arise from a diverging infinite discounted sum. Instead, in these papers one-period-ahead expected consumption utility already fails to be finite.

REE literature. Timmermann (1993, 1996), for example, ignores the posterior uncertainty about a and uses instead only the posterior mean for a to evaluate the discounted sum in equation (29) to set

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

where

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

so that only the posterior mean matters for equilibrium prices. As should be clear from the results in this section, this shortcut can strongly alter the asset pricing implications.²⁷

7 Relation to the Literature

The concept of an Internally Rational Expectations Equilibrium (IREE) developed in this paper is a generalization of the Bayesian REE concept. The latter emerges as a special case of IREE when agents possess sufficient knowledge about the market so that they can deduce the equilibrium market outcome associated with any possible sequence of fundamentals.

The IREE is also related to the private information REE analyzed in Allen, Morris and Shin (2006). Both equilibrium concepts relax the common knowledge assumptions of standard models, in the case of Allen, Morris and Shin due to the assumption of private information, in our case due to the assumption of imperfect market knowledge. Also, it appears relatively straightforward to extend the IREE presented in the present paper so as to incorporate private information. This extension would cause private information REE to be a special case of private information IREE.

The relationship between IREE and Arrow-Debreu equilibrium has been discussed in section 3.4 before. The absence of Arrow securities in our model appears to be an important ingredient giving rise to the possibility that IREE outcomes can differ from the Arrow-Debreu equilibrium outcomes. It seems worthwhile investigating this issue in greater detail in future research.

Our work is also related to a number of papers in the learning literature that attempt to construct a full set of beliefs over long-horizons, see important work by Preston (2005) and Eusepi and Preston (2010). The main difference is that our agents' beliefs take the form of a well

²⁷Similarly, Pastor and Veronesi (2003) assume a finite asset price after some fixed terminal date $T < \infty$. Sensitivity then arises with respect to the chosen date T .

defined probability measure over a stochastic process while these papers use the anticipated utility framework of Kreps (1998). As a consequence, agents in these models construct each period a new probability measure, but one that is almost surely inconsistent with the measure 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). In rational beliefs models agents' probability densities are assumed to be shifting in response to the realization of an extrinsic 'generating sequence'. In our model belief revisions are triggered by model intrinsic factors, i.e., market outcomes and fundamentals. Moreover, in rational belief models, agents entertain a standard probability space over fundamentals. Thus, unlike in the present setup, agents' joint beliefs about prices and dividends incorporates a singularity.

The IREE is also related to the model-consistent equilibrium concept of Anderson and Sonnenschein (1985) who assume that agents have a parameterized econometric model that defines a probability density over prices. Anderson and Sonnenschein, however, impose a kind of rational expectations structure on the beliefs of agents: within the class of models considered, agents are assumed to employ the parameter values that best fit the actual outcome of the data. Thus, unlike in our setting there is no learning from market outcomes because agents are assumed to know the best fitting model from the start.

The paper is also related to recent work on self-confirming equilibrium (SCE), see Fudenberg and Levine (1993), Sargent (1999), and Cho and Sargent (2008). A SCE is a stationary equilibrium in which agents know the objective probability distributions as they emerge in equilibrium, although agents may entertain wrong beliefs about what would happen off the equilibrium path.²⁸ In an SCE agents' joint beliefs over fundamentals and market outcomes thus features a singularity, as in the case with rational expectations. The singularity again reflects the fact that agents have perfect knowledge about equilibrium market behavior, possibly because they can observe an infinite history of past equilibrium outcomes and have learned to optimally condition on the past. This contrasts with the situation described in an IREE, where agents are uncertain about equilibrium market behavior and their beliefs do not contain a singularity, possibly because they do not observe infinite amounts of past (stationary equilibrium) market outcomes. An IREE can nevertheless converge asymptotically to a SCE or an REE, as is actually the case in our example in section 5. And to avoid situations in

²⁸This generates gaps between REE and SCE, as a REE is a SCE but not vice versa.

which agents have infinite amounts of past data available but still forecast suboptimally, our definition of an IREE in section 3.4 could have added the following requirement: if an IREE converges to a stationary outcome, then this outcome must be an SCE or REE. The difficulty with this and similar restrictions on asymptotic rationality is that they have no implication for any finite amounts of data, i.e., no empirically testable hypotheses arise from such asymptotic requirements on forecast optimality. And requiring full optimality of beliefs along the convergence process leads one back to the concept of a Bayesian REE. In the light of this discussion, the development of a general metric allowing to quantify intermediate degrees of forecast rationality over a convergence process appears to be an important task for future research.

