

A Theory of Asset Prices based on Heterogeneous Information *

Elias Albagli
USC Marshall

Christian Hellwig
Toulouse School of Economics

Aleh Tsyvinski
Yale University

December 14, 2011

Abstract

We propose a theory of asset prices that emphasizes heterogeneous information as the main element determining prices of different securities. With only minimal restrictions on security payoffs and trader preferences, noisy aggregation of heterogeneous information drives a systematic wedge between the impact of fundamentals on an asset price, and the corresponding impact on cash flow expectations. From an ex ante perspective, this information aggregation wedge leads to a systematic gap between an asset's expected price and its expected dividend, whose sign and magnitude depend on the asymmetry between upside and downside payoff risks, and on the importance of information heterogeneity. Moreover, when information frictions are sufficiently severe, the model is consistent with arbitrarily high levels of excess price variability as well as low return predictability. Importantly, these results do not rely on traders' risk aversion and thus offer an alternative theory of expected asset returns and price volatility. As applications of our theory, we first highlight how heterogeneous information leads to systematic departures from the Modigliani-Miller theorem and provide a new theory of debt versus equity. Second, in a dynamic extension we provide conditions under which price bubbles are sustainable.

*We thank Bruno Biais, John Geanakoplos, Narayana Kocherlakota, Stephen Morris, Guillaume Plantin, Jeremy Stein, Jean Tirole, Martin Weber, Xavier Vives, Eric Young, and audiences at IIES (Stockholm), Yale, UCLA, the 2nd French Macro-Finance Summer workshop (Sciences Po.), and ESSET Gerzensee for helpful comments. Hellwig gratefully acknowledges financial support from the European Research Council (starting grant agreement 263790). Tsyvinski is grateful to NSF for support and EIEF for hospitality.

1 Introduction

Dispersed investor information, and disagreement among investors about the expected cash-flows of different securities is a common feature of many, if not most financial markets. In this paper, we develop a parsimonious, flexible model of asset pricing in which heterogeneity of information and its aggregation in the market emerges as the core force determining asset prices and expected returns. Our model delivers novel insights and sharp predictions that link the asset’s predicted prices and returns to features of the market environment and the distribution of the underlying cash-flow risk. We further show that heterogeneous information provides a natural source of excess price volatility. Finally, our model can easily be adapted to address a variety of questions. Using our model, we reconsider two classical results in a heterogeneous information setting: the Modigliani-Miller Theorem, and the sustainability of bubbles in a dynamic environment.

Specifically, we consider an asset market along the lines of Grossman and Stiglitz (1980), Hellwig (1980), and Diamond and Verrecchia (1981).¹ An investor pool is divided into informed traders who have observed a noisy signal about the value of an underlying cash flow, and uninformed noise traders. The traders all submit orders to buy shares in the cash flow at the going price. The price serves as a noisy signal of the state, which traders use along with their private signals to form an update about the cash flows. Using the market structure first introduced in Hellwig, Mukherji and Tsyvinski (2006), we assume that traders are risk neutral but face limits on their asset positions. This enables us to derive a closed-form characterization for prices and expected dividends conditional on the price, with no restriction on cash flows other than monotonicity in the underlying fundamental shocks.²

In our model, the asset price is equal to the expectation of cash flows for a “marginal investor” who is just indifferent between investing and not investing in the asset. We compare the marginal investor’s posterior belief to the belief of an objective outsider, who uses the observation of the price to update beliefs about dividends, or equivalently an “econometrician” who uses a sample of price-dividend observations to estimate this relationship. Compared to the outsider, the marginal trader treats the information contained in the price as if he over-estimated its information content. His posterior expectations thus attach a higher weight to the market signal and a lower residual uncertainty to the fundamental than would be justified by its true information content. We label this discrepancy the “information aggregation wedge”.

¹See Brunnermeier (2001), Vives (2008), and Veldkamp (2011) for textbook discussions.

²With risk neutrality, our characterization of equilibrium demand is connected to models of competitive bidding in common value auctions (Milgrom, 1981), where traders have private information.

Despite its appearance, the information aggregation wedge is not the result of non-Bayesian updating or irrational trading decisions. Instead, it results from compositional shifts under investor heterogeneity: to maintain market-clearing, the identity of the marginal trader has to change with the observed price in a way that amplifies the impact of the price on the marginal trader's expectations. For example, consider either an increase in the informed traders' demand coming from a more favorable realization of the fundamentals (and hence their aggregate signal distribution), or an increase in the noise traders' demand. These shifts both result in a higher price and a higher expectation of future dividends, because of the information conveyed through the price. In addition, since the demand by informed traders has become larger (or the pool of available securities smaller), the marginal investor's private signal has to become more optimistic just to maintain market-clearing. This further increases the price, but not expected dividends, over and above the direct signal effect. The asset price thus appears to respond more to the market signal than would be justified on the basis of its true information content.

From an *ex ante* perspective, we characterize the average price and dividends in closed form as a function of the cash flow distribution and a parameter that summarizes the severity of the informational friction. This information friction parameter depends on the accuracy of informed traders' private signals, and the variance of noise trading shocks. Intuitively, the unconditional wedge is the expected value of a mean-preserving spread of the underlying distribution of the payoffs, i.e. from an *ex ante* perspective the market puts a higher weight on the tails than the objective distribution of the fundamental. Moreover, the unconditional wedge has increasing differences between the informational noise parameter, and the asymmetry between upside and downside risks, where the latter is defined as a partial order on payoff risks that compares the marginal gains and losses at fixed distances from the prior mean of the fundamental.

From this characterization, Theorem 1 then provides several general implications for expected returns. Regardless of the informational parameters, the unconditional wedge is zero when payoff risk is symmetric. The wedge is positive (meaning that the expected price exceeds expected dividends) for risks that are dominated by the upside, and negative for risks that are dominated by the downside. Moreover, in absolute value this wedge becomes more pronounced for more asymmetric payoff risks, or for a higher degree of information aggregation frictions. Our model thus offers sharp, novel predictions that link the occurrence, size and direction of price premia and discounts, both unconditionally and conditionally on the realization of shocks, to specific characteristics of the market and the underlying cash-flow risk.

Theorem 2 characterizes the variability of prices relative to expected and realized dividends. We

show that prices are always more variable than expected dividends. If the information aggregation wedge is sufficiently important, prices may even be more variable than realized dividends. In the limiting cases, the variability of prices exceeds that of realized dividend by any arbitrarily large factor. Moreover, the correlation between price and realized dividends may be arbitrarily close to zero. This stands in sharp contrast with the standard result in the asset pricing literature that price volatility coming from dividend expectations is bounded above by the volatility of realized dividends (as in West, 1988). Since dividend volatility in the data falls short in explaining variability of prices (LeRoy and Porter, 1981; Shiller, 1981), the consensus explanation stresses variation in stochastic discount rates (Campbell and Shiller, 1988; Cochrane, 1992). Our theory instead suggests that high price volatility could result from volatile market expectations about dividends in a fully rational environment despite low variability in observed dividends, as long as the informational frictions are strong enough.

