ASSET PRICES IN A TIME SERIES MODEL WITH *PERPETUALLY* DISPARATELY INFORMED, COMPETITIVE TRADERS

KENNETH KASA, TODD B. WALKER, AND CHARLES H. WHITEMAN

ABSTRACT. This paper develops a dynamic asset pricing model with persistent heterogeneous beliefs. The model features competitive traders who receive idiosyncratic signals about an underlying fundamentals process. We adapt Futia's (1981) frequency domain methods to derive conditions on the fundamentals that guarantee noninvertibility of the mapping between observed market data and the underlying shocks to agents' information sets. When these conditions are satisfied, agents must 'forecast the forecasts of others'. The paper provides an explicit analytical characterization of the resulting higher-order belief dynamics. These additional dynamics can explain apparent violations of variance bounds and rejections of cross-equation restrictions.

JEL Classification Numbers: G12, D82

1. INTRODUCTION

The standard present value model has a difficult time explaining several features of observed asset prices. From the perspective of this model, prices seem to be excessively volatile. The model's cross-equation and Granger causality restrictions are typically rejected as well. As a result, linear present value models have all but disappeared from serious academic research on asset pricing.[1] Instead, attention has shifted to models with time-varying risk premia. Unfortunately, these models offer little improvement empirically, although Constantinides and Duffie (1996) achieve some success by introducing heterogeneity, in the form of nondiversifiable labor income risk.

This paper returns to the linear constant discount rate framework, and argues that *informational heterogeneity* can account for many of the model's apparent empirical

Date: August 20, 2007.

We would like to thank Tim Cogley, Martin Evans, Ron Michener, Bart Taub, Pierre-Olivier Weill and seminar participants at NYU, the Federal Reserve Board, UC Davis, Indiana University, VPI, UNC, and the Atlanta Fed for helpful discussions. Support from the National Science Foundation under grant SES85-10505 is gratefully acknowledged. This paper subsumes and extends Walker and Whiteman's *Equilibrium Volatility and Heterogeneity in a Simple Asset Pricing Model* and Kasa's *Asset Pricing with Heterogeneous Beliefs: A Frequency Domain Approach*.

[1]Cochrane (2001) discusses the empirical failings of constant discount rate models. He emphasizes that many apparently distinct anomalies, such as predictability and excess volatility, can be interpreted as manifestations of the same underlying problem; namely, a misspecification of the discount rate. He also emphasizes that the same problems show up in all asset markets, e.g., stocks, bonds, foreign exchange, real estate, etc.

shortcomings. In particular, we make one simple change to the standard present value model–we assume fundamentals consist of a sum of orthogonal components, and that individuals observe different pieces of this sum. The presence of asymmetric information places rational investors into a situation where they must 'forecast the forecasts of others' (see, Townsend (1983), Singleton (1987)). We demonstrate how the resulting higher-order belief dynamics can reconcile standard present value models with several apparent empirical anomalies.

Of course, this is not the first paper to study the role of asymmetric information in asset markets, nor is it the first to study higher-order beliefs.[2] However, our approach *is* the first to combine several key ingredients. First, our model is *dynamic*, it features *persistent* heterogeneous beliefs, and the equilibrium is *stationary*. Following Grossman and Stiglitz (1980), most existing work on asset pricing with asymmetric information is confined to static, or finite-horizon models. Although this is a useful abstraction for some theoretical questions, it is obviously problematic when it comes to empirical applications. There has been some work devoted to dynamic extensions of the Grossman-Stiglitz framework (see, e.g., Wang (1993)), but following Grossman and Stiglitz, this literature postulates hierarchical information structures, with 'informed' and 'uninformed' traders. Again, this assumption has its uses, but from our perspective it 'throws the baby out with the bath water', since it eliminates the forecasting the forecasts of others problem (Townsend (1983)). Our model postulates a more natural symmetric information structure.

Second, our approach features *signal extraction from endogenous prices*. This distinguishes our work from the recent work on global games and imperfect common knowledge (Morris and Shin (1998, 2000, 2003)). Although this literature has made important contributions to our understanding of higher-order beliefs, it is not directly applicable to asset pricing, since it abstracts from asset markets. As Atkeson (2000) notes, prices play an important role in aggregating information, and it remains to be seen how robust the work on global games is to the inclusion of asset markets.[3]

Third, our approach delivers an *analytical solution*, with explicit closed-form expressions for the model's higher-order belief dynamics. Although this may seem like a minor contribution given the power of computation these days, it turns out that analytical solutions are extremely useful in models featuring an infinite regress of higher-order beliefs. Numerical methods in this setting are fraught with dangers. In particular, they require prior knowledge of the relevant state vector. As first noted by Townsend (1983), it is not at all clear what the state is when agents forecast the forecasts of others. Townsend argued that the logic of infinite regress produces

[2]See Brunnermeier (2001) for a nice overview of the existing literature.

[3]Recent work by Angeletos and Werning (2006) incorporates signal extraction from prices into the Morris-Shin framework. However, their model is essentially static. Moreover, their informational assumptions preclude even the possibility that prices could be fully revealing. In contrast, our work focuses directly on the conditions *required* for a non-revealing equilibrium.

an infinite-dimensional state. He short-circuited the infinite regress and obtained a tractable numerical solution by assuming that information becomes common knowledge after a (small) number of periods. This truncation strategy has since been applied by a number of subsequent researchers (see, e.g., Singleton (1987), Bacchetta and van Wincoop (2006), and Nimark (2007)). However, recent work by Pearlman and Sargent (2005) and Walker (2007) demonstrates that numerical approaches can be quite misleading. Pearlman and Sargent, employing a clever 'guess and verify' strategy based on the incorporation of lagged forecast errors in the state, showed that Townsend's model in fact produces an equilibrium that is fully revealing, and that Townsend's higher-order belief dynamics are entirely an artifact of his numerical methods. Walker (2007) does the same thing for Singleton's asset pricing version of Townsend's model. Using the same approach as this paper, he obtains an analytical solution *without* truncation, and shows that the equilibrium is in fact fully revealing.

Our approach adapts and extends the frequency domain methods of Futia (1981). These methods exploit the power of the Riesz-Fischer Theorem. This theorem allows us to transform a difficult time-domain/sequence-space signal extraction problem into a much easier function space problem. Rather than guess a state vector and then solve a Kalman filter's Riccati equation, a frequency domain approach leads to the construction of Blaschke factors. Finding these Blaschke factors is the key to solving the agents' signal extraction problems.[4] In general, finding Blaschke factors is no easier than solving Riccati equations. However, a key innovation in our approach is to work *backwards* from postulated Blaschke factors to the supporting set of fundamentals. This reverse engineering strategy allows us to isolate necessary conditions for the existence of a heterogeneous beliefs equilibrium. The advantages gained from knowing these conditions cannot be overestimated. For example, recent work by Makarov and Rytchkov (2006) also applies frequency domain methods to a linear present value asset pricing model. In contrast to our approach, they begin by postulating a time series process for fundamentals, and then search for an equilibrium price process. Interestingly, they argue that a finite-state equilibrium does not exist. However, their fundamentals specification does not satisfy our existence condition, which perhaps explains why they are unable to find a finite-state equilibrium.

Besides Makarov and Rytchkov (2006), the only other paper we are aware of that applies frequency domain methods to asset pricing is the recent work of Bernhardt, Seiler, and Taub (2005). Like us, they analyze a dynamic model with symmetric, heterogeneously informed traders. However, their work differs from ours in two important respects. First, following Kyle (1985), they focus on the strategic use of information when individual traders influence asset prices, which are set in a competitive dealership market. In contrast, our model is Walrasian. While strategic issues

[4] Kasa (2000) uses frequency domain methods to solve Townsend's model. He comes to the same conclusion as Pearlman and Sargent (2005).

are interesting, and of great practical importance, they add an additional layer of complexity to the already difficult problem of characterizing higher-order belief dynamics. Second, the complexity of their model requires numerical methods. This gives rise to the above noted difficulties associated with guessing correct functional forms and appropriate state vectors. It also makes it difficult to distinguish higher-order belief dynamics from the strategic use of private information.

Our model's solution takes the form of a nonfundamental (i.e., noninvertible) moving-average representation, mapping the underlying shocks to agents' information sets to observed prices and fundamentals. Using Blaschke factors, one can easily convert this to a Wold representation. The (statistical) innovations of the Wold representation turn out to be complicated moving averages of the entire histories of the underlying (economic) shocks. These moving averages encode the model's higher-order belief dynamics. A key virtue of our approach is that this equilibrium representation can be taken to the data in a direct, *quantitative* way. In contrast, existing work on higher-order beliefs is purely qualitative. This allows us to revisit past empirical failures of linear present value asset pricing models. We ask the following question - Suppose asset markets feature heterogeneous beliefs, but an econometrician mistakenly assumes agents have homogeneous beliefs. What will he conclude? One might think, based on the conditioning down arguments of Hansen and Sargent (1991a) and Campbell and Shiller (1987), that this would not create any problems. Interestingly, this is not the case. Conditioning down does not work here. The arguments of Hansen-Sargent and Campbell-Shiller apply to situations where agents and econometricians have different information sets. They do *not* apply in general to situations where there is informational heterogeneity among the agents themselves. This is because the law of iterated expectations does not apply to the average beliefs operator (Allen, Morris, and Shin (2006), Morris and Shin (2003)). We show that present value models with heterogeneous beliefs can easily produce violations of standard variance bounds and rejections of cross-equation restrictions. This sounds a note of caution when interpreting previous rejections of present value models. Perhaps it is not the constant discount rate that is the problem, but rather the (usually implicit) assumption of homogeneous beliefs, or equivalently, a fully revealing equilibrium.

The remainder of the paper is organized as follows. The next section outlines Futia's model. Futia showed how to solve the model in two cases: (i) when the equilibrium is fully revealing, and (ii) when information sets are hierarchical, so that some agents know strictly more than others. He showed that the hierarchical equilibrium may or may not be fully revealing. Section 3 shows how to solve the model in the more realistic case of symmetric, yet disparate, information sets. Rather than working from posited laws of motion for the fundamentals, our strategy is to work backwards from an assumed nonrevealing equilibrium to the supporting stochastic process for fundamentals. For comparison purposes, Section 4 briefly considers a full information benchmark. Section 5 constructs the equilibrium Wold representation and discusses the model's empirical implications. Section 6 contains an application to the foreign exchange market, and Section 7 concludes by discussing some extensions and applications.

2. A Model of Trade in Shares of a Risky Asset

2.1. **Model.** The model we work with below follows Futia's (1981) 'Land Speculation in Hilbert Space' setup closely. By working within the context of a well-defined Hilbert space, we are ruling out phenomena like bubbles and sunspots. All stochastic processes are restricted to be square-summable (potentially with discounting), and all equilibria are restricted to lie in the space spanned by square-summable linear combinations of past innovations to market fundamentals, which may, however, be a larger space than the space spanned by the history of observed market data. In his model, Futia considers investment in a single durable asset in fixed total supply. Demand for the asset arises from two sources; a time- and state-varying nonspeculative component (i.e., liquidity traders), and from competitive, price-taking speculators. The presence of liquidity traders adds noise to the model, and serves to break the no-trade theorem that would otherwise apply in Futia's hierarchical information setup. It is assumed that nonspeculative demand never exceeds total supply, and the residual, denoted f_t, is interpreted as 'market fundamentals.' Each investor has a demand for the asset given by,

$$q_t^i = E_t^i p_{t+1} - \beta^{-1} p_t \tag{2.1}$$

where $\beta^{-1} = (1+r) > 1$ is interpreted as the opportunity cost of funds. The important thing to note here is that conditional expectations are indexed by agents, recognizing the fact that information sets differ. Equation (2.1) simply says that demand is an increasing function of the difference between the expected capital gain on the asset and the opportunity cost of the funds. It is not infinite, however, due to risk aversion. In fact, equation (2.1) can be derived from a simple portfolio choice problem in which agents have exponential (CARA) preferences and a one-period investment horizon. From this perspective, equation (2.1) implicitly normalizes to unity the product of the coefficient of absolute risk aversion and the (constant) conditional variance of the price.[5]

Equating aggregate speculative demand to aggregate speculative supply delivers the following market-clearing condition:

$$p_t = \beta \int_0^1 E_t^i p_{t+1} di - \beta f_t \tag{2.2}$$

where it has been assumed that there is a measure one continuum of speculative traders. Notice, following Singleton (1987), that the market price of the risky asset depends on a *weighted average* of the market participants' forecasts of p_{t+1}. Thus

[5]Whiteman (1989) shows how to solve the model with the conditional variances retained. Doing so would complicate the algebra, but would not change the results qualitatively.

each agent's forecast of p_{t+1} depends on his forecast of the market-wide weighted average forecast of p_{t+2}, and so on.[6] Evidently, there is an infinite regress in expectations, a problem encountered in a different context by Townsend (1983). Townsend and Singleton "broke" the infinite regress by completely revealing the state of the economy, albeit with a lag. Such divine revelation leaves only a few objects unknown at any date, and makes the regress problem manageable. In this paper, the infinite regress problem is never broken (it is treated just like any other equilibrium problem) and there is no divine revelation; special assumptions about the nature of the informational heterogeneity keep the problem manageable.