8 Conclusions

We formulate a model with internally rational agents that fail to be externally rational because they possess limited knowledge about the market. Lack of market knowledge naturally gives rise to learning from market outcomes, so that expectations regarding the future market outcomes become an important determinant of the current market outcome, independently from agents' expectations about fundamentals. Since market outcomes feed back into agents beliefs, imperfect market knowledge can give rise to additional propagation in economic models.

The present paper shows - perhaps surprisingly - that some of the modeling choices in the adaptive learning literature, e.g., the use of one-step-ahead Euler equations, are less ad-hoc than might initially appear. Adam, Marcet and Nicolini (2010) show that this generalizes to a setting with risk averse agents, provided stock market wealth is a negligible part of agents' total wealth. Yet, in the more appealing case with non-negligible stock market wealth this ceases to be true. Long horizon asset price forecasts then matter for current equilibrium prices, as shown in Adam and Marcet (2010). This is so because future prices affect future consumption abilities and agents dislike consumption volatility when they are risk averse. The microfoundations of the model are thus informative about which beliefs matter for the equilibrium outcomes in models of learning.

We believe that internal rationality is an interesting approach to describe agent behavior in models of learning. It gives rise to models of learning that are fully consistent with optimizing behavior and provides a rationale for equilibrium equations with very different implications than those emerging from the standard REE approach. We expect the concept of internal rationality to have a wide range of fruitful applications in many field of economics dealing with dynamic decision making.

A Appendix

A.1 Existence of a Maximum

Strictly speaking one has to guarantee existence of a maximum before the first order conditions (8) can be used. With arbitrary price beliefs and risk neutrality existence of a maximum is in question, since an agent that assigns positive probability to prices growing at a rate larger than δ^{-1} achieves infinite utility. We outline below how to modify slightly the agents' utility in a way that existence is guaranteed and that the pricing implications of the model are virtually unchanged. We only consider beliefs \mathcal{P} such that $P, D > 0$ almost surely.

Consider the family of utility functions that is indexed by a finite constant \bar{C}

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

where g is a strictly increasing, strictly concave, differentiable and bounded function satisfying $g(0) = 0$, $g'(0) = 1$ and $g(\cdot) \leq \bar{g}$. We have $U_{\bar{C}} = 1$ for $C < \bar{C}$ but $U_{\bar{C}} < 1$ for $C \geq \bar{C}$.

The consumer's problem can be rewritten as a function only of the choice for contingent stock holding plans $S = \{S_0, S_1, \dots\}$ with $S_t : \Omega^t \rightarrow [0, \bar{S}]$. Since for any given \bar{C} the objective function is bounded and since the action space S is compact, an expected utility maximizing plan does exist. We now show that for any finite number of periods $T < \infty$, the first order conditions with this bounded utility function are arbitrarily close to those used in the main text as $\bar{C} \rightarrow \infty$. Hence, the pricing implications are arbitrarily close to those in the main text.

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 is close to one when \bar{C} is sufficiently high. Therefore, with arbitrarily high probability the agent's first order conditions in $t = 1, \dots, T$ are given by

$$P_t = \delta^i E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] \quad \text{and } S_t^i \in [0, \bar{S}] \quad (32)$$

plus the slackness conditions for $S_t^i = \bar{S}$ or $S_t^i = 0$. From Lebesgue's Dominated Convergence Theorem

$$\lim_{\bar{C} \rightarrow \infty} E_t^{\mathcal{P}^i} [U'_{\bar{C}}(C_{t+1}^i) (P_{t+1} + D_{t+1})] = E_t^{\mathcal{P}^i} [(P_{t+1} + D_{t+1})] \quad (33)$$

This implies that for $\bar{C} \rightarrow \infty$ the first order conditions (32) approximate with arbitrary precision the first order conditions (8) used in the main text.