We consider two applications of our theory. The first revisits the Modigliani-Miller Theorem, which establishes that under conditions of no arbitrage the total market value of any given cash flow is not influenced by how it is divided into separate securities. Absent distortions inside the firm, the optimal capital structure is indeterminate and disconnected from the firm's market valuation (Modigliani and Miller, 1958). Capital structure theories then focus mostly on trade-offs that affect the generation of cash flows inside the firm, such as agency costs, information frictions or tax distortions, assuming that the market value of the resulting cash flow is not affected by its split into different securities. Here instead we take the view that capital structure and firm value may also be influenced by heterogeneous information and financial market frictions.

We consider a seller who is splitting a given cash flow into two pieces which are sold to separate investor pools in two different markets, and suppose that at least one of the pieces is dominated either by upside or by downside risk. We show that the expected revenue of the seller is not affected by the split, if and only if the two markets are characterized by identical informational characteristics. However, when the investor pools differ, the seller can manipulate her expected revenue by selling downside risks in the market with smaller information aggregation frictions, and upside risks in the market with larger information aggregation frictions. The seller maximizes expected revenue by completely separating upside and downside risks, splitting the cash flow into a debt claim for the downside, and an equity claim for the upside, with a default point for debt at the prior median.

Second, we consider the sustainability of rational bubbles. A well-known result shows that the absence of arbitrage eliminates the possibility of persistent over-pricing of securities (Tirole, 1982;

Santos and Woodford, 1997). While the anticipation of higher future prices would, in principle, induce agents to increase the price bid in the current period, the combination of no arbitrage with transversality conditions (or backwards induction, in case of assets with finite horizons) rules out the possibility of any security trading at a price that exceeds the net present value of expected future cash flows.

We consider a simple, infinitely repeated version of our trading model with constant discounting, and give conditions under which a security may be permanently over- or under-priced, regardless of current market conditions. As usual, we can break down the current price, expected dividends and wedge into a component resulting from expectations about current cash-flows, and a component resulting from expected discounted future cash-flows and prices. The former inherits the same properties as the static conditional information aggregation wedge, while the latter inherits the properties of the unconditional wedge. If the cash-flow has a bounded downside risk and is dominated by the upside, and traders are sufficiently patient, then the positive expected future wedge more than offsets any negative current wedge. The asset then trades at a premium over its expected dividend value regardless of the current state realization. The flipside of these conditions shows that securities that have bounded upside and are dominated by the downside risk may be permanently underpriced.

We then generalize the characterization of the information aggregation wedge to almost completely general specifications of preferences and general distributional assumptions, relaxing in particular the core assumptions of risk-neutrality and on the bounds of trader positions. This result confirms that the wedge is a very general property of asset pricing models with noisy information aggregation, and not due to the specific structure of preferences that we imposed. The formal result we prove is the following: for any noisy rational expectations equilibrium that satisfies a pair of regularity conditions on equilibrium posterior beliefs and demand functions by informed traders, there exists a sufficient statistic such that the equilibrium price can be represented as a function only of that statistic. This is the key insight that allows us to show the existence of the information aggregation wedge even in this general model. In other words, the core features of the equilibrium characterization we work with throughout the paper turn out to be extremely general properties of price formation with noisy information aggregation. Another generalization we consider is to modify our benchmark model to show how the magnitude of the wedge is inversely related to the extent of arbitrage activity by risk-neutral, uninformed traders.

Our paper contributes to a large literature on noisy information aggregation in asset markets, including the papers cited above. Much of this literature works within a canonical preference struc-

ture of CARA utility and normally distributed signals and dividends. Remarkably, the information aggregation wedge appears to have received little attention in this literature, even though it is present in these canonical models, and, as we show, is the source of rich implications for prices, trading activities, and market volatility.³ By avoiding the restrictive functional form assumptions on cash-flow distributions, we are able to provide a characterization of this wedge for a general class of securities and draw implications that link average returns and return volatility to features of the cash-flow distribution and the importance of information frictions.

Within the CARA-normal framework, several authors have studied risk-aversion as a source of excess volatility. In a dynamic OLG economy with symmetric information, Spiegel (1998) shows that there are 2^n equilibria in an n -asset economy. In the high volatility equilibria, traders expect prices to vary greatly in response to noisy supply innovations in the next period, becoming reluctant to absorb current supply shock. Hence, supply shocks have large impact on current prices, sustaining a high volatility equilibrium. Watanabe (2008) extends this framework to an environment with heterogeneous information, providing additional results on the correlation properties between assets in each equilibria and the impact of information precision. In our model excess price volatility is neither coming from risk aversion as traders are risk neutral, nor from the dynamic OLG structure since our baseline model is static, but from the shift in identity of the trader pricing the asset which lies at the core of the information aggregation wedge.

Another influential literature emphasizes heterogeneous beliefs and short sales constraints as potential sources of bubbles, mis-pricing, and market anomalies (Harrison and Kreps, 1978; Allen, Morris and Postlewaite, 1993; Chen, Hong and Stein, 2002; Scheinkman and Xiong, 2003; Hong and Stein, 2007; Hong and Sraer, 2011). Mispricing is sustained by the option to resell an over-valued security to an even more optimistic buyer in the future. This option becomes valuable in the presence of (one-sided) short-sales constraints, and implies a channel for over-valuation. Heterogeneity in prior beliefs is taken as exogenous, and with the exception of Allen, Morris and Postlewaite (1993), traders do not update from the observation of prices. We touch on similar themes, but stay within the REE tradition in which traders' beliefs result from exogenous signals, and information aggregation through prices imposes tight restrictions on the heterogeneity in beliefs. Furthermore, our limits to arbitrage are not explicitly asymmetric, give rise to over- as well as under-valuation results, and our market environment is static, so the resale option doesn't play an important role (except in the dynamic application to bubbles). The mechanism that gives rise to

³The only written statement of this observation that we have found appears in Vives (2008), where it is only mentioned in passing.

mis-pricing and bubbles is therefore quite different.⁴

The literature on “over-confidence” explores how asset prices and financial markets are influenced by the degree to which investors over-estimate the accuracy of their own information (e.g. Odean, 1998, and Daniel, Hirshleifer and Subrahmanyam, 1998). When viewed from the perspective of a representative investor, the market price that emerges in our model is perfectly consistent with these same over-confidence biases, yet all investors are fully rational and not mistaken about the quality of their signals. What may look like an over-confidence bias in the aggregate can thus be accounted for by heterogeneity and aggregation from the micro level.