Given a stochastic process for the fundamentals, and assuming rational expectations, equation (2.2) determines the equilibrium stochastic process for prices. Note, however, that with the appropriate definition of f_t, (2.2) is quite general. For example, with $-f_t$ defined as dividends, it becomes a present value model for stock prices; with $-f_t$ defined as the difference between national money supplies and income levels it becomes the monetary model of exchange rates; with $-f_t$ defined as a short-term interest rate it becomes the expectations hypothesis of the term structure; and so on.

In the analysis that follows, we consider two cases of this model. The first case assumes that fundamentals are latent from (not observed by) traders. This is the setup of Futia (1981). The second case assumes that net supply is identically equal to zero (i.e., no liquidity traders), and fundamentals are observable. This model follows the interpretation of the present value model for stock prices. The models have similar informational structures and lead to the same empirical conclusions for asset prices.

2.2. **Information.** In specifying the nature of uncertainty and the structure of information, we assume that the world is driven by an m-vector of serially and mutually independent Gaussian N(0,1) random variables $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, ..., \varepsilon_{mt})$. *Admissible* random variables are linear combinations of current, past, and future values of $\{\boldsymbol{\varepsilon_t}\}$ that have square-summable coefficients. The set H of all admissible random variables is a well-known Hilbert space; $f_t, p_t, E_t^i(\cdot) \in H$ for all t. We also restrict equilibria to lie in H, which explicitly rules out bubbles and sunspot equilibria.

The common information at date t consists of past values of the price of the risky asset. The space spanned by square-summable linear combinations of past values of p_t is denoted by $H_p(t)$. The exogenous private information set of agent i at time t is a subset U_t^i of H satisfying $U_t^i \subseteq U_{t+1}^i$; U_t^i is the space spanned by square-summable linear combinations of current and past values of the variables other than

[6]As recently emphasized by Allen, Morris, and Shin (2006), the law of iterated expectations does *not* in general apply to the average expectations operator. One can interpret this result using Hansen and Sargent (1991a) notion of an 'exact' rational expectations model. Traditional homogenous beliefs present value models are examples of exact models. Although agents may have more information than econometricians, there are no 'missing fundamentals' from these models. In contrast, models with heterogeneous beliefs are 'inexact', since higher-order beliefs in effect become missing fundamentals.

p_t seen by agent i. Let J_t denote the space spanned by all exogenous information at time t contained in the model $(\cup_{i=1}^{\infty} U_t^i)$. Given these assumptions, the conditional expectations are given by

$$E_t^i p_{t+1} = \Pi[p_{t+1} | U_t^i \bigvee H_p(t)], \tag{2.3}$$

where Π denotes linear least squares projection, and $X \bigvee Y$ is standard notation for "the linear space spanned by X and Y." We now define a rational expectations equilibrium (REE).

Definition 2.1. *A rational expectations equilibrium is a stochastic process* $\{p_t\}$ *for the price of the risky asset which satisfies* (2.2) *with conditional expectations formed according to* (2.3) *and* $p_t \in J_t \bigvee H_p(t)$.

The last condition is what Futia referred to as the 'no divine revelation' clause. That is, the equilibrium price cannot rely on information that originates from outside of the model. The equilibrium price must lie in the space spanned by past prices and the exogenous information known by the traders.

Definition 2.2. *The REE is symmetric if it is a rational expectations equilibrium in which all agents make identical forecasts,*

$$E_t^i p_{t+1} = E_t^j p_{t+1}, \qquad \forall \;\; i \text{ and } j.$$

Notice that by observing "action" in the market, agents may glean information not in their own private information sets that helps predict market fundamentals. To preserve the asymmetric information structure in equilibrium, it must be the case that privately held information is not revealed by observation of current and past prices.

Agents in the model have asymmetric information regarding the stochastic process of fundamentals, f_t. Without loss of generality, we assume fundamentals consist of the sum of orthogonal components:

$$f_t = \sum_{i=1}^{m} a_i(L) \varepsilon_{it}. \tag{2.4}$$

It will be assumed that $a_i(L)$ is a polynomial (of possibly infinite order) in nonnegative powers of the lag operator L, with square summable coefficients, and that $a_i(z) \neq 0$ for any $|z| \leq 1$ and $a_i(z) \neq a_j(z)$ for any $i \neq j$. The orthogonal shocks are assumed to be serially uncorrelated with $E\varepsilon_{it}\varepsilon_{js} = 0$ for all $i \neq j$, and constitute the fundamental economic building blocks of the model, implying the price process will be square summable sequences of $\{\varepsilon_i\}$.

We assume there are two types of speculative traders. Type 1 traders costlessly observe realizations of ε_{1t}, while type 2 traders costlessly observe realizations of ε_{2t}. Without loss of generality, we assume equal shares of the two types of traders. This allows us to write (2.2) as

$$p_t = \beta\left\{\frac{1}{2}E^1[p_{t+1}|H_p(t)\bigvee H_{\varepsilon_1}(t)] + \frac{1}{2}E^2[p_{t+1}|H_p(t)\bigvee H_{\varepsilon_2}(t)]\right\} - \beta f_t. \tag{2.5}$$

Of course when fundamentals are observable, each traders' conditional expectation will include $H_f(t)$.

3. Constructing a Nonrevealing Equilibrium

One of the main contributions of the paper is to establish conditions under which the disparate expectations of (2.5) are preserved in a dynamic equilibrium. As mentioned above, the usual approach for solving rational expectations models (i.e., parameterize a conjectured law of motion, apply the Kalman filter to evaluate the conditional expectations, and then match coefficients) cannot be employed due to the complications of infinite regress. In models with asymmetric information, other traders' forecasts of future prices affect the current price of the asset, and therefore these forecasts are relevant state variables. But in a dynamic setting, this relevant state variable would become infinitely large because traders must forecast the average forecast of the average forecast of ..., *ad infinitum*.

In solving the model, we apply the following solution method. First, each trader uses all available information at time t to form beliefs about the current price process. Second, every trader behaves optimally and the conditional expectation of p_{t+1} will be calculated via Wiener-Kolmogorov optimal prediction formulas. Third, the appropriate form of equation (2.5) is then used to impose market clearing. In solving the subsequent fixed-point problem, we appeal to the Riesz-Fischer Theorem and derive the solution in the frequency domain. However, this process will only generate a *candidate* equilibrium price process. Traders will surely condition on past prices, so the candidate is in fact a rational expectations equilibrium provided it does not reveal any additional information beyond what was initially assumed.

We consider two separate assumptions about the fundamentals. First we assume aggregate fundamentals are unobservable. This is perhaps most descriptive of macroeconomic applications, e.g., present value models of the exchange rate, where relevant aggregate fundamentals may not be known or reported.[7] It turns out that prices must be revealing in the case when there are just two trader-types. To support a heterogeneous beliefs equilibrium with observed aggregate fundamentals, we therefore extend the analysis to three trader-types.

[7]Engel and West (2005) argue that unobserved fundamentals appear to be necessary to reconcile present value models with observed exchange rates. Hamilton and Whiteman (1985) argue that the mere *possibility* of unobserved fundamentals vitiates standard bubbles tests. Interestingly, our results suggest a likely candidate for 'missing fundamentals', i.e., higher-order beliefs.

3.1. **Latent Fundamentals.** We begin with Futia's set-up. Aggregate fundamentals are

$$f_t = a_1(L)\varepsilon_{1t} + a_2(L)\varepsilon_{2t},$$

where the polynomials $a_1(L)$ and $a_2(L)$ are taken as given (subject to some restrictions given below). While no trader sees aggregate fundamentals directly, each type of trader sees a stochastic process that is *correlated* with f_t; specifically, type 1 traders see realizations of ε_{1t}, while type 2 traders see ε_{2t}. In equilibrium, the information set of type 1 traders is given by current and past values of the stochastic process ϵ_{1t}, p_t, having the moving average representation

$$\begin{bmatrix} \varepsilon_{1t} \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \pi_1(L) & \pi_2(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix}$$
$$\mathbf{x}_t = \mathbf{M}(L)\boldsymbol{\epsilon}_t. \tag{3.1}$$

where the $\pi_1(L)$ and $\pi_2(L)$ polynomials are to be determined from the equilibrium conditions of the model.

If the equilibrium is to be nonrevealing, it must be the case that the $\pi_i(L)$ polynomials are noninvertible (in non-negative powers of L). Otherwise, an agent of one type equipped with *basic* statistical knowledge (e.g., knowledge of VARs) could infer the other type of agent's information from observations of the price and his own shock realizations. That is, if the $\pi_i(L)$ polynomials were invertible, $H_{\varepsilon_1}(t) \bigvee H_p(t)$ and $H_{\varepsilon_1}(t) \bigvee H_{\varepsilon_2}(t)$ would coincide, and the equilibirum would reveal agent 2's information to agent 1. This noninvertibility restriction corresponds to the assumption that $\pi_1(z)$ and $\pi_2(z)$ have zeroes inside the unit circle. Thus to preserve asymmetric information in equilibrium we must seek and find equilibrium pricing polynomials of the form

$$\pi_i(L) = (L - \lambda)[\rho_i + Lg_i(L)] \tag{3.2}$$

where we now require $|\lambda| < 1$, and that the $(\rho_i + zg_i(z))$ functions are analytic and without zeroes inside the unit circle.[8] That is, we assume the pricing function has a single zero inside the unit circle and seek a λ, ρ_i, and $g_i(L)$ that satisfies the above conditions.

The following lemma ensures that pricing functions of the form (3.2) imply that by observing current and past prices and current and past realizations of ε_{1t}, type 1 traders will not be able to infer ε_{2t}, the private information of type 2 traders.

Lemma 3.1. *The moving average representation given by* (3.1) *and* (3.2) *is not a fundamental (Wold) representation.*

Proof. A necessary and sufficient condition for (3.1) to be a Wold representation is that the space spanned by past observables $\mathbf{x}_t$ must be equivalent, in the mean square

[8]We can relax the assumption that π_1 and π_2 have the same zero (λ) at the cost of losing Proposition 1 below. The upside is that the analogue of Assumption 1 would be more easily met.

sense, to the space spanned by $\boldsymbol{\epsilon}_t$. This requires $\mathbf{M}(L)$ to have a one-sided inverse in non-negative powers of L. A necessary condition for the existence of this inverse is that the determinant of $\mathbf{M}(L)$ cannot have any zeros inside the unit circle. By direct calculation,

$$\det \mathbf{M}(L) = (L-\lambda)[\rho_2 + Lg_2(L)]$$

has a zero inside the unit circle at λ. □

If type 1 traders do not observe ε_{2t} directly, what information do they possess in evaluating the conditional expectation? Type 1 traders have information set $H_p(t) \bigvee H_{\varepsilon_1}(t)$, and therefore any information gleaned from the past sequences of $\{\varepsilon_{1t}\}$ and the price process $\{p_t\}$ will be used to evaluate the conditional expectation. While p_t and ε_{1t} can be expressed as a square-summable linear combination of current and past values of ε_{1t} and ε_{2t}, the converse is not true; $H_{\varepsilon_1}(t) \bigvee H_{\varepsilon_2}(t)$ spans a larger space than $H_{\varepsilon_1}(t) \bigvee H_p(t)$. Therefore, in order to evaluate the expectations of type 1 traders, we need to restrict attention to the subspace generated by $H_p(t) \bigvee H_{\varepsilon_1}(t)$.