A.2 Approximation of Beliefs

We prove (26), ‘ \approx ’ denotes an equality correct up to first order.

$$\begin{aligned} E_t^{\mathcal{P}} \left(e^{\log \beta^P} e^{\log \varepsilon_{t+1}^P} \right) &\approx E_t^{\mathcal{P}} \left(1 + \log \beta^P + \log \varepsilon_{t+1}^P \right) = 1 + \log \beta_t^P \\ &= 1 + \frac{1}{t + \nu_0} \sum_{j=1}^t \log \frac{P_j}{P_{j-1}} + \frac{\nu_0}{t + \nu_0} \ln \beta_0^P \\ &\approx \frac{1}{t + \nu_0} \sum_{j=1}^t \frac{P_j}{P_{j-1}} + \frac{\nu_0}{t + \nu_0} \beta_0^P = \widehat{\beta}_t^P \end{aligned}$$

the first line approximates $\log \beta^P$ and $\log \varepsilon_{t+1}^P$ around $\log \beta^P = \log \varepsilon_{t+1}^P = 0$, uses $E_t^{\mathcal{P}} \log \varepsilon_{t+1}^P = 0$ and that $E_t^{\mathcal{P}} [\log \beta^P] = \log \beta_t^P$;²⁹ the second line uses the update rule (25a); the third line approximates P_j/P_{j-1} and β_0^P around $P_j/P_{j-1} = 1$ and $\log \beta^P = 0$.

A.3 Proof of Propositions

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

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

so that the partial discounted sum can be expressed as

$$\sum_{j=1}^T \delta^j D_{t+j}(\omega_D) = \sum_{j=1}^T \delta^j a(\omega_D)^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t \quad (34)$$

To prove the first part of the proposition notice

$$\begin{aligned} &E^{\mathcal{P}} \left(\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j}(\omega_D) \middle| D^t \right) \geq E^{\mathcal{P}} \left(\sum_{j=1}^T \delta^j D_{t+j}(\omega_D) \middle| D^t \right) \\ &= E^{\mathcal{P}} \left(\sum_{j=1}^T \delta^j a(\omega_D)^j \prod_{\tau=1}^j \eta_{t+\tau}(\omega_D) D_t \middle| D^t \right) = D_t \int_0^{\infty} \left(\sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} \\ &\geq D_t \int_{\delta^{-1}}^{\infty} \left(\sum_{j=1}^T \delta^j (\tilde{a})^j \right) Post_t(\tilde{a}) d\tilde{a} \geq D_t \cdot T \cdot \int_{\delta^{-1}}^{\infty} Post_t(\tilde{a}) d\tilde{a} = \infty \end{aligned}$$

²⁹This follows from the fact that the marginal posterior for $(\log \beta^P, \log \beta^D)'$ is Student t -distributed, location vector $(\log \beta_t^P, \log \beta_t^D)'$. See chapter 9.11 in DeGroot (1970).

where the first inequality uses the fact $D_t \geq 0$, the first equality the expression (34), the next equality the independence of future η 's from D^t and a , the next inequality uses the fact that dividends are positive and the second the fact that $\delta \tilde{a} \geq 1$ over the considered range of integration. The last equality uses that since part of the support is higher than δ^{-1} implies $\int_{\delta^{-1}}^{\infty} Post_t(\tilde{a}) d\tilde{a} > 0$.

To prove the second part of the proposition, given t , define the function

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

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 since $E^{\mathcal{P}}(\mathcal{F}(\omega_D)|D^t) = \frac{\delta B}{1-\delta B} D_t < \infty$. Moreover, for all T and for given D^t

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

Therefore, the partial sums (34) 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 below:

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

The second equality follows from using (34) and taking expectations as when deriving (35). This proves the second part of the proposition. ■

Proof of Proposition 2.: The proof is by construction. Let us first prove that $P_0 \rightarrow \infty$ for a sequence of priors when $k \rightarrow \infty$. Given any prior f with upper bound for the support $B < \delta^{-1}$ we can construct an alternative sequence of priors

$$f^k(a) = \begin{cases} (1 - \frac{B^k - B}{k}) f(a) & \text{for } a \in [0, B] \\ \frac{1}{k} & \text{for } a \in]B, B^k] \end{cases}$$

where $B^k = \max\{B, \delta^{-1}(1 - \frac{1}{k})\}$.