More generally, any theory of mispricing must rely on some source of noise affecting the market, coupled with some limits to the traders’ ability to exploit the resulting arbitrage opportunity (see Gromb and Vayanos, 2010, for an overview and numerous references). In our model, the combination of noise trading and limits to arbitrage with heterogeneous information leads not just to random errors in the price, but to systematic, predictable departures of the price from the asset’s fundamental value. The exact nature of our limits to arbitrage assumptions (embedded in the position limits and the noise trading formulation) is not central for our results, but guarantees the tractability of the updating, with virtually no assumptions imposed on cash-flows.

In section 2, we describe our model and provide the equilibrium characterization of asset prices. In section 3, we define the information aggregation wedge and discuss at length the resulting prediction for conditional and unconditional asset returns. Section 4 uses the insight offered by these two results to revisit the Modigliani-Miller theorem, and the existence of bubbles in the dynamic version of the model. Section 5 concludes the analysis with the robustness discussion.

2 Model

2.1 Agents, assets, information structure and financial market

The market is set as a Bayesian trading game with a unit measure of risk-neutral, informed traders, a stochastic measure of uninformed “noise traders”, and a ‘Walrasian auctioneer’. There is a risky asset whose supply is normalized to a unit measure, and whose dividend is a strictly increasing and twice continuously differentiable function $\pi(\cdot)$ of a stochastic “fundamental” θ .

⁴For example, the difference between our work and theories of bubbles based on short-sales constraints becomes clear if one considers the case of debt instruments, as in Hong and Sraer (2011). Whereas in their model, short-sales constraints lead to over-valuation of debt, but with less volatility and trading volume than equity bubbles, our model predicts that debt may naturally be under-priced.

At the start, nature draws θ according to a normal distribution with mean 0, and unconditional variance σ_θ^2 , $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Each informed trader i then receives a noisy private signal x_i which is normally distributed with a mean θ and a variance β^{-1} , and is i.i.d. across traders (conditional on θ), $x_i \sim \mathcal{N}(\theta, \beta^{-1})$. Each trader decides whether to purchase up to one share of the asset at the prevailing price P , in exchange for cash. Formally, trader i submits a price-contingent demand schedule $d_i(\cdot)$ to maximize her expected wealth $w_i = d_i \cdot (\pi(\theta) - P)$. Traders cannot short-sell the asset or buy additional shares, restricting demand to $[0, 1]$. Individual trading strategies are then a mapping $d : \mathbb{R}^2 \rightarrow [0, 1]$ from signal-price pairs (x_i, P) into asset holdings. Aggregating traders' decisions leads to the aggregate demand by informed traders, $D : \mathbb{R}^2 \rightarrow [0, 1]$,

$$D(\theta, P) = \int d(x, P) d\Phi(\sqrt{\beta}(x - \theta)), \quad (1)$$

where $\Phi(\cdot)$ denotes a cumulative standard normal distribution, and $\Phi(\sqrt{\beta}(x - \theta))$ represents the cross-sectional distribution of private signals x_i conditional on the realization of θ .⁵ In addition, there is stochastic demand for the asset from noise traders, which takes the form $\Phi(u)$, where u is normally distributed with mean zero and variance σ_u^2 , $u \sim \mathcal{N}(0, \sigma_u^2)$, independently of θ . This specification is adapted from Hellwig, Mukherji, and Tsyvinski (2006), and allows us to preserve the tractability of Bayesian updating with normal posterior beliefs.⁶

Once all traders have submitted their orders, the auctioneer selects a price P to clear the market. Formally, the market-clearing price function $P : \mathbb{R}^2 \rightarrow \mathbb{R}$ selects P from the correspondence $\hat{P}(\theta, u) = \{P \in \mathbb{R} : D(\theta, P) + \Phi(u) = 1\}$, for all $(\theta, u) \in \mathbb{R}^2$.⁷ After all trades are completed, the dividends $\pi(\theta)$ are realized and disbursed to the owners of the asset.

Let $H(\cdot|x, P) : \mathbb{R} \rightarrow [0, 1]$ denote the traders' posterior cdf of θ , conditional on observing a private signal x , and conditional on the market price P . A *Perfect Bayesian Equilibrium* consists of demand functions $d(x, P)$ for informed traders, a price function $P(\theta, u)$, and posterior beliefs $H(\cdot|x, P)$ such that (i) $d(x, P)$ is optimal given $H(\cdot|x, P)$; (ii) the asset market clears for all (θ, u) ; and (iii) $H(\cdot|x, P)$ satisfies Bayes' rule whenever applicable, i.e., for all p such that $\{(\theta, u) : P(\theta, u) = p\}$ is non-empty.

⁵We assume that the Law of Large Numbers applies to the continuum of traders, so that conditional on θ the cross-sectional distribution of signal realizations ex post is the same as the ex ante distribution of traders' signals.

⁶We generalize this demand specification in Section 4.2 allowing for price-elastic demands by noise traders.

⁷We can without loss of generality restrict the range of $P(\cdot)$ to coincide with the range of $\pi(\cdot)$.

2.2 Equilibrium Characterization

We begin by characterizing informed traders' demand. With risk-neutrality, the trader's expected value of holding the asset is $\int \pi(\theta)dH(\theta|x, P)$. Since private signals are log-concave and $\pi(\cdot)$ is increasing in θ , posterior beliefs $H(\cdot|x, P)$ are first-order stochastically increasing in x , and $\int \pi(\theta)dH(\theta|x, P)$ is strictly increasing in x , for any P that is observed in equilibrium (Milgrom, 1981). The traders' decisions are therefore characterized by a signal threshold function $\hat{x} : \mathbb{R} \rightarrow \mathbb{R} \cup \{\pm\infty\}$, such that $d(x_i, P) = \mathbb{I}_{x_i \geq \hat{x}(P)}$, that is, the trader places an order to buy a share at price P , if and only if $x_i \geq \hat{x}(P)$. We call the trader who observes the signal equal to the threshold, $x = \hat{x}(P)$, and who is therefore indifferent, the *marginal trader*. The signal threshold is uniquely defined by

$$\begin{aligned}\hat{x}(P) &= +\infty \text{ if } \lim_{x \rightarrow +\infty} \int \pi(\theta)dH(\theta|x, P) \leq P, \\ \hat{x}(P) &= -\infty \text{ if } \lim_{x \rightarrow -\infty} \int \pi(\theta)dH(\theta|x, P) \geq P, \\ P &= \int \pi(\theta)dH(\theta|\hat{x}(P), P) \text{ otherwise.}\end{aligned}\tag{2}$$