There are two ways to do this. The first is to work directly with the subspace seen by the agent by employing Blaschke factors to "flip zeros" outside the unit circle. The second is to work with the larger information set involving current and past values of ε_{2t} and then project this into $H_{\varepsilon_1}(t) \bigvee H_p(t)$. Both ways are instructive.

The direct method involves employing Blaschke factors to find the unique (up to a constant) fundamental representation associated with (3.1), which is given by

$$\underbrace{\begin{bmatrix} \varepsilon_{1t} \\ p_t \end{bmatrix}}_{\mathbf{x}_t} = \underbrace{\begin{bmatrix} 1 & 0 \\ (L-\lambda)[\rho_1 + Lg_1(L)] & (L-\lambda)[\rho_2 + Lg_2(L)] \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \frac{1-\lambda L}{L-\lambda} \end{bmatrix}}_{\mathbf{M}^*(L)} \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & \frac{L-\lambda}{1-\lambda L} \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix}}_{\boldsymbol{\epsilon}_t^*} \tag{3.3}$$

The Blaschke factor $[(L-\lambda)/(1-\lambda L)]$ transforms representation (3.1) into a fundamental representation. If we define

$$e_{2t} \equiv \left[\frac{L-\lambda}{1-\lambda L}\right] \varepsilon_{2t}, \tag{3.4}$$

then for type 1 traders, knowledge of $\{p_t\}$ is equivalent to knowledge of $\{e_{2t}\}$, and not $\{\varepsilon_{2t}\}$. Moreover, notice that from (3.4), it is apparent that knowledge of current and past ε_{2t} is sufficient for e_{2t}, but that the inverse of the Blaschke factor does not possess a valid expansion in and on the unit circle in L due to the pole at $L = |\lambda|$. However, by setting $F = L^{-1}$, it is easy to see that the Blaschke factor does have a valid inverse in the forward operator F

$$\left[\frac{F-\lambda}{1-\lambda F}\right] e_{2t} = \varepsilon_{2t}, \qquad \varepsilon_{2t} = (L^{-1} - \lambda) \sum_{j=0}^{\infty} \lambda^j e_{2,t+j}.$$

In other words, ε_{2t} carries information about *future* e_2's.

The parameter λ may be interpreted as an *information wedge.* Notice that if $|\lambda| > 1$, (3.1) becomes a fundamental representation and asymmetric information will not be preserved in equilibrium. By observing current and past prices, traders of both types will be able to infer the information of the other type by applying VAR analysis. If $|\lambda| < 1$, then the equilibrium prices will not reveal information. The subsequent excess volatility result discussed in Section 5 hinges upon the price process being a nonrevealing equilibrium.

The conditional expectations (2.5) can then be found by using Wiener-Kolmogorov optimal prediction formulas. That is,

$$\begin{aligned} E(\mathbf{x}_{t+1}) &= L^{-1}[\mathbf{M}^*(L) - \mathbf{M}(0)^*]\boldsymbol{\epsilon}_t^* \\ E^1[p_{t+1}|H_p(t) \bigvee H_{\varepsilon_1}(t)] &= [\rho_1 + (L-\lambda)g_1(L)]\varepsilon_{1t} \\ &+ \left[\rho_2 + (L-\lambda)g_2(L) - \frac{\rho_2(1-\lambda^2)}{1-\lambda L}\right]\varepsilon_{2t} \qquad (3.5) \end{aligned}$$

The second way to determine the needed conditional expectation involves using a 'conditioning down' argument. First, project p_{t+1} onto the space $H_{\varepsilon_1}(t) \bigvee H_{\varepsilon_2}(t)$. That is, assume (counterfactually) that agents directly observe realizations of the underlying shocks, ε_{1t} and ε_{2t}. Subsequently, we will condition down by 'subtracting off' the appropriate orthogonal complements.[9] From the orthogonality of ε_1 and ε_2, and the Wiener-Kolmogorov prediction formula we have the projection onto the larger space given by:

$$\begin{aligned} E[p_{t+1}|H_{\varepsilon_1}(t) \bigvee H_{\varepsilon_2}(t)] = &\quad L^{-1}[\pi_1(L) - \pi_1(0)]\varepsilon_{1t} + L^{-1}[\pi_2(L) - \pi_2(0)]\varepsilon_{2t} \\ = &\quad [\rho_1 + (L-\lambda)g_1(L)]\varepsilon_{1t} + [\rho_2 + (L-\lambda)g_2(L)]\varepsilon_{2t} \quad (3.6) \end{aligned}$$

As before, since $|\lambda| < 1$, then $H_{\varepsilon_1}(t) \bigvee H_{\varepsilon_2}(t)$ is a larger space than $H_{\varepsilon_1}(t) \bigvee H_p(t)$, and we need to condition down (i.e., project onto the subspace $H_{\varepsilon_1}(t) \bigvee H_p(t)$). Clearly, the first term on the right-hand side of (3.6) can be retained, since type 1 traders observe ε_{1t}. The second term, however, needs to be modified. Due to Beurling's Theorem, Blaschke factors play a fundamental role in constructing orthogonal projections and invariant subspaces of analytic functions.[10] Our required projection follows as a special case of the following theorem:

[9] Although we cautioned earlier that the law of iterated expectations does not apply to the average expectations operator, it certainly *does* apply to each individual's forecasting problem.

[10] Invariant subspaces can be thought of as playing the role of eigenvectors in infinite dimensional function spaces. Just as it helps to visualize the action of a matrix by visualizing its projections onto orthogonal eigenvectors, it helps to visualize the action of a z-transform by visualizing its (shift) invariant subspaces. Loosely speaking, Beurling's theorem tells us that shift invariant subspaces of analytic functions consist of Blaschke products. (See theorem 3.9 of Radjavi and Rosenthal (1973) for a statement and proof of Beurling's theorem).

Theorem 3.2 (Theorem 3.14 in Radjavi and Rosenthal (1973))**.** *Let D denote the open unit disk, $\mathcal{H}^2$ denote the Hardy space of square integrable analytic functions on D. If $\lambda_1 \cdots \lambda_n$ are in D and if*

$$\phi(z) = \prod_{j=1}^{n} \frac{\lambda_j - z}{1 - \bar{\lambda}_j z} \tag{3.7}$$

then the orthogonal complement of $\phi\mathcal{H}^2$ in $\mathcal{H}^2$, $\mathcal{H}^2 \ominus \phi\mathcal{H}^2$, has dimension n. Conversely, every invariant subspace of S (i.e., a shift operator) of co-dimension n has this form.

The function $\phi(z)$ in (3.7) is an example of a Blaschke product. Note that $|\phi(z)| = 1$ on D, implying that multiplication by $\phi(z)$ is norm-preserving. The theorem implies that in our case, with $n = 1$, the part of $H_{\varepsilon_2}(t)$ that cannot be written as a linear combination of current and past e_{2t} is unidimensional–a single square summable linear combination of current and past values of ε_{2t}. In particular, Theorem 3.2 yields the following result:

Lemma 3.3. *The projection $E[p_{t+1}|H_{\varepsilon_2}(t)] = [\rho_2 + (L-\lambda)g_2(L)]\varepsilon_{2t}$ has the following orthogonal decomposition:*

$$[\rho_2 + (L - \lambda)g_2(L)]\varepsilon_{2t} = \left[h(L)\frac{L-\lambda}{1-\lambda L}\right]\varepsilon_{2t} + \frac{constant}{1-\lambda L}\varepsilon_{2t} \tag{3.8}$$

where $h(L)$ is an analytic function in D with zeroes outside D.

Proof. We begin with the decomposition of ε_{2t} itself into a component that can be written in terms of $e_{2t} = \left[\frac{L-\lambda}{1-\lambda L}\right]\varepsilon_{2t}$ and another orthogonal to it:

$$\varepsilon_{2t} = K(L)\left[\frac{L-\lambda}{1-\lambda L}\right]\varepsilon_{2t} + M(L)\varepsilon_{2t} \tag{3.9}$$

where $K(L)$ and $M(L)$ are one-sided polynomials in the lag operator with square-summable coefficients. Orthogonality is enforced by the requirement that $M(L)\varepsilon_{2t} = \varepsilon_{2t} - K(L)\left[\frac{L-\lambda}{1-\lambda L}\right]\varepsilon_{2t}$ be orthogonal to $e_{2t} = \frac{L-\lambda}{1-\lambda L}\varepsilon_{2t}, e_{2t-1} = \frac{L-\lambda}{1-\lambda L}\varepsilon_{2t-1}$, e_{2t-2}, etc. Thus

$$EM(L)\varepsilon_{2t}e_{2t-j} = \frac{1}{2\pi i}\oint \frac{M(z)z^{-j}(z-\lambda)}{1-\lambda z}\frac{dz}{z} = 0, \qquad j = 0, 1, 2, ...$$

Direct calculation using the residue calculus yields the restrictions $M(\lambda) = M_0/(1-\lambda^2)$ and $M_j = \lambda M_{j-1}$ for $j \geq 1$. This implies $M(L) = M_0/(1-\lambda L)$. This immediately gives (3.8) for some $h(z)$. It is straightforward to verify the orthogonality of the two components on the RHS of (3.8) for any analytic $h(z)$[11]:

$$\begin{aligned}\left(h(z)\frac{z-\lambda}{1-\lambda z}, \frac{1}{1-\lambda z}\right) &= \frac{1}{2\pi i}\oint h(z)\frac{z-\lambda}{1-\lambda z}\cdot\frac{1}{1-\lambda z^{-1}}\frac{dz}{z} \\ &= \frac{1}{2\pi i}\oint h(z)\frac{z-\lambda}{1-\lambda z}\cdot\frac{z}{z-\lambda}\frac{dz}{z} \\ &= \frac{1}{2\pi i}\oint h(z)\frac{1}{1-\lambda z}dz \\ &= 0 \qquad \text{(by Cauchy's integral formula)}\end{aligned}$$

□

The usefulness of the decomposition in Lemma 3.3 derives from the fact that it isolates exactly what type 1 traders *cannot* infer about ε_{2t} from observations of market data. To determine the constant, simply equate both sides at $L = \lambda$, which gives the constant as $\rho_2(1-\lambda^2)$. The orthogonal decomposition yields the conditional expectations.

Lemma 3.4. *Given the hypothesized pricing functions in* (3.2)*, the conditional expectations of type 1 traders are given by:*

$$E[p_{t+1}|H_p(t)\bigvee H_{\varepsilon_1}(t)] = [\rho_1+(L-\lambda)g_1(L)]\varepsilon_{1t}+\left[\rho_2+(L-\lambda)g_2(L)-\frac{\rho_2(1-\lambda^2)}{1-\lambda L}\right]\varepsilon_{2t} \tag{3.10}$$

and the conditional expectations of type 2 traders are given by:

$$E[p_{t+1}|H_p(t)\bigvee H_{\varepsilon_2}(t)] = \left[\rho_1+(L-\lambda)g_1(L)-\frac{\rho_1(1-\lambda^2)}{1-\lambda L}\right]\varepsilon_{1t}+[\rho_2+(L-\lambda)g_2(L)]\varepsilon_{2t} \tag{3.11}$$

At this point it is useful to compare (3.10) and (3.11) to (3.6). Notice how conditioning down onto the observable subspaces requires us to 'subtract' the orthogonal complement of the subspace generated by a Blaschke factor constructed with the presumed noninvertible root, λ. It is also useful to compare (3.10) and (3.11). The following proposition demonstrates how the asymmetric information structure can lead to 'overreaction' in financial markets.