The density f^k distributes probability mass $\frac{1}{k}$ uniformly on the interval $]B, B^k]$, here B^k is the upper bound for the support of f^k and we

have $\lim_{k \rightarrow \infty} B^k = \delta^{-1}$. As required this sequence of alternative priors satisfies $\int |f^k - f| \rightarrow 0$ as $k \rightarrow \infty$ because

$$\int |f^k - f| = \int_0^B \left| \frac{B^k - B}{k} f \right| + \int_B^{B^k} \frac{1}{k} = 2 \frac{B^k - B}{k} < 2 \frac{\delta^{-1} - B}{k}$$

Moreover, it follows from part 1 of proposition 1 and simple derivations that

$$\begin{aligned} P_0 &= E_0^{\mathcal{P}^k} \left(\sum_{j=1}^{\infty} \delta^j D_j \right) = D_0 E_{f^k} \left(\frac{\delta a}{1 - \delta a} \right) \\ &\geq D_0 \int_B^{B^k} \frac{\delta \tilde{a}}{1 - \delta \tilde{a}} f^k(\tilde{a}) d\tilde{a} \geq D_0 \frac{B\delta}{k} \int_B^{B^k} \frac{1}{1 - \delta \tilde{a}} d\tilde{a} \end{aligned}$$

Using the change of variables $x = 1 - \delta \tilde{a}$ the integral in the last line can be expressed as

$$\int_B^{B^k} \frac{1}{1 - \delta \tilde{a}} d\tilde{a} = \int_{1 - \delta B^k}^{1 - \delta B} \frac{1}{x} dx = \frac{1}{(1 - \delta B^k)^2} - \frac{1}{(1 - \delta B)^2}$$

and we have

$$\lim_{k \rightarrow \infty} \frac{B\delta}{k} \int_B^{B^k} \frac{1}{1 - \delta \tilde{a}} d\tilde{a} = \lim_{k \rightarrow \infty} \frac{B\delta}{k} \left(k^2 - \frac{1}{(1 - \delta B)^2} \right) = \infty$$

which establishes the claim for P_0 .

For P_t , all that changes is that the posterior for each f^k is scaled by the likelihood of the observed realization according to the model at hand. As long as the likelihood puts positive weight on all parameter values below δ^{-1} the derivations above work in the same way. ■

References

- ADAM, K. (2003): “Learning and Equilibrium Selection in a Monetary Overlapping Generations Model with Sticky Prices,” *Review of Economic Studies*, 70, 887–908.
- (2005): “Learning to Forecast and Cyclical Behavior of Output and Inflation,” *Macroeconomic Dynamics*, 9(1), 1–27.
- ADAM, K., AND A. MARCET (2010): “Booms and Busts in Asset Prices,” *London School of Economics mimeo*.
- ADAM, K., A. MARCET, AND J. P. NICOLINI (2010): “Learning and Stock Market Volatility,” *Mannheim University Mimeo*.
- ALLEN, F., S. MORRIS, AND H. S. SHIN (2006): “Beauty Contest and Iterated Expectations in Asset Markets,” *Review of Financial Studies*, 19, 719–752.

- ANDERSON, R., AND H. SONNENSCHN (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.
- BLUME, L., AND D. EASLEY (2006): "If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica*, 74(4), 929–966.
- 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.
- CHAKRABORTY, A., AND G. EVANS (2008): "Can Perpetual Learning Explain the Forward Premium Puzzle?," *Journal of Monetary Economics*, 55(3), 477–490.
- CHO, I.-K., AND T. J. SARGENT (2008): "Self-Confirming Equilibria," in *The New Palgrave Dictionary of Economics, Second Edition*.
- 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 (2010): "Expectations, Learning and Business Cycle Fluctuations," *American Economic Review (forthcoming)*, 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.
- FUDENBERG, D., AND D. LEVINE (1993): "Self-Confirming Equilibrium," *Econometrica*, 61, 523–545.
- GEWEKE, J. (2001): "A Note on some Limitations of CRRA Utility," *Economics Letters*, 71, 341–345.
- HANSEN, L. P., AND T. J. SARGENT (2009): *Robustness*. Princeton University Press, Princeton, Stanford University, mimeo.
- 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.
- MAR CET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93, 1476–1498.
- MARIMON, R. (1997): “Learning from Learning in Economics,” in *Advances in Economic Theory*, ed. by D. Kreps, and K. F. Wallis, Cambridge. Cambridge University Press.
- MOLNAR, K., AND S. SANTORO (2007): “Optimal Monetary Policy When Agents are Learning,” *Norwegian School of Economics Working Paper No. 7*.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315–335.
- 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.
- (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98(1), 5–37.
- TIMMERMANN, A. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1996): “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.
- WEITZMAN, M. L. (2007): “Subjective Expectations and Asset-Return Puzzles,” *American Economic Review*, 97(4), 1102–1130.
- WOODFORD, M. (2010): “Robustly Optimal Monetary Policy with Near-Rational Expectations,” *American Economic Review (forthcoming)*.