Expression (2) illustrates three cases: (i) if the most optimistic trader's expected dividend is lower than the price, no trader buys, so the signal threshold becomes $+\infty$; (ii) if the most pessimistic trader's expected dividend exceeds the price, all traders buy, and the threshold for buying is $-\infty$; (iii) only some traders buy, and the threshold $\hat{x}(P)$ takes an interior value at which the marginal trader's posterior expectation of the dividend must equal the price. Aggregating the individual trading decisions, the informed demand is $D(\theta, P) = \int_{\hat{x}(P)}^{\infty} 1 \cdot d\Phi(\sqrt{\beta}(x - \theta)) = 1 - \Phi(\sqrt{\beta}(\hat{x}(P) - \theta))$, which equals 0 if $\hat{x}(P) = +\infty$, and 1 if $\hat{x}(P) = -\infty$.

Next, we analyze the market-clearing condition. Since $\Phi(u) \in (0, 1)$, in equilibrium, $\hat{x}(\cdot)$ must be finite for all P on the equilibrium path, and satisfy the third condition in (2). From the market-clearing condition, we then have $\Phi(\sqrt{\beta}(\hat{x}(P) - \theta)) = \Phi(u)$, which allows us to characterize the correspondence of market-clearing prices:

$$\hat{P}(\theta, u) = \left\{ P \in \mathbb{R} : \hat{x}(P) = \theta + \frac{1}{\sqrt{\beta}}u \right\}.\tag{3}$$

From now on, we focus on equilibria in which the price is conditioned on (θ, u) through the observable state variable $z \equiv \theta + 1/\sqrt{\beta} \cdot u$. The next lemma characterizes the resulting equilibrium beliefs. All proofs are provided in the appendix.

Lemma 1 (Information Aggregation) (i) *In any equilibrium with conditioning on z , the equilibrium price function $P(z)$ is invertible.* (ii) *Equilibrium beliefs for price realizations observed*

along the equilibrium path are given by

$$H(\theta|x, P) = \Phi \left(\frac{\sqrt{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \left(\theta - \frac{\beta x + \beta\sigma_u^{-2} \cdot \hat{x}(P)}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \right)}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \right). \quad (4)$$

Part (ii) of the Lemma exploits the invertibility to arrive at a complete characterization of posterior beliefs $H(\cdot|x, P)$. With invertibility, we can summarize information conveyed by the price through z . Conditional on θ , z is normally distributed with mean θ and variance σ_u^2/β . Thus, the price is isomorphic to a normally distributed signal of θ , with a precision that is increasing in the precision of private signals, and decreasing in the variance of demand shocks.

Using Lemma 1 we rewrite (2), the indifference condition that defines the signal threshold $\hat{x}(P)$:

$$P = \int \pi(\theta) d\Phi \left(\frac{\sqrt{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \left(\theta - \frac{\beta + \beta\sigma_u^{-2}}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \hat{x}(P) \right)}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \right). \quad (5)$$

This condition equates P to the marginal trader's expectation of dividends. The latter also depends on P through its effect on posterior beliefs. Using the market-clearing condition $\hat{x}(P) = z$, Proposition 1 uniquely characterizes the market equilibrium.⁸

Proposition 1 (Asset market equilibrium) *For any increasing dividend function $\pi(\cdot)$, an asset market equilibrium exists, is unique, and is characterized by the price function $P_\pi(z)$ and the traders' threshold function $\hat{x}(p) = z = P_\pi^{-1}(p)$, where*

$$P_\pi(z) = \mathbb{E}(\pi(\theta)|x = z, z) = \int \pi(\theta) d\Phi \left(\frac{\sqrt{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \left(\theta - \frac{\beta + \beta\sigma_u^{-2}}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} z \right)}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}} \right). \quad (6)$$

The price function $P_\pi(z)$ is uniquely defined and strictly monotone, and therefore defines the unique market equilibrium.⁹

3 The Information Aggregation Wedge

3.1 Conditional Information aggregation wedge

We now discuss how noisy information affects equilibrium prices and expected dividend values. To be precise, we form expectations of dividends from the perspective of an outside observer (or “econometrician”) who does not have access to any private signal about θ , but knows the parameters

⁸Notice that this only implies the uniqueness of the equilibrium that conditions on the summary statistic z , not overall uniqueness of the equilibrium characterized in proposition 1.

⁹We index an equilibrium function or variable by π to make explicit that it is derived from a specific dividend function $\pi(\cdot)$, i.e. $P_\pi(\cdot)$ is the equilibrium price function that is derived from dividend function $\pi(\cdot)$ by equation (6).

of the game and observes the realization of the price P , or equivalently the state z . This outsider holds a conditional belief that $\theta|z \sim \mathcal{N}(\beta\sigma_u^{-2} / (\sigma_\theta^{-2} + \beta\sigma_u^{-2}) \cdot z, (\sigma_\theta^{-2} + \beta\sigma_u^{-2})^{-1})$, and therefore has an expectation of dividends conditional on public information z , denoted $V_\pi(z)$:

$$V_\pi(z) = \mathbb{E}(\pi(\theta)|z) = \int \pi(\theta) d\Phi \left(\sqrt{\sigma_\theta^{-2} + \beta\sigma_u^{-2}} \left(\theta - \frac{\beta\sigma_u^{-2}}{\sigma_\theta^{-2} + \beta\sigma_u^{-2}} z \right) \right). \quad (7)$$

The main observation from comparing Proposition 1 with equation (7) is that at the interim stage –when the share price is observed but before dividends are realized– the equilibrium price differs from the expected dividend, conditional on the public information. This difference is due to the impact of private information on equilibrium prices. We label this difference the *information aggregation wedge*, $W_\pi(z) \equiv P_\pi(z) - V_\pi(z)$.

The choice of $V_\pi(\cdot)$ as a natural benchmark of comparison for $P_\pi(\cdot)$ follows from the fact that $V_\pi(\cdot)$ also corresponds to the expected dividend (or in a sufficiently large data set, to the average dividend), conditional on the observation of P (recall that $P_\pi(\cdot)$ is invertible). This benchmark differs from the one chosen e.g. by Harrison and Kreps (1978), who compare an asset’s value to the dividend expectation of any trader in their market, or to an average of those expectations, as in Allen, Morris and Shin (2006) and Bacchetta and van Wincoop (2006).¹⁰ In our formulation, $V_\pi(\cdot)$ and $P_\pi(\cdot)$ both have direct empirical counterparts in any set of price-return data, and this formulation therefore allows us to directly focus on the empirical, testable implications of our model.