Proposition 3.5. *Traders respond 'more aggressively' to realizations of other traders' (unobserved) signals.*

Proof. This can be seen by differencing (3.10) and (3.11),

$$E_t^1 p_{t+1} - E_t^2 p_{t+1} = \frac{1-\lambda^2}{1-\lambda L}\left(\rho_1\varepsilon_{1t}-\rho_2\varepsilon_{2t}\right). \tag{3.12}$$

[11]For Hilbert space aficionados, notice that the second term on the right-hand side of (3.8) is a reproducing kernel for $\mathcal{H}^2$, since for any $f \in \mathcal{H}^2$ we have $(f(z),(1-\lambda z)^{-1}) = f(\lambda)$. This is not accidental.

□

At any given point in time, the forecasts of the two traders differ as a function of the histories of their observed signals. Hence, the trader who has received higher signals (on average) tends to forecast higher prices. Interestingly, traders respond 'more aggressively' to realizations of other traders' (unobserved) signals. For example, if trader 2's signals have been larger on average, trader 1 will have lower price forecasts (remember, prices respond negatively to the ε_i's, since they represent shocks to supply). Moreover since knowledge of the other traders' signals is only obtainable by observing publicly-available prices, this result is tantamount to overreaction to public signals (Allen, Morris, and Shin (2006)). As we will see, this overreaction to public signals generates 'excess volatility'.

The third step in the solution process is to impose the equilibrium condition (2.5) and solve the subsequent fixed-point problem. Doing this yields:

Proposition 3.6. *Under Assumption 3.7 (given below), there exists a unique heterogeneous beliefs rational expectations pricing function for the model given in* (2.5), *with z-transforms given by:*

$$\pi_1(z) =(z-\lambda)\left[2a_1(\lambda)+\frac{z}{z-\beta}\left\{-2a_1(\lambda)+\frac{\beta}{z-\lambda}\left[2a_1(\lambda)-a_1(z)-\frac{a_1(\lambda)(1-\lambda^2)}{1-\lambda z}\right]\right\}\right] \tag{3.13}$$

$$\pi_2(z) =(z-\lambda)\left[2a_2(\lambda)+\frac{z}{z-\beta}\left\{-2a_2(\lambda)+\frac{\beta}{z-\lambda}\left[2a_2(\lambda)-a_2(z)-\frac{a_2(\lambda)(1-\lambda^2)}{1-\lambda z}\right]\right\}\right] \tag{3.14}$$

and $|\lambda|<1$ *given implicitly by the equation:* $2\lambda a_1(\lambda)=\beta[a_1(\beta)+a_1(\lambda)(1-\lambda^2)/(1-\lambda\beta)]$.

The proof is by construction. Notice that it is sufficient to verify the result for $\pi_1(z)$, due to symmetry. The equilibrium condition (2.5) and conditional expectations (3.10) and (3.11) gives

$$(L-\lambda)(\rho_1+Lg_1(L))\varepsilon_{1t}=\beta\left[\rho_1+(L-\lambda)g_1(L)-\frac{1}{2}\frac{\rho_1(1-\lambda^2)}{1-\lambda L}\right]\varepsilon_{1t}-\beta a_1(L)\varepsilon_{1t}.$$

Assuming that this expression holds for all realizations of ε_{1t}, the coefficients on ε_{1s} must match for every s. In lieu of solving this infinite sequential problem, one can solve an equivalent functional problem by invoking the Riesz-Fischer Theorem and examining the corresponding power series equalities[12]

$$(z-\lambda)(\rho_1+zg_1(z))=\beta\left[\rho_1+(z-\lambda)g_1(z)-\frac{1}{2}\frac{\rho_1(1-\lambda^2)}{1-\lambda z}\right]-\beta a_1(z). \tag{3.15}$$

[12]The Appendix provides more detail concerning this solution method.

Evaluating (3.15) at $z = \lambda$ immediately delivers the unknown constant, $\rho_1 = 2a_1(\lambda)$. To determine λ in terms of the exogenous parameter, plug in ρ_1, divide both sides by $z - \lambda$, and then collect terms. This yields,

$$(z-\beta)g_1(z) = -2a_1(\lambda) + \frac{\beta}{z-\lambda}\left[2a_1(\lambda) - a_1(z) - \frac{a_1(\lambda)(1-\lambda^2)}{1-\lambda z}\right] \tag{3.16}$$

Notice that the right-hand side is analytic by construction (i.e., the singularity at λ has been 'removed'). Since $g_1(z)$ has been assumed to be analytic, the right-hand side of (3.16) must be zero when evaluated at $z = \beta$. Evaluating the right-hand side at $z = \beta$ and setting it to zero gives us the following equation characterizing λ

$$2\lambda = \beta\left[\frac{a_1(\beta)}{a_1(\lambda)} + \frac{1-\lambda^2}{1-\lambda\beta}\right], \tag{3.17}$$

which is a slight re-arrangement of the equation given in 3.6. Notice that in general λ will depend on $a_1(z)$, and thus we can expect a *different* λ when solving the fixed point equation for $\pi_2(z)$. Since we've postulated a common λ for both $\pi_1(z)$ and $\pi_2(z)$, we impose the following assumption.[13]

Assumption 3.7. *There exists a unique $|\lambda| < 1$ with $\lambda \neq \beta$, that solves the two equations:*

$$2\lambda = \beta[a_i(\beta)/a_i(\lambda) + (1-\lambda^2)/(1-\lambda\beta)] \qquad i = 1, 2$$

A trivial case where this is generally satisfied is when the dynamics of the two unobserved components are the same (i.e., when $a_1(L) = a_2(L)$). However in order to avoid a stochastic singularity in the bivariate representation for prices and fundamentals, this case is ruled out. Moreover $a_1(L)$ and $a_2(L)$ cannot both be AR(1) processes, since in this case there is only one unique AR root for any given λ. However, it is not difficult to find representations satisfying Assumption 3.7. For example, an ARMA(1,1) will satisfy the condition.

Finally, we can determine $g_1(z)$ by dividing both sides of (3.16) by $z-\beta$ (remember, by construction, the singularity at β has just been removed by the appropriate choice of λ). Given $g_1(z)$, λ, and ρ_1, the expression for $\pi_1(z)$ given by (3.13) follows from plugging into $\pi_1(z) = (z-\lambda)(\rho_1 + zg_1(z))$.

3.2. **Observable Fundamentals.** Although some asset markets might be well described by unobserved aggregate fundamentals, in other cases it makes more sense to assume aggregate fundamentals *are* observed. A leading example would be the prices of individual stocks and bonds, where earnings and dividends are widely reported.[14]

[13] A common λ in the price process is not required for the results presented in Section 5; thus Assumption 1 is not, in general, restrictive. Assumption 3.7 is the bivariate generalization of Futia's (1981) Theorem 6.1.

[14] Remember, we are imposing a constant discount rate from the outset, so the usual difficulty of identifying the macroeconomic determinants of stochastic discount factors do not apply here.

The next result shows that when aggregate fundamentals are observed, equilibrium prices *must* be revealing, i.e., a heterogeneous beliefs equilibrium does not exist.

Proposition 3.8. *With just two trader types, there does not exist a heterogeneous beliefs Rational Expectations Equilibrium if aggregate fundamentals are observable. Prices must be fully revealing.*

Proof. Suppose Type 1 traders observe $(p_t, f_t, \varepsilon_{1t})$ and Type 2 traders observe $(p_t, f_t, \varepsilon_{2t})$. Then Type 1 effectively observes $a_2(L)\varepsilon_{2t}$. Since p_t also depends on the history of ε_{2t}, the first corollary on p. 101 of Hoffman (1962) implies $\{p_t, a_2(L)\varepsilon_{2t}\}$ spans $H_{\varepsilon_2}(t)$ unless $a_2(z)$ and $\pi_2(z)$ have identical noninvertible roots. However, $a_2(\lambda) \neq 0$ by the existence condition given in Assumption 3.7. □

There are several ways we could modify the model to support a heterogeneous beliefs equilibrium. The most natural is to simply add more trader-types. The two trader-type specification is very special, since in this case the number of types exactly matches the number of aggregate observables. We now show that with more than two types a heterogeneous beliefs equilibrium can be robustly supported even when aggregate fundamentals are common knowledge.

Suppose now that fundamentals (2.4) are comprised of *three* orthogonal components,

$$f_t = a_1(L)\varepsilon_{1t} + a_2(L)\varepsilon_{2t} + a_3(L)\varepsilon_{3t}$$

and that the asymmetric information is the same as the previous section. That is, type 1 traders costlessly observe ε_{1t}, type 2 traders costlessly observe ε_{2t}, and type 3 traders costlessly observe ε_{3t}. Conjecturing the same pricing function for each component as before then implies the following observer system for type 1 traders,

$$\begin{bmatrix} \varepsilon_{1t} \\ f_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ a_1(L) & a_2(L) & a_3(L) \\ (L-\lambda)[\rho_1 + Lg_1(L)] & (L-\lambda)[\rho_2 + Lg_2(L)] & (L-\lambda)[\rho_3 + Lg_3(L)] \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \\ \varepsilon_{3t} \end{bmatrix}$$
$$\mathbf{y}_t = \mathbf{H}(L)\boldsymbol{\epsilon}_t$$

Type 1 traders use this system to compute,

$$E^1(p_{t+1}|H_p(t) \bigvee H_f(t) \bigvee H_{\varepsilon_1}(t)) = E^1(p_{t+1}|H_{\varepsilon_1}(t) \bigvee H_{e_2}(t) \bigvee H_{e_3}(t)),$$

Notice that this now entails the use of *two* (identical) Blaschke factors, one mapping ε_{2t} to e_{2t}, and another mapping ε_{3t} to e_{3t}. A completely symmetric argument applies to traders 2 and 3.

Of course, the existence condition in Assumption 3.7 must be modified to reflect the presence of the additional component. We now have

Assumption 3.9. *There exists a unique $|\lambda| < 1$ with $\lambda \neq \beta$, that solves the three equations:*

$$3\lambda = \beta[a_i(\beta)/a_i(\lambda) + 2(1-\lambda^2)/(1-\lambda\beta)] \qquad i = 1,2,3$$

In general, this is a more restrictive condition, but again, it is not hard to find examples that satisfy it. Given Assumption 3.9, the entire analysis goes through as before, with nearly identical pricing functions.[15]. The only real difference is that now the overall price process consists of the sum of *three* symmetric pricing functions. Hence, in what follows we focus on the latent fundamentals case, recognizing that any results obtained can be translated to the observed fundamentals case, assuming the somewhat more restrictive existence condition in Assumption 3.9.

4. Equilibrium with Homogeneous Beliefs

To interpret the heterogeneous beliefs equilibrium given by equations (3.13) and (3.14), it is useful to consider the benchmark case of homogeneous beliefs. The appropriate modification of Assumption 3.7 allows for a fully revealing equilibrium.