The price $P_\pi(z)$, and the expected dividend conditional on public information, $V_\pi(z)$, differ in how expectations of θ are formed. The price equals the dividend expectation of the marginal trader who is indifferent between keeping or selling her share. This trader conditions on the market signal z , as well as a private signal whose realization must equal the threshold $\hat{x}(P)$ in order to be consistent with the trader’s indifference condition. The trader treats these two sources of information as mutually independent signals of θ . At the same time, the market-clearing condition implies that $\hat{x}(P)$ must equal z in order to equate demand and supply of shares. The marginal trader’s expectation $\mathbb{E}(\pi(\theta)|x = z, z)$ thus behaves as if she received one signal z of precision $\beta + \beta\sigma_u^{-2}$ instead of $\beta\sigma_u^{-2}$. In contrast, the expected dividends $\mathbb{E}(\pi(\theta)|z)$ conditional on P (or equivalently z) weighs z according to its true precision $\beta\sigma_u^{-2}$.

The difference in the responsiveness of the price relative to the expected dividend conditional in the price results from the compositional shift in the identity of the traders holding the shares.

¹⁰Like these antecedents, however, our characterization of the wedge, as well as its implications for returns and volatility, can also be understood in terms of a failure of the law of iterated expectations, i.e. generically, $\mathbb{E}(\mathbb{E}(\pi(\theta)|x = z, z)) \neq \mathbb{E}(\mathbb{E}(\pi(\theta)|z)) = \mathbb{E}(\pi(\theta))$.

This is depicted in figure 1. Any increase in z shifts the identity of the marginal trader's private signal one-for-one. If θ increases, the distribution of private signals shifts up, so for a given signal threshold, demand for the asset by informed traders increases, but demand from uninformed traders is unchanged. If instead s increases, uninformed demand increases, but informed demand remains the same. In both cases, the asset is relatively scarcer for informed traders, so the threshold for the informed trader's private signal has to increase in order to clear the market. In addition to this compositional shift (which only appears in the expectation of the marginal trader, e.g. the price), all traders, as well as the uninformed outsider, recognize that an increase in z , as revealed through P shifts up their expectation of the state θ . This is reflected in the weight $\beta\sigma_u^{-2}$ attributed to z in both $P_\pi(z)$ and $V_\pi(z)$.

Figure 1: Marginal Trader Identity Shift



Belief heterogeneity and limits to arbitrage are both necessary ingredients to obtain the wedge. If instead all informed traders have access to a common signal z of fundamentals, they all hold identical expectations and must be indifferent between buying and not buying to clear the market. But this requires that the price equals the common expectation of the dividend, i.e. $P_\pi(z) = V_\pi(z)$. The same result applies with free entry of uninformed arbitrageurs (Kyle, 1985).

The remainder of this subsection describes properties of the wedge, conditional on z , which will form the basis for our main results on expected returns and price volatility. To this end, we define

$$\gamma_P \equiv \frac{\beta + \beta\sigma_u^{-2}}{\sigma_\theta^{-2} + \beta + \beta\sigma_u^{-2}}, \quad \text{and} \quad \gamma_V \equiv \frac{\beta\sigma_u^{-2}}{\sigma_\theta^{-2} + \beta\sigma_u^{-2}}$$

as the response coefficients of the expectations of θ entering $P_\pi(\cdot)$ and $V_\pi(\cdot)$ to innovations in z .

The equilibrium price and expected dividends are then rewritten as:

$$\begin{aligned} P_\pi(z) &= \int \pi(\gamma_P z + \sigma_\theta \sqrt{1 - \gamma_P u}) \phi(u) du \\ V_\pi(z) &= \int \pi(\gamma_V z + \sigma_\theta \sqrt{1 - \gamma_V u}) \phi(u) du \end{aligned}$$

This formulation summarizes the difference between the price and the expected dividend by the response parameters γ_P and γ_V , which measure the marginal trader's and outsider's update of θ to z . These parameters enter $P_\pi(\cdot)$ and $V_\pi(\cdot)$ in two ways: the marginal trader's expectation responds more strongly to z , and his residual uncertainty about θ (after observing z) is lower: $\sigma_\theta^2(1 - \gamma_P)$ instead of $\sigma_\theta^2(1 - \gamma_V)$. Using a third-order Taylor expansion, we approximate the wedge by

$$W_\pi(z) \approx \pi(\gamma_P z) - \pi(\gamma_V z) + \frac{\sigma_\theta^2}{2} [\pi''(\gamma_P z)(1 - \gamma_P) - \pi''(\gamma_V z)(1 - \gamma_V)]. \quad (8)$$