Assumption 4.1. *There exists a $|\lambda| > 1$ that solves the two equations:*

$$2\lambda = \beta[a_i(\beta)/a_i(\lambda) + (1-\lambda^2)/(1-\lambda\beta)] \qquad i = 1,2$$

Given Assumption 4.1, the mapping in (3.1) has a one-sided inverse and by observing equilibrium prices, traders are able to infer the fundamental shocks ε_{1t} and ε_{2t}. Traders will then guess the equilibrium price to be of the form

$$\pi_i^s(L)\varepsilon_{it} = [\rho_i + Lh_i(L)]\varepsilon_{it}$$

where the superscript s reminds us that we are solving for a fully revealing, *symmetric*, equilibrium. Using this in (3.2) now delivers the fixed point conditions:

$$\rho_i + zh_i(z) = \beta h_i(z) - \beta a_i(z) \qquad i = 1,2 \tag{4.1}$$

Collecting terms gives,

$$(z-\beta)h_i(z) = -\rho_i - \beta a_i(z) \tag{4.2}$$

Removing the singularity at $z = \beta$ then determines, $\rho_i = -\beta a_i(\beta)$. Substituting this back into (4.2) then gives

$$h_i(z) = \frac{-\beta}{z-\beta}\left[a_i(z) - a_i(\beta)\right]$$

which finally gives us:

[15]The 2's appearing in (3.13) and (3.14) become 3's, and the $a_i(\lambda)(1-\lambda^2)$ terms become $2a_i(\lambda)(1-\lambda^2)$

Proposition 4.2. *Given Assumption 4.1, there exists a homogeneous beliefs rational expectations equilibrium given by:*

$$\pi_i^s(z) = -\beta\left\{a_i(\beta) + \frac{z}{z-\beta}\left[a_i(z) - a_i(\beta)\right]\right\} \qquad i = 1, 2 \tag{4.3}$$

Although it may not be immediately apparent, equation (4.3) is a familiar result– in a fully revealing, homogeneous expectations equilibrium, asset prices have an innovation variance that is increasing in the persistence of the fundamentals. To see this, note that

$$\begin{aligned} E_{t-1}p_t &= h_1(L)\varepsilon_{1t-1} + h_2(L)\varepsilon_{2t-1} \\ &= \frac{-\beta L}{L-\beta}\left\{[a_1(L) - a_1(\beta)]\varepsilon_{1t} + [a_2(L) - a_2(\beta)]\varepsilon_{2t}\right\} \end{aligned}$$

Using this with (4.3) then yields

$$p_t = E_{t-1}p_t - \beta a_1(\beta)\varepsilon_{1t} - \beta a_2(\beta)\varepsilon_{2t}. \tag{4.4}$$

Hence, price innovations represent the capitalized value of the innovations to fundamentals.

Equation (4.3) is useful because it facilitates interpretation of the heterogeneous expectations equilibrium in equations (3.13) and (3.14). Maintaining comparability between the two equilibria requires a given stochastic process that is consistent with *both* a heterogeneous and a homogeneous expectations equilibrium. Clearly, this will not be true in general because what is needed is a specification that is simultaneously consistent with Assumptions 3.7 and 4.1. We state this explicitly as:

Assumption 4.3. *There exist $a_1(L)$ and $a_2(L)$ polynomials that simultaneously satisfy Assumptions 3.7 and 4.1.*

This delivers the following relationship between heterogeneous expectations and homogeneous expectations equilibria.

Proposition 4.4. *Given Assumption 4.3, there exists both a heterogeneous expectations equilibrium and a homogeneous expectations equilibrium, with z-transforms related as follows:*

$$\pi_i(z) = \pi_i^s(z) + a_i(\lambda)(1-\lambda^2)\cdot\frac{\beta}{z-\beta}\left(\frac{\beta}{1-\lambda\beta} - \frac{z}{1-\lambda z}\right) \qquad i = 1, 2 \tag{4.5}$$

where the $\pi_i^s(z)$ are given by (4.3), and $|\lambda| < 1$ is given by (3.17).

The proof is again by construction. By using the equation characterizing λ in (3.17), one can simplify equations (3.13) and (3.14) to obtain (4.5). As a consistency check, one can verify that $\pi_i(\lambda) = 0$.

The first term on the right-hand side of (4.5) tells us how prices respond to observable shocks to fundamentals. The second term then exhibits the additional dynamics induced when the shocks to fundamentals are unobservable, and traders must 'forecast the forecasts of others'. That is, the second term captures in a clear and precise way the higher-order belief dynamics associated with a heterogeneous beliefs equilibrium. By canceling the common root at $z = \beta$, it is clear that higher-order beliefs exhibit AR(1) dynamics, with a persistence given by λ. Interestingly, Woodford (2003) obtains a qualitatively similar result in a quite different setup.[16]

It is clear from (4.5) that higher-order beliefs generate additional price volatility. One manifestation of this is the following,

Corollary 4.5. *Heterogeneous beliefs amplify the initial response of asset prices to innovations in fundamentals.*

Proof. Evaluate (4.5) at $z = 0$. This yields

$$\pi_i(0) = \pi_i^s(0) - \beta a_i(\lambda) \cdot \frac{1-\lambda^2}{1-\lambda\beta} < \pi_i^s(0)$$

which verifies the result since responses to supply shocks are negative. □

Before turning to empirical implications, it is worthwhile working through an explicit numerical example. If nothing else, this will at least verify that the various assumptions imposed can be satisfied with reasonable specifications of the fundamentals. To illustrate the heterogeneous beliefs dynamics, we plot asset price impulse response functions. The orthogonality of the two fundamentals components allows us to proceed on a shock-by-shock basis. Without loss of generality, we consider the case of ε_{1t} shocks.

When solving for a heterogeneous beliefs equilibrium it is easier to work backwards from a pre-specified λ to a supporting fundamentals process than it is to start with fundamentals, and then check whether they are consistent with the existence of a heterogeneous beliefs equilibrium. Therefore, let $\lambda = 0.5$, and assume that $a_1(L)$ takes the form $(L - \phi_1)/(1 - \gamma_1 L)$, with $|\phi_1| < 1$ and $|\gamma_1| < 1$. Hence, $a_1(L)$ is noninvertible. Since we are confining our attention to ε_{1t} shocks, we don't need to take a stand on a precise specification of $a_2(L)$, other than assume it satisfies Assumption 2. At a minimum, this means its noninvertible root cannot equal ϕ_1. To be specific, we assume $\phi_1 = 0.83$ and $\beta = 0.90$. Plugging these into (3.17), one can readily verify that γ_1 must equal 0.476. Finally, given these values for $(\lambda, \phi_1, \gamma_1, \beta)$, we can use (4.5) to generate and compare the impulse response functions for the heterogeneous and homogeneous beliefs equilibria. We can also plot their differences, which are the higher-order belief dynamics associated with the heterogeneous beliefs equilibrium.

The following plots illustrate the asset price response to a one-unit shock in ε_{1t}. To make the results more comparable to standard asset pricing models, where dividends

[16]When aggregate fundamentals are observed (and so there are 3 trader types), the higher-order beliefs term in (4.5) is slightly altered, i.e., $a(\lambda)$ becomes $2a(\lambda)$.

are the fundamentals, we've multiplied the responses by minus one, so that prices *increase* following an innovation.

Figure 1: Impulse Response and Higher-Order Belief Dynamics

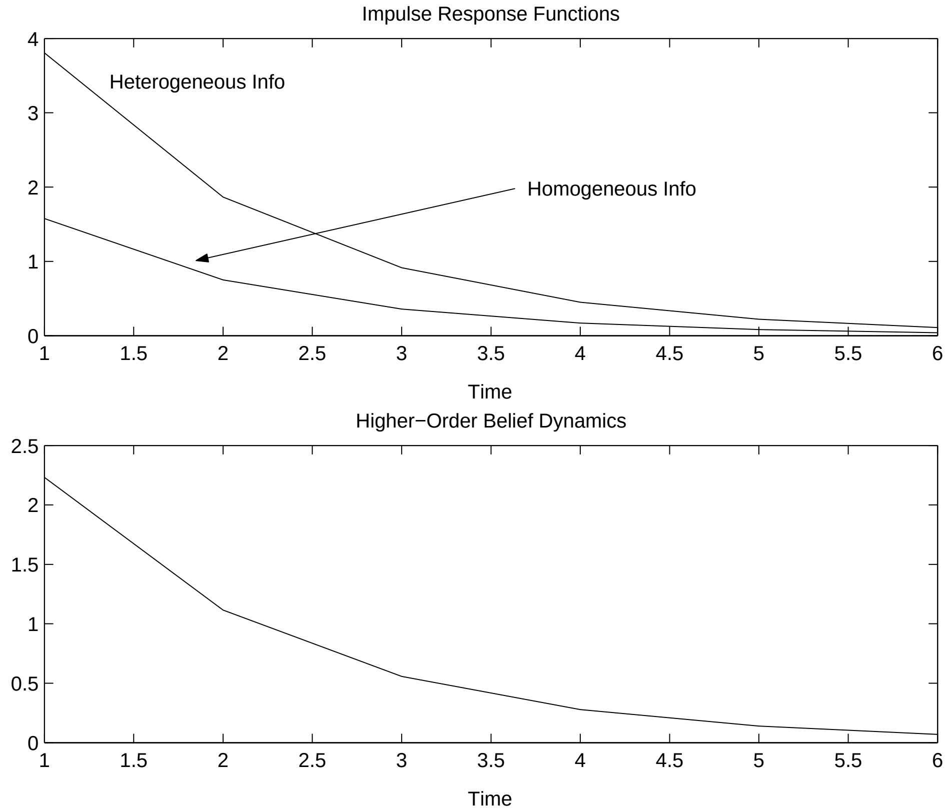

These plots clearly reveal the additional volatility and persistence induced by heterogeneous information and higher-order belief dynamics. Notice that the initial price response is more than twice as large in the heterogeneous beliefs equilibrium. In addition, the effects are persistent.

5. Empirical Implications

This section addresses the following question - Suppose the world is described by a heterogeneous expectations equilibrium, but an econometrician who, unlike the agents, observes fundamentals and interprets the data as if it were generated from a homogeneous expectations equilibrium. What kind of inferential errors could result? We focus on two empirical results that have been common in the asset pricing literature: (1) violations of variance bounds, and (2) rejections of cross-equation restrictions.[17]

[17]It may seem odd to assume the econometrician observes fundamentals but traders do not. However, as noted in section 3, by adding more traders, the unobserved fundamentals case becomes

5.1. **Variance Bounds.** Figure 1 and Corollary 4.5 suggest that higher-order belief dynamics will generally make asset prices appear to be 'too volatile' relative to their fundamentals. One of the main contributions of this paper is the ability to *quantify* the degree of excess volatility associated with higher-order belief dynamics. One way of doing this is to show that heterogeneous beliefs equilibria can violate standard variance bounds inequalities. Violations of these bounds are a robust empirical finding.

Variance bounds are based on the idea that observed asset prices should be less volatile than their perfect foresight counterparts (i.e., the subsequent realization of discounted future fundamentals). Since prices represent *expectations* of discounted future fundamentals, it makes sense that they should be smoother than the realizations of discounted future fundamentals. To show that heterogeneous beliefs equilibria can violate variance bounds, it therefore suffices to show that the variance of observed prices can exceed the variance of perfect foresight prices. The following proposition shows that this is indeed possible if λ (the persistence of higher-order belief dynamics) is sufficiently close to β.

Proposition 5.1. *If fundamentals are ARMA(1,1) (ie., $a_i(L) = (1-\phi_i L)/(1-\gamma_i L)$), then asset prices violate the standard variance bound whenever λ is sufficiently close to β and ϕ_i and γ_i are sufficiently small.*

Proof. Given orthogonality, it is sufficient to consider only one of the two components. Without loss of generality, we focus on the variance associated with ε_{1t}. First note that the z-transform of the perfect foresight price associated with this component is given by:

$$\pi_1^{pf}(z) = -\beta(1-\beta z^{-1})^{-1} a_1(z)$$

Using Parseval's formula, we can evaluate its variance as

$$\begin{aligned}
\mathrm{var}(p_1^{pf}) &= \frac{1}{2\pi i}\oint \pi_1^{pf}(z)\pi_1^{pf}(z^{-1})\frac{dz}{z} \\
&= \frac{\beta^2}{2\pi i}\oint \frac{(1-\phi_1 z)(1-\phi_1 z^{-1})}{(1-\gamma_1 z)(1-\gamma_1 z^{-1})(1-\beta z)(1-\beta z^{-1})}\frac{dz}{z} \\
&= \frac{\beta^2}{(\beta-\gamma_1)(1-\gamma_1\beta)}\left[\frac{(\beta-\phi_1)(1-\phi_1\beta)}{1-\beta^2} - \frac{(\gamma_1-\phi_1)(1-\phi_1\gamma_1)}{1-\gamma_1^2}\right] \\
&= \beta^2\frac{(1-\phi_1\gamma_1)(1-\phi_1\beta)+(\phi_1-\beta)(\phi_1-\gamma_1)}{(1-\gamma_1\beta)(1-\beta^2)(1-\gamma_1^2)}
\end{aligned}$$

---

essentially equivalent to the observed fundamentals case. From the perspective of an econometrician who does not observe any of the ε_{it}'s, and who only observes market data (f_t, p_t), the observed fundamentals case can be made to mimic the unobserved fundamentals representation by appropriate choice of the of the $a_i(L)$ functions (which produces a convenient solution of a spectral factorization). Section 6 provides an example.