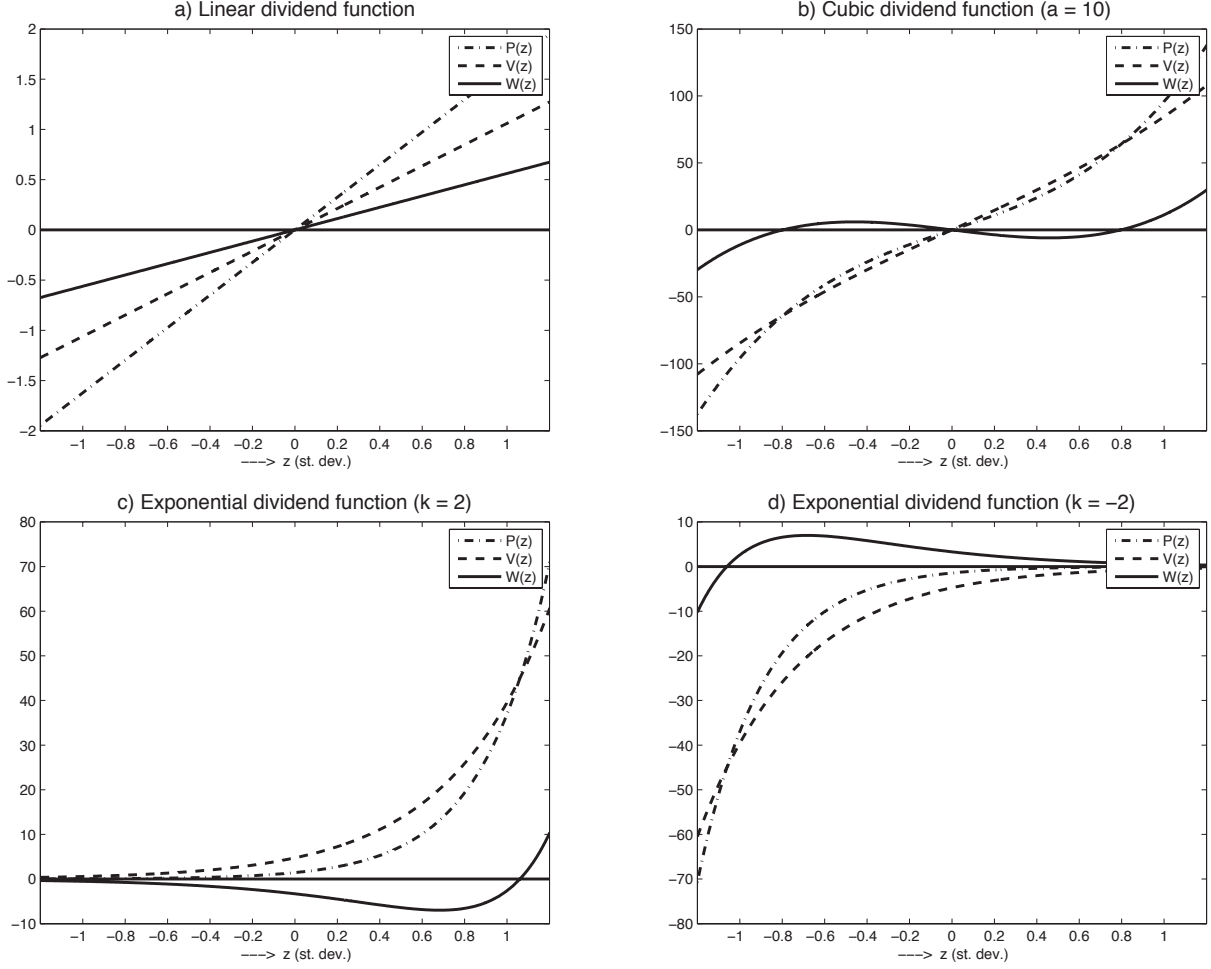
The term $\pi(\gamma_P z) - \pi(\gamma_V z)$ captures the shift in expectations, while the second term in squared brackets captures the role of residual uncertainty. The latter plays a role only if $\pi(\cdot)$ is non-linear, and in that case matters through second- and higher derivatives. The shift in expectations from $\gamma_V z$ to $\gamma_P z$ amounts to a mean-preserving spread from an ex ante perspective, and is therefore a source of increased variability in the price, relative to expected fundamentals: $\pi(\gamma_P z) - \pi(\gamma_V z)$ crosses 0 at a single point where $z = 0$, is negative when $z < 0$, and positive when $z > 0$.

When $\pi(\cdot)$ is linear, the higher sensitivity of expectations to z is the only effect determining the wedge, while the residual uncertainty plays no role. Panel a) of figure 2 plots the price (solid line), the expected dividend (dashed line) and conditional wedge (dashed-dotted line) as a function of the state variable z , for $\pi(\theta) = \theta$. The price is more sensitive to innovations in z than the expected dividend, resulting in a wedge $W_\pi(z) = (\gamma_P - \gamma_V)z$ that is negative for $z < 0$, zero for $z = 0$, and positive for $z > 0$.

For non-linear dividends, residual uncertainty shifts the level of the wedge up or down, depending on a comparison between the residual uncertainty levels $1 - \gamma_P$ relative to $1 - \gamma_V$, and the second derivatives $\pi''(\gamma_P z)$ and $\pi''(\gamma_V z)$. At the prior mean $z = 0$, the second derivatives are comparable, so the reduction of uncertainty implies a negative wedge if $\pi''(0) > 0$, and a positive wedge, if $\pi''(0) < 0$. Away from $z = 0$, the third- and higher derivatives may reduce or even overturn this effect, and therefore make it impossible to offer precise results on the shape of $W_\pi(\cdot)$ without additional restrictions. We illustrate these possibilities with two parametric examples that follow.

Example 1: Exponential dividend function

Figure 2: Conditional Price, Expected Dividend, and Wedge



Suppose that $\pi(\theta) = \frac{1}{k}e^{k\theta}$, with $k \neq 0$. Expected dividends, prices and the wedge are then characterized by:

$$V_{\pi}(z) = \frac{1}{k}e^{k\gamma_V z + \frac{k^2}{2}\sigma_{\theta}^2(1-\gamma_V)}, \quad P_{\pi}(z) = \frac{1}{k}e^{k\gamma_P z + \frac{k^2}{2}\sigma_{\theta}^2(1-\gamma_P)}$$

$$W_{\pi}(z) = P_{\pi}(z) \left(1 - e^{-k(\gamma_P - \gamma_V)z + \frac{k^2}{2}(\gamma_P - \gamma_V)\sigma_{\theta}^2} \right).$$

In this case, the price and expected dividend are both exponential functions in z , with a stronger reaction of prices to z . The residual uncertainty affects both $V_{\pi}(z)$ and $P_{\pi}(z)$ multiplicatively, but the factor is larger for $V_{\pi}(z)$, reflecting the fact that residual uncertainty is greater for the outsider.

If $k > 0$, we then have a dividend function that is increasing, convex, and bounded below by zero (figure 2, panel c). The wedge is negative at $z = 0$ and non-monotone. It decreases at first,

reaches its lowest value at some intermediate point, and is increasing and convex from there on, crossing 0 at $z = \frac{k}{2}\sigma_\theta^2 > 0$. The reverse image obtains when $k < 0$, in which case π is increasing, concave, and bounded above by zero (figure 2, panel d). For negative z , the wedge is negative at first and increasing in z , crossing 0 at $z = \frac{k}{2}\sigma_\theta^{-2} < 0$. It reaches its maximum value at a negative z and then monotonically converges towards 0. This example thus confirms the intuitions from the shift in means which makes $P_\pi(z)$ more responsive to a shift in z , and the shift in residual uncertainty that is captured by the multiplicative factors. The curvature parameter k governs the shape of the wedge function, and whether the residual uncertainty increases or decreases the wedge.

We use this example to illustrate our two main results. First, we show that the expected wedge is positive if and only if $k > 0$, and negative if $k < 0$. That is, the security trades at a premium in the case with convex dividends and upside risks, and at a discount in the case with concave dividends and downside risks. Taking expectations, we have

$$\begin{aligned} \mathbb{E}(V_\pi(z)) &= 1/k \cdot e^{\frac{k^2}{2}\sigma_\theta^2}, & \mathbb{E}(P_\pi(z)) &= 1/k \cdot e^{\frac{k^2}{2}\sigma_\theta^2} [1 + (\frac{\gamma_P}{\gamma_V} - 1)\gamma_P] \\ \text{and } \mathbb{E}(W_\pi(z)) &= 1/k \cdot e^{\frac{k^2}{2}\sigma_\theta^2} \left\{ e^{\frac{k^2}{2}\sigma_\theta^2(\gamma_P/\gamma_V - 1)\gamma_P} - 1 \right\}, \end{aligned}$$

which is positive whenever $k > 0$, and negative for $k < 0$ (and can be checked to approach 0 continuously as $k \rightarrow 0$).