Next, letting $\kappa(z)$ denote the z-transform of the higher-order belief dynamics (i.e., the second term on the right-hand side of (4.5)), we can write the variance of observed asset prices as:

$$\begin{aligned}
\mathrm{var}(p_1) &= \frac{1}{2\pi i}\oint \pi_1^s(z)\pi_1^s(z^{-1})\frac{dz}{z} + \frac{1}{2\pi i}\oint \kappa(z)\kappa(z^{-1})\frac{dz}{z} + \frac{2}{2\pi i}\oint \pi_1^s(z)\kappa(z^{-1})\frac{dz}{z} \\
&= \left(\frac{\beta(1-\phi_1\beta)}{1-\gamma_1\beta}\right)^2 \frac{(1-x\gamma_1)(1-x/\gamma_1)}{1-\gamma_1^2} + \left(\frac{\beta(1-\phi_1\lambda)}{(1-\lambda\beta)(1-\gamma_1\lambda)}\right)^2 (1-\lambda^2) \\
&\quad +2\frac{\beta^2(1-\lambda^2)(1-\phi_1\beta)(1-\phi_1\lambda)(1-x\lambda)}{(1-\gamma_1\beta)(1-\lambda\beta)(1-\gamma_1\lambda)^2}
\end{aligned}$$

where $x = \phi_1(1-\beta\gamma_1)/(1-\beta\phi_1)$, and where the second line again uses the residue theorem. Setting $\lambda = \beta$ and $\gamma_1 = \phi_1 = 0$ yields $\mathrm{var}(p_1^{pf}) = \beta^2/(1-\beta^2)$ and $\mathrm{var}(p_1) = \beta^2/(1-\beta^2) + 3\beta^2 > \mathrm{var}(p_1^{pf})$. The proof follows by continuity. □

This result has a very intuitive interpretation. The heterogenous beliefs equilibrium (4.5) contributes an additional component to the asset price which constitutes additional 'fundamentals.' Traders don't just care about their own expectations of future fundamentals, they also care about, and try to forecast, *other traders'* expectations about fundamentals. It is this additional component that leads to the violation of the variance bound. If these higher-order beliefs were incorporated into fundamentals, then asset prices would indeed satisfy the variance bound.

FIGURE 2: VARIANCE RATIO $\mathrm{VAR}(p_1)/\mathrm{VAR}(p_1^{pf})$

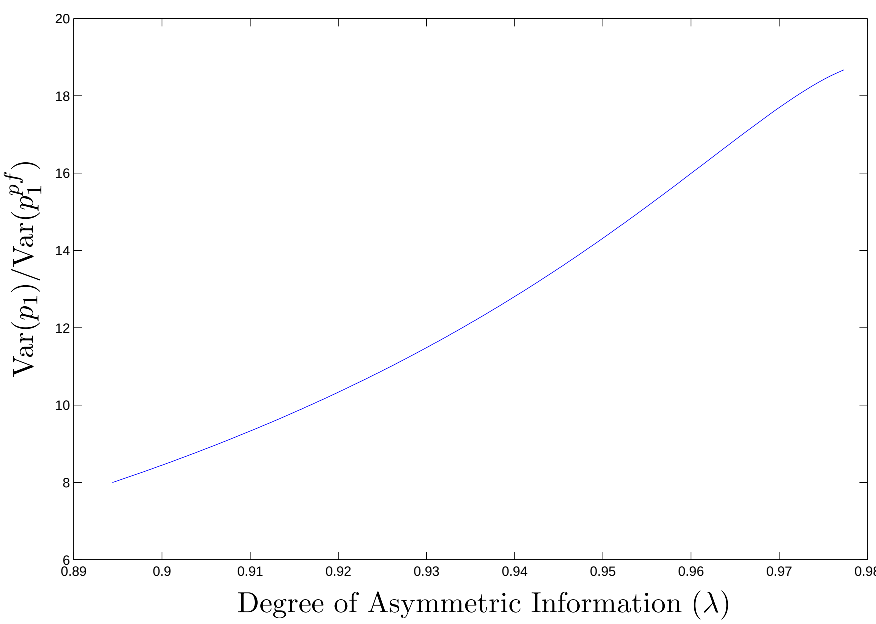

Figure 2 demonstrates violations of the variance bound for specific parameter values. Assuming fundamentals follow an ARMA(1,1) process with AR coefficient $\rho = 0.5$, Figure 2 plots the variance of the heterogenous beliefs price divided by the variance of the perfect foresight price by changing the MA component of fundamentals to satisfy Assumption 1 (i.e., to ensure a heterogeneous beliefs equilibrium exists). As the degree of asymmetric information (λ) approaches $\beta = 0.98$, the variance of the heterogeneous beliefs price process is 19 times as large as the perfect foresight price.

5.2. **Cross-Equation Restrictions.** The Rational Expectations revolution ushered in many methodological changes. One of the most important concerned the way econometricians identify their models. Instead of producing zero restrictions, the Rational Expectations Hypothesis imposes cross-equation restrictions. Specifically, parameters describing the laws of motion of exogenous forcing processes enter the laws of motion of endogenous decision processes. In fact, in a oft-repeated phrase, Sargent dubbed these restrictions the 'hallmark of Rational Expectations'. Hansen and Sargent (1991b) and Campbell and Shiller (1987) proposed useful procedures for testing these restrictions. It so happens that when these tests are applied to present value asset pricing models, they are almost uniformly rejected, and in a resounding way. There have been many responses to these rejections. Some interpret them as evidence in favor of stochastic discount factors. Others intepret them as evidence against the Rational Expectations Hypothesis. Looking on the bright side, Campbell and Shiller (1987) argue that a model can still be useful even when its cross-equation restrictions are statistically rejected. We offer yet another response. We show that rejections of cross-equation restrictions may reflect an informational misspecification, one that presumes a revealing equilibrium and homogeneous beliefs when in fact markets are characterized by heterogeneous beliefs.

To study the model's cross-equation restrictions, we need to derive its Wold representation. We can do this by following the steps outlined in Hansen and Sargent (1991b), which are based on the results in Rozanov (1967). The theoretical moving-average representation is given by,

$$\begin{bmatrix} f_t \\ p_t \end{bmatrix} = \begin{bmatrix} a_1(L) & a_2(L) \\ \pi_1(L) & \pi_2(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix} \tag{5.1}$$

where the pricing functions, $\pi_1(L)$ and $\pi_2(L)$, are given by equations (3.13) and (3.14). Assuming a heterogeneous beliefs equilibrium, and defining the vectors $\mathbf{x}_t = (f_t, p_t)'$ and $\boldsymbol{\epsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t})'$, write the MA representation as

$$\mathbf{x}_t = \mathcal{A}(L)\boldsymbol{\epsilon}_t.$$

By construction, $\mathcal{A}(L)$ does not have a one-sided inverse in positive powers of L due to the root inside the unit circle. Although not essential, we assume λ is the *only* root inside the unit circle, and impose the following assumption.

Assumption 5.2. *Write the determinant of $\mathcal{A}(z)$ as $|\mathcal{A}(z)| = (z-\lambda)\Delta(z)$, where $\Delta(z)$ is given by*

$$\Delta(z) = a_1(z)\tilde{\pi}_2(z) - a_2(z)\tilde{\pi}_1(z)$$

where $\tilde{\pi}_i(z) = \rho_i + zg_i(z)$. Then $\Delta(z)$ is a nonzero analytic function with roots outside the unit circle.

Since equal $a_i(z)$'s generate equal $g_i(z)$'s, this assumption requires the components to have different stochastic structures, which avoids stochastic singularities in the bivariate representation.

The lack of a one-sided inverse prevents traders from inferring the signals of other traders. The basic idea behind a Wold representation is to 'flip' this root outside the unit circle. This can be accomplished by post-multiplying $\mathcal{A}(L)$ by the product of two orthogonal matrices, and then pre-multiplying the shock vector by the transposes of these same matrices. This gives us,

$$\begin{aligned} \mathbf{x}_t &= [\mathcal{A}(L)WB(L)][B(L^{-1})'W'\boldsymbol{\epsilon}_t] \\ &= \mathcal{A}^*(L)\boldsymbol{\epsilon}_t^* \end{aligned} \tag{5.2}$$

where W is a scalar orthogonal matrix and $B(L)$ is a Blaschke matrix that flips the root from $z=\lambda$ to $z=\lambda^{-1}$. These matrices are given by

$$B(z) = \begin{bmatrix} \frac{1-\lambda z}{z-\lambda} & 0 \\ 0 & 1 \end{bmatrix} \qquad W = \frac{1}{\sqrt{1+\eta^2}} \begin{bmatrix} -\eta & 1 \\ 1 & \eta \end{bmatrix} \tag{5.3}$$

where $\eta = a_2(\lambda)/a_1(\lambda)$. Notice that $B(z)B(z^{-1})' = I$ on $|z| = 1$ and $WW' = I$. By construction, $\mathcal{A}^*(L)$ is invertible, so the observable VAR representation is $\mathcal{A}^*(L)^{-1}\mathbf{x}_t = \boldsymbol{\epsilon}_t^*$. A key point here is that the residuals, $\boldsymbol{\epsilon}_t^*$, are *not* the innovations to agents' information sets, $\boldsymbol{\epsilon}_t$. Instead, what is estimated are the linear combinations defined by $B(L^{-1})'W'\boldsymbol{\epsilon}_t$. These encode the model's higher-order belief dynamics. Although these linear combinations are mutually and serially uncorrelated by construction, they span a strictly smaller information set. Hence, the variance of $\boldsymbol{\epsilon}_t$ is smaller than the variance of $\boldsymbol{\epsilon}_t^*$.

Performing the matrix multiplication in (5.2) delivers the following Wold representation:

$$\begin{bmatrix} f_t \\ p_t \end{bmatrix} = \begin{bmatrix} \left(\frac{1-\lambda L}{L-\lambda}\right)[w_{11}a_1(L) + w_{21}a_2(L)] & w_{12}a_1(L) + w_{22}a_2(L) \\ \left(\frac{1-\lambda L}{L-\lambda}\right)[w_{11}K(a_1(L)) + w_{21}K(a_2(L))] & w_{12}K(a_1(L)) + w_{22}K(a_2(L)) \end{bmatrix} \begin{bmatrix} \varepsilon_{1t}^* \\ \varepsilon_{2t}^* \end{bmatrix} \tag{5.4}$$

where w_{ij} are the elements of the W matrix in (5.3) and $K(\cdot)$ defines the Rational Expectations pricing operator given by Proposition 3.6, i.e., $\pi_i(z) = K(a_i(z))$. We can use this to state two results, one pertaining to the case where the econometrician is alert to the possible existence of heterogeneous beliefs, and one pertaining to the case where the econometrician ignores the potential for heterogeneous beliefs.

As noted by Hansen and Sargent (1991a), Rational Expectations can be interpreted as placing restrictions across the rows of a model's moving average representation. In our case this means elements of the second row are exact functions of the corresponding elements of the first row. From the results in Hansen and Sargent (1991b), it is not too surprising that these restrictions continue to apply even in models with heterogeneous beliefs, *as long as the econometrician employs the correct pricing functions.*

Proposition 5.3. *Standard cross-equation restriction tests are valid even when the model features heterogeneous beliefs, as long as the econometrician is aware of this possibility, and uses the correct pricing functions.*

Proof. The proof follows directly from the fact that K is a linear operator. That is,

$$K(w_{ij}a_1(L) + w_{mn}a_2(L)) = w_{ij}K(a_1(L)) + w_{mn}K(a_2(L))$$

where w_{ij} and w_{mn} are arbitrary scalars. Hence, the bottom row of (5.4) is an exact function of the first. (The Blaschke factors in the first column can be folded into the definitions of the $a_i(L)$ polynomials). □

This is good news in the sense that it suggests standard testing procedures can be employed when evaluating models with heterogeneous beliefs.[18] The bad news is that ignoring the presence of heterogeneous beliefs can produce misleading results.