Second, we show that the model exhibits excess price volatility. Focusing on log variances for analytical convenience, we have: $Var(\log \pi(\theta)) = k^2\sigma_\theta^2$, $Var(\log V_\pi(z)) = \gamma_V k^2\sigma_\theta^2$, and $Var(\log P_\pi(z)) = \gamma_P^2/\gamma_V \cdot k^2\sigma_\theta^2$. Therefore, we observe that $Var(\log P_\pi(z)) > Var(\log V_\pi(z))$, for any parameter set. Moreover, if the information aggregation wedge becomes sufficiently important, then we may have $\gamma_P^2/\gamma_V > 1$, and therefore $Var(\log P_\pi(z)) > Var(\log \pi(\theta))$. In particular, this is a result of the following two limiting scenarios: (i) if for given $\gamma_V < 1$, γ_P approaches 1, i.e. the informed traders have very precise signals for given level of information in the price, or (ii) if $\gamma_V \rightarrow 0$, while γ_P is bounded away from 0. In this case, the market becomes very noisy, for a given level of private information. On the other hand, $Var(\log V_\pi(z))$ is always less than $Var(\log \pi(\theta))$, which is a direct application of Blackwell's Theorem on comparison of information structures (Blackwell, 1951, 1953).

Our main two theorems that follow generalize these observations about the unconditional wedge and excess price volatility. Theorem 1 below establishes that the sign and magnitude of the average wedge on the comparison of upside vs. downside risks. Theorem 2 generalizes the result that prices are more variable than expected dividends, and in some cases even more variable than realized dividends. Our second example, however, reinforces the observation that the conditional wedge

$W_\pi(\cdot)$ need not be monotone in general, and may also cross 0 at multiple points, which rules out conditional or local versions of these results without imposing additional assumption on dividends.

Example 2: Cubic dividend function

Suppose that $\pi(\theta) = \theta + a\theta^3$, with $a > 0$ to ensure monotonicity of π . For a cubic function (figure 2, panel b), the above approximation holds exactly, so that

$$W_\pi(z) = (\gamma_P - \gamma_V)z + a(\gamma_P^3 - \gamma_V^3)z^3 + 3az\sigma_\theta^2[\gamma_P(1 - \gamma_P) - \gamma_V(1 - \gamma_V)],$$

where the first two terms correspond to the shift in means, and the last to the shift in residual uncertainty. If $\gamma_P + \gamma_V > 1$ and a sufficiently large, $W'_\pi(0) < 0$. Since $W_\pi(0) = 0$, it follows immediately that $W_\pi(\cdot)$ is non-monotone and crosses 0 in three different locations.

3.2 Unconditional information aggregation wedge

To obtain general results, we focus on unconditional moments of prices and expected dividends. Let $W_\pi = \mathbb{E}(W_\pi(z))$ denote the expected information aggregation wedge associated with a payoff function $\pi(\cdot)$. The next lemma provides a characterization of W_π which forms the basis for the subsequent comparative statics results.

Lemma 2 (Unconditional Wedge) *Define σ_P as $\sigma_P^2 = \sigma_\theta^2(1 + (\gamma_P/\gamma_V - 1)\gamma_P)$. The unconditional information aggregation wedge W_π is characterized by*

$$W_\pi = \int_0^\infty (\pi'(\theta) - \pi'(-\theta)) \left(\Phi\left(\frac{\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\theta}{\sigma_P}\right) \right) d\theta. \tag{9}$$

This characterization shows how the wedge depends separately on both the curvature the payoff function, and the parameters describing the informational environment. The parameter $\sigma_P > \sigma_\theta$ corresponds to the prior variance of θ , as assessed by the marginal trader, and summarizes the importance of informational frictions in the market. By taking ex ante expectations over z , the shifts in mean and residual uncertainty combine into a mean-preserving spread between the weights that the marginal trader and the outsider associate with each realizations of θ .

The marginal trader places more weight on the tails of the fundamental distribution, from an ex ante perspective (i.e., $\sigma_P > \sigma_\theta$). This result can intuitively be understood as follows: the marginal trader’s posterior of θ , conditional on z , is normal with mean $\gamma_P z$ and variance $(1 - \gamma_P)\sigma_\theta^2$. The prior over z is normal with mean 0 and variance σ_θ^2/γ_V . Compounding the two distributions, the marginal trader’s prior over θ is characterized as a normal distribution with mean 0 and variance

$(1 - \gamma_P) \sigma_\theta^2 + \gamma_P^2 \sigma_\theta^2 / \gamma_V = \sigma_P^2$. The outsider, on the other hand, holds the posterior that conditional on z , θ is normal with mean $\gamma_V z$ and variance $(1 - \gamma_V) \sigma_\theta^2$. His compounded distribution then corresponds to the actual prior distribution of θ , as the prior variance is just $(1 - \gamma_V) \sigma_\theta^2 + \gamma_V \sigma_\theta^2 = \sigma_\theta^2$. Hence, the information frictions summarized by the distance of σ_P from σ_θ will be large whenever the market signal is noisy relative to private signals, or the ratio γ_P / γ_V is high, as this leads to a large discrepancy between the posterior beliefs held by the marginal trader and the outsider.

We use Lemma 2 to sign the unconditional wedge as a function of the shape of the dividend function, and to offer comparative statics with respect to π and the informational parameters γ_P and γ_V . Our next definition provides a partial order on payoff functions that we will use for the comparative statics.

Definition 1 (i) A dividend function π has symmetric risks if $\pi'(\theta) = \pi'(-\theta)$ for all $\theta > 0$.

(ii) A payoff function π is dominated by upside risks, if $\pi'(\theta) \geq \pi'(-\theta)$ for all $\theta > 0$. A payoff function π is dominated by downside risks, if $\pi'(\theta) \leq \pi'(-\theta)$ for all $\theta > 0$.

(iii) A dividend function π_1 has more upside (less downside) risk than π_2 if $\pi_1'(\theta) - \pi_1'(-\theta) \geq \pi_2'(\theta) - \pi_2'(-\theta)$ for all $\theta > 0$.