Proposition 5.4. *Standard cross-equation restriction tests, which falsely presume a common information set, can produce spurious rejections.*

Proof. From the results in section 4, we can decompose the heterogeneous beliefs pricing operator, K, into a traditional symmetric pricing operator, K^s, and a higher-order beliefs operator, K^h, so that $K = K^s + K^h$. Neglecting heterogeneous beliefs amounts to dropping the K^h component of the pricing operator. Evidently, if there are heterogeneous beliefs, so that $\pi_i(L) = K(a_i(L))$, then cross-equation restriction tests based on the false assumption of homogeneous beliefs can produce strong rejections when $K^h(a_i(L))$ is 'big' (in the operator sense). □

6. An Application to the Foreign Exchange Market

One of the intriguing findings of empirical asset pricing research is that similar results are obtained in all asset markets, whether it be stocks, bonds, foreign exchange, or real estate. This suggests that something basic, or fundamental, is missing from standard asset pricing models. We've argued that higher-order beliefs can *in principle* provide this missing element. To make this argument persuasive, however, it is important to go beyond the hypothetical, and show that this story holds water

[18] However, note that without modification the clever VAR testing strategy of Campbell and Shiller (1987) will *not* be valid, since it relies on the validity of the law of iterated expectations, which does not apply in models featuring higher order beliefs.

quantitatively, for empirically plausible specifications of fundamentals. To do this we need to narrow our focus to a particular asset market.

Foreign exchange has been an especially challenging market to understand, and so in this section we use our model to address some puzzling features of exchange rates. Engel and West (2005) provide a recent discussion of the empirical shortcomings of linear present value exchange rate models. They focus on the fact that most exchange rates are well approximated by random walks, even though monetary fundamentals are not. They show that this puzzle can be explained if fundamentals have a unit root component, and the model's discount rate is close to unity. They also argue that once allowance is made for the possibility of 'missing fundamentals', the model can also account reasonably well for exchange rate volatility and Granger causality findings. This second result is particularly interesting from the perspective of our work, since our model suggests that higher-order belief dynamics provide a natural candidate for these missing fundamentals. It turns out that higher-order beliefs can also generate near-random walk exchange rate behavior, even when fundamentals are stationary.

Our starting point is the following Uncovered Interest Parity condition, modified to allow for heterogeneous beliefs:

$$\begin{aligned} s_t &= \int_0^1 E_t^i s_{t+1} di - (i_t - i_t^*) \\ &\equiv \int_0^1 E_t^i s_{t+1} di - f_t \end{aligned} \tag{6.1}$$

The standard monetary model then uses money demand equations and a PPP condition to substitute out the interest differential in terms of relative money supplies and income. Although we could do this as well, we're going to instead regard the interest differential as the exogenous fundamental. This is empirically plausible, given the fact that Central Banks target the interest rate rather than the money supply, although the exogeneity assumption rules out feedback from the exchange rate to the interest rate. This is an attractive assumption also because it maps readily into our previous framework, and as we'll see, generates straightforward predictions for Uncovered Interest Parity regressions.

Since it is implausible to assume that interest rates are unobserved, from our previous results we know that for a heterogeneous beliefs equilibrium to exist, there must be at least three trader types. Thus, we postulate the following unobserved components specification for f_t,

$$f_t = a_1(L)\varepsilon_{1t} + a_2(L)\varepsilon_{2t} + a_3(L)\varepsilon_{3t}$$

where it is assumed that a type-i trader observes $(s_t, f_t, \varepsilon_{it})$. Given that $\beta = 1$ here, the existence condition in Assumption 3.9 takes the following form

$$\lambda = 2 + a_i(1)/a_i(\lambda) \qquad \text{i=1,2,3} \tag{6.2}$$

To make things as simple as possible we assume that the noninvertible root, λ, coincides with a common autoregressive root in the fundamentals process. We further suppose that each component is ARMA(1,2), so that for example we have:

$$a_1(L) = \frac{1 + a_1 L + a_2 L^2}{1 - \lambda L}$$

To avoid singularities, we need to assume the MA coefficients differ across components. However, given orthogonality, we can just focus on one of the components. Although it is not essential to assume a common autoregressive root, which also coincides with the noninvertiable root of the price process, it does simplify the analysis, and highlights the role of higher-order beliefs in generating random walk exchange rate behavior.

The existence condition in (6.2) implies the following restriction between the MA coefficients, a_1 and a_2

$$a_1 = -\left[\frac{3 + a_2[1 + \lambda^2(2-\lambda)] - \lambda}{1 + \lambda(2-\lambda)}\right] \tag{6.3}$$

Given this, we obtain the following heterogeneous beliefs pricing function

$$\pi(L) = -\left[\frac{(1 + a_1 + a_2)/(1-\lambda) + 2a(\lambda)\cdot(1+\lambda) + (a_1 + a_2)L + a_2 L^2}{1 - \lambda L}\right] \tag{6.4}$$

Now, the empirical challenge is to specify a fundamentals process that is persistent (since observed interest differentials are persistent), but not as persistent as a random walk, and at the same time produce a near-random walk exchange rate process. To do this, we let the higher-order beliefs parameter be close to unity, i.e., $\lambda = .99$, and then try setting $a_2 = -0.3$. The existence condition (6.3) then implies $a_1 = -.707$. Given these parameter values, we get the following spectral densities for the exchange rate and fundamentals:

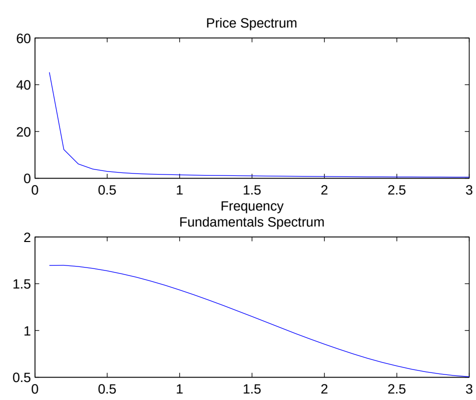

It is apparent that the exchange rate is far more persistent than the interest differential. This becomes clearer if plot the spectral densities of their first-differences:

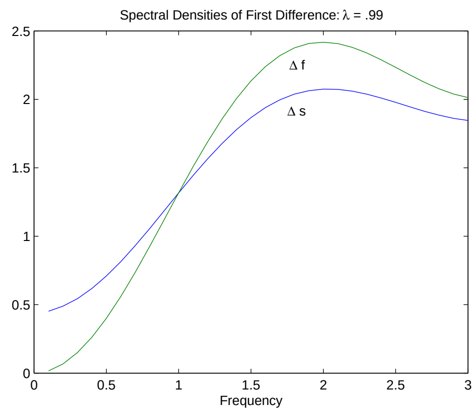

Note that the predicted exchange rate is not exactly a random walk, since its first-differenced spectrum is not flat, but at the same time, it is clear that higher-order beliefs have the effect of allocating spectral power to the lower frequency ranges. In fact, we can see just how important higher-order beliefs are by exploiting the decomposition in equation (4.5), which allows us to plot separately the higher-order beliefs spectrum and the overall price spectrum:

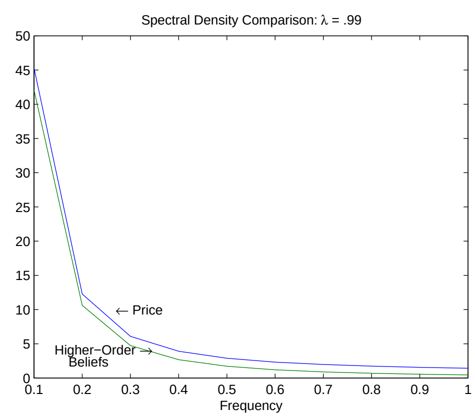

Clearly, the exchange rate is being driven almost entirely by higher-order belief dynamics. We think this result provides a nice complement to the near-random walk result of Engel and West (2005).

Besides the random walk nature of exchange rates, another troublesome feature of the foreign exchange market is that fact that Uncovered Interest Parity does not hold, and attempts to link these deviations to observable risk premia have not been very successful. In fact, not only does Uncovered Interest Parity not hold, it doesn't even get the sign right! That is, regressions of ex post exchange rate changes (defined as the price of foreign currency) on domestic minus foreign interest rate differentials produce *negative* coefficient estimates. Interestingly, our model can explain deviations from Uncovered Interest Parity. To see this, we must derive the model's Wold representation for $(s_t, i_t - i_t^*)$. Since the econometrician does not observe *any* of the ε_{it}'s, this involves *two* steps. First, we must perform a spectral factorization to reduce the 3-dimensional equilibrium representation of section 3.2 into a 2-dimensional representation. Second, we must then convert this to a Wold representation by flipping roots, as described in section 5.2. Note that if we assume $a_1(L) \equiv a_2(L)$, then the factorization step becomes trivial. We just need to define a new shock, $\varepsilon_{1t} + \varepsilon_{2t}$. From the perspective of the econometrician, the system remains nonsingular, since he only observes $(s_t, i_t - i_t^*)$. The advantage from doing this is that we can now directly apply the results from section 5.2 to obtain the Wold representation. In particular, letting $\mathbf{x}_t = (i_t - i_t^*, s_t)$, we get (5.2) as the Wold representation, with $B(z)$ and W defined exactly as in (5.3).

To derive the model's implications for UIP regressions, we just need to calculate the following projection

$$P[\mathbf{x}_{t+1}|H_{\mathbf{x}}(t)] = [\mathcal{A}^*(L)L^{-1}]_+ \mathcal{A}^*(L)^{-1}\mathbf{x}_t \tag{6.5}$$

with $\mathcal{A}^*(L)$ given by (5.4). UIP imposes the restriction that the bottom row of this matrix is equal to $(1, 1)$. Exploiting the linearity of the pricing operator, we can write $\mathcal{A}^*(L)$ as

$$\mathcal{A}^*(L) = \begin{bmatrix} \alpha_1(L) & \alpha_2(L) \\ \\ K(\alpha_1(L)) & K(\alpha_2(L)) \end{bmatrix}$$

where $K(\cdot)$ defines the Rational Expectations pricing operator as before, and the $\alpha_i(L)$ functions are the linear combinations of the $a_i(L)$ functions defined in (5.4). Note that given our previous assumptions, each of these is ARMA(1,2), with identical autoregressive roots. After a little algebra, we then get the following two restrictions:

$$\begin{aligned} \alpha_1 K(\alpha_2) - \alpha_2 K(\alpha_1) &= K(\alpha_2)[K(\alpha_1) - K(\alpha_1)_0]L^{-1} - K(\alpha_1)[K(\alpha_2) - K(\alpha_2)_0]L^{-1} \\ \alpha_1 K(\alpha_2) - \alpha_2 K(\alpha_1) &= -\alpha_2[K(\alpha_1) - K(\alpha_1)_0]L^{-1} + \alpha_1[K(\alpha_2) - K(\alpha_2)_0]L^{-1} \end{aligned}$$

One can readily verify that with homogeneous beliefs these conditions are satisfied, since from the (symmetric) Rational Expectations fixed point condition we have $[K(\alpha) - K_0]L^{-1} = K(\alpha) + \alpha$. However, by comparing (3.6) to (3.10) and (3.11), we can see that this is no longer true when higher-order beliefs are present, due to the additional $\rho(1 - \lambda^2)/(1 - \lambda L)$ functions. Remember, with private information, agents try to forecast *other agents'* forecasts, which play the role of unobserved fundamentals. Since other agents' forecasts are never directly revealed ex post, there is no guarantee that errors in forecasting them are uncorrelated with fundamentals. The implications for UIP are not obvious. What is clear is that standard UIP regressions will be misspecified, since the bottom-row elements are in general not simply scalars, but lag polynomials. This produces a complex omitted variable bias. Intuitively though, since risk is not an issue here, to explain a negative slope coefficient traders must on average expect an *appreciation* following an interest rate increase. This could occur if the omitted lag polynomial produces a hump-shaped impulse response function.