This definition classifies payoff functions by comparing marginal gains and losses at fixed distances from the prior mean to determine whether the payoff exposes its owner to bigger payoff fluctuations on the upside or the downside. Any linear dividend function has symmetric risks, any convex function is dominated by upside risks, and any concave dividend function is dominated by downside risks. The classification however also extends to non-linear functions with symmetric gains and losses, as well as non-convex functions with upside risk or non-concave functions with downside risk. Figure 3 plots examples of payoff functions dominated by different types of risk.

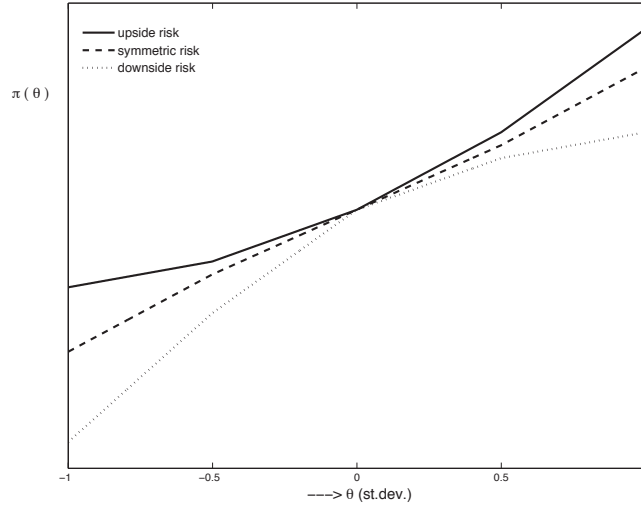
The following Theorem summarizes the comparative statics implications that follow directly from this partial order, and the characterization in lemma 2.

Theorem 1 (Average prices and returns) (i) **Sign:** If π has symmetric risk, then $W_\pi = 0$. If π is dominated by upside risk, then $W_\pi \geq 0$. If π is dominated by downside risk, then $W_\pi \leq 0$.

(ii) **Comparative Statics w.r.t. π :** For given σ_P^2 , if π_1 has more downside and less upside risk than π_2 , then $W_{\pi_2} \geq W_{\pi_1}$.

(iii) **Comparative Statics w.r.t. σ_P^2 :** If π is dominated by upside or downside risk, then $|W_\pi|$ is increasing in σ_P . Moreover, $\lim_{\sigma_P \rightarrow \sigma_\theta} |W_\pi| = 0$, and $\lim_{\sigma_P \rightarrow \infty} |W_\pi| = \infty$, whenever there exists $\varepsilon > 0$, such that $|\pi'(\theta) - \pi'(-\theta)| > \varepsilon$ for all $\theta \geq \varepsilon$.

Figure 3: Dividend risk types



(iv) **Increasing differences:** If π_1 has more upside risk than π_2 , then $W_{\pi_1}(\sigma_P) - W_{\pi_2}(\sigma_P)$ is increasing in σ_P .

This theorem summarizes how the shape of the dividend function and the informational parameters combine to determine the sign and magnitude of the unconditional information aggregation wedge. It shows that unconditional price premia or discounts arise as a combination of two elements: upside or downside risks in the dividend profile π , and an impact of private information on market prices ($\gamma_P > \gamma_V$). The latter requires that updating from prices is noisy ($\gamma_V < 1$). This Theorem forms the first part of our core theoretical contribution, and shows that noisy information aggregation may influence conditional and unconditional returns of assets through their payoff profile and the informational characteristics of the market.

The result is easily understood from our interpretation of the wedge as the expected value of a symmetric, mean-preserving spread of the true underlying fundamental distribution.

Part (i) shows that the sign of the wedge is determined by whether π is dominated by upside, downside, or symmetric risk. When the dividend function has symmetric risk, the gains from this spread on the upside exactly cancel the expected losses on the downside, and the total effect is 0. When the dividend is dominated by upside risks, the expected upside gains dominate and the value of the mean-preserving spread is positive, leading to a positive unconditional wedge. Conversely, when the dividend is dominated by downside risks, the expected losses on the downside dominate and the expected value of the spread is negative.

Parts (ii), (iii), and (iv) complement the first result on the possibility of price premia or discounts with specific predictions on how its magnitude depends on cash flow and informational characteristics.

Part (ii) shows that an asset with more upside or less downside risk on average has a higher price premium or a lower price discount, all else equal. Thus, returns on average are lower (and prices higher) for securities that represent more upside risks. Simply put, the mean-preserving spread becomes more valuable when the payoff function shifts towards more upside risk.

Part (iii) shows the role of informational parameters. For a given payoff function, the unconditional wedge increases in absolute value as the information aggregation friction has bigger effects (higher σ_P). For a given set of upside or downside risks, a bigger mean-preserving spread generates bigger gains or losses. Moreover, a wedge obtains only if $\gamma_P > \gamma_V$, i.e., if the heterogeneous beliefs have an impact on price. The wedge is increasing in γ_P and decreasing in γ_V , as the precision of market information and private information move the wedge in opposite directions. Under regularity conditions, which ensure that the payoff asymmetry doesn't disappear in the tails, the absolute value of the wedge approaches infinity when $\gamma_V \rightarrow 0$. This obtains if for a given value of β , the market noise becomes infinitely large. In this limiting case, the marginal trader remains responsive to z , even though the z is infinitely noisy.

Part (iv) shows that the unconditional wedge has increasing differences between the dominance of upside risk and the level of market noise. This implies that the effects of market noise and asymmetry in dividend risk are mutually reinforcing on the magnitude of the wedge.

Importantly, our results on differences between expected prices and dividends are not a consequence of irrational trading strategies, behavioral biases of investors, or agency conflicts. Nor are such differences accounted for by risk premia (since traders are risk neutral). Our model thus offers a theory in which average prices can differ systematically from expected dividends as a result of the interplay between the dividend structure and the partial aggregation of information into prices, in a context where traders hold heterogeneous beliefs in equilibrium and arbitrage is limited. To our knowledge, this channel is new to the literature.

3.3 Excess Price Variability

Our second main result concerns the variability of prices, relative to expected dividends and realized dividends. As can readily be seen from the above characterizations, if $W'_\pi(\cdot) > 0$, the unconditional variance of prices (prior to realization of z) exceeds the variability of expected dividends. Consider furthermore the limiting case where $\gamma_P \rightarrow 1$, in which $P_\pi(z) \rightarrow \pi(z)$. Since the variance of z

exceeds that of θ , it follows immediately that in this limit, where the informed traders' signals become arbitrarily precise, the variability of prices can exceed the variability of dividends. This result is illustrated in the linear and the log-normal examples discussed in section 3.1.

Our second theorem generalizes these observations. To do so, we will need to impose some restrictions to handle the non-linearities and