The following two figures show what happens when you regress Δs_t on lagged interest differentials and lagged exchange rates when fundamentals are specified as above. The first figure is for the case $\lambda = .95$, and the second is for the case $\lambda = .50$.

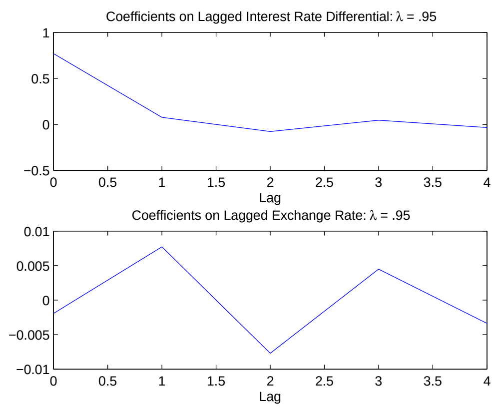

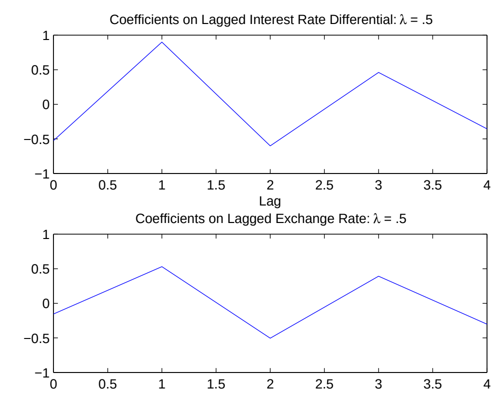

One can see that for $\lambda = .95$ the model delivers a downward bias in the UIP regression coefficient, but it is not negative, as in the data. However, if $\lambda = .50$, then the regression coefficient *is* negative, but unfortunately, the other coefficients become counterfactually large. It remains to be seen whether empirically realistic specifications of fundamentals can reconcile observed UIP regressions.

7. Conclusion

For more than twenty years now, economists have been rejecting linear present value models of asset prices. These rejections have been interpreted as evidence in favor of time-varying risk premiums. Unfortunately, linking these risk premiums to observable data has proven to be quite challenging. Promising approaches for meeting the challenge involve introducing incomplete markets and agent heterogeneity into the models.

This paper has suggested that a different sort of heterogeneity, an *informational heterogeneity*, offers an equally promising route toward reconciling asset prices with observed fundamentals. Unfortunately, heterogeneous information does not automatically translate into heterogeneous beliefs, and it is only the latter that generates the 'excess volatility' that is so commonly seen in the data. The hard work in the analysis, therefore, is deriving the conditions that prevent market data from fully revealing the private information of agents in dynamic settings. We have argued that frequency-domain methods possess distinct advantages over time-domain methods in this regard. The key to keeping information from leaking out through observed asset prices is to ensure that the mappings between the two are 'noninvertible'. These noninvertibility conditions are easy to derive and manipulate in the frequency domain.

Our results demonstrate how informational heterogeneity can in principle explain well-known empirical anomalies, such as excess volatility and rejections of cross-equation restrictions. Ever since Townsend (1983) and Singleton (1987), (or in fact, ever since Keynes!) economists have suspected that higher-order beliefs could be responsible for the apparent excess volatility in financial markets. Our results at last confirm these suspicions. Although we believe we have made substantial progress, there are still many avenues open for future research. Two seem particularly important. First, our existence conditions place restrictions on fundamentals. It remains to be seen, however, whether heterogeneous beliefs equilibria can be supported with empirically plausible specifications for fundamentals. We are optimistic that they can, since the restrictions are fairly generic, but verification of this conjecture remains the subject of future research. Second, the analysis here rests heavily on linearity. However, most macroeconomic models feature nonlinearities of one form or another. It is not at all clear whether standard linearization methods are applicable in models featuring higher-order beliefs. Resolving this issue will be important for future applications.

8. Appendix: Frequency Domain Techniques

This appendix offers a brief introduction to the frequency domain techniques used to solve the model. In lieu of matching the infinite sequence associated with the fixed-point (3.15), we employ the following theorem and solve for a functional fixed point.

Theorem (Riesz-Fischer): *Let $\{c_n\}$ be a square summable sequence of complex numbers (i.e., $\sum_{n=-\infty}^{\infty} |c_n|^2 < \infty$). Then there exists a complex-valued function, $g(\omega)$, defined for $\omega \in [-\pi, \pi]$, such that*

$$g(\omega) = \sum_{j=-\infty}^{\infty} c_j e^{-i\omega j} \tag{8.1}$$

where convergence is in the mean-square sense

$$\lim_{n\to\infty} \int_{-\pi}^{\pi} \left| \sum_{j=-n}^{n} c_j e^{-i\omega j} - g(\omega) \right|^2 d\omega = 0$$

and $g(\omega)$ is square (Lebesgue) integrable

$$\int_{-\pi}^{\pi} |g(\omega)|^2 d\omega < \infty$$

Conversely, given a square integrable $g(\omega)$ there exists a square summable sequence such that

$$c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} g(\omega) e^{i\omega k} d\omega \tag{8.2}$$

The Fourier transform pair in (8.1) and (8.2) defines an isometric isomorphism (i.e., a one-to-one onto transformation that preserves distance and linear structure) between the space of square summable sequences, $\ell^2(-\infty, \infty)$, and the space of square integrable functions, $L^2[-\pi, \pi]$. The sequence space, ℓ^2, is referred to as the 'time domain' and the function space, L^2, is referred to as the 'frequency domain'. The equivalence between these two spaces allows us to work in whichever is most convenient. A basic premise of this paper is that in models featuring heterogeneous beliefs, the frequency domain is analytically more convenient.

In the context of linear prediction and signal extraction, it is useful to work with a version of Riesz-Fischer theorem that is generalized in one sense and specialized in another. In particular, it is possible to show, via Poisson's integral formula, that the statement of the theorem applies not only to functions defined on an interval (the boundary of the unit circle), but to analytic functions defined within the entire unit circle of the complex plane. However, when extending the theorem in this way we exclude functions with Fourier coefficients that are nonzero for negative k. That is, we limit ourselves to functions where $c_{-k} = 0$ in equations (8.1) and (8.2). This turns out to be useful, since it is precisely these functions that represent the 'past' in the time domain. A space of analytic functions in the unit disk defined in this way is called a Hardy space, with an inner product defined by the contour integral,

$$(g_1, g_2) = \frac{1}{2\pi i} \oint g_1(z)\overline{g_2(z)}\, \frac{dz}{z}.$$

Rather than postulate a functional form and match coefficients, we solve for a single analytic function which represents, in the sense of the Riesz-Fischer theorem, this unknown pricing function. The approach is still ‘guess and verify’, but it takes place in a function space, and it works because the Riesz-Fischer theorem tells us that two stochastic processes are ‘equal’ if and only if their z-transforms are identical as analytic functions inside the (open) unit disk. The real advantage of this approach stems from the ease with which it handles noninvertibility (i.e., nonrevealing information) issues. Invertibility hinges on the absence of zeroes inside the unit circle of the z-transform of the observed market data. By characterizing these zeroes, we characterize the information revealing properties of the equilibrium.

References

Allen, F., S. Morris, and H. S. Shin (2006): "Beauty Contests and Iterated Expectations in Asset Markets," *Review of Financial Studies*, 19, 719–752.

Angeletos, G.-M., and I. Werning (2006): "Crises and Prices: Information Aggregation, Multiplicity and Volatility," *American Economic Review*, 96, 1720–36.

Atkeson, A. (2000): "Comment on 'Rethinking Multiple Equilbria in Macroeconomic Modeling'," in *NBER Macroeconomics Annual 2000*, ed. by B. S. Bernanke, and K. Rogoff. MIT Press.

Bacchetta, P., and E. van Wincoop (2006): "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?," *American Economic Review*, 96(3), 552–576.

Bernhardt, D., P. Seiler, and B. Taub (2005): "Speculative Dynamics," Working Paper, University of Illinois.

Brunnermeier, M. K. (2001): *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis, and Herding*. Oxford University Press, Oxford.

Campbell, J. Y., and R. J. Shiller (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

Cochrane, J. H. (2001): *Asset Pricing*. Princeton University Press, Princeton.

Constantinides, G. M., and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.

Engel, C., and K. D. West (2005): "Exchange Rates and Fundamentals," *Journal of Political Economy*, 113(3), 485–517.

Futia, C. A. (1981): "Rational Expectations in Stationary Linear Models," *Econometrica*, 49(1), 171–192.

Grossman, S. J., and J. E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393–408.

Hamilton, J. D., and C. H. Whiteman (1985): "The Observable Implications of Self-Fulfilling Expectations," *Journal of Monetary Economics*, 16, 353–373.

Hansen, L. P., and T. J. Sargent (1991a): "Exact Linear Rational Expectations Models: Specification and Estimation," in *Rational Expectations Econometrics*, ed. by L. P. Hansen, and T. J. Sargent, pp. 45–76. Westview Press.

——— (1991b): "Two Difficulties in Interpreting Vector Autoregressions," in *Rational Expectations Econometrics*, ed. by L. P. Hansen, and T. J. Sargent, pp. 77–119. Westview Press.

Hoffman, K. (1962): *Banach Spaces of Analytic Functions*. Prentice-Hall, Englewood Cliffs, New Jersey.

Kasa, K. (2000): "Forecasting the Forecasts of Others in the Frequency Domain," *Review of Economic Dynamics*, 3, 726–756.

Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

Makarov, I., and O. Rytchkov (2006): "Forecasting the Forecasts of Others: Implications for Asset Pricing," Working Paper, MIT.

Morris, S., and H. S. Shin (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88, 587–597.

——— (2000): "Rethinking Multiple Equilibria in Macroeconomics," *NBER Macro Annual*, 15, 20–40.

——— (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics: Theory and Applications*, ed. by S. T. M. Dewatripont, L.P. Hansen. Cambridge University Press, Cambridge.

Nimark, K. P. (2007): "Dynamic Higher Order Expectations," Working Paper, Research Department, Reserve Bank of Australia.

Pearlman, J. G., and T. J. Sargent (2005): "Knowing the Forecasts of Others," *Review of Economic Dynamics*, 8(2), 480–497.

Radjavi, H., and P. Rosenthal (1973): *Invariant Subspaces*. Springer-Verlag, New York.

Rozanov, Y. A. (1967): *Stationary Random Processes*. Holden-Day, San Francisco.

Singleton, Kenneth, J. (1987): "Asset Prices in a Time Series Model with Disparately Informed, Competitive Traders," in *New Approaches to Monetary Economics*, ed. by W. Barnett, and K. Singleton. Cambridge University Press, Cambridge.

Townsend, R. M. (1983): "Forecasting the Forecasts of Others," *Journal of Political Economy*, 91, 546–588.

Walker, T. B. (2007): "How Equilibrium Prices Reveal Information in Time Series Models with Disparately Informed, Competitive Traders," forthcoming, *Journal of Economic Theory*.

Wang, J. (1993): "A Model of Asset Prices under Asymmetric Information," *Review of Economic Studies*, 60(2), 249–282.

Whiteman, C. H. (1989): "Equilibrium Volatility and Heterogeneity in a Simple Asset Pricing Model," Working Paper, University of Iowa.

Woodford, M. (2003): "Imperfect Common Knowledge and the Effects of Monetary Policy," in *Knowledge, Information, and Expectations in Modern Macroeconomics*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford. Princeton University Press, Princeton, N.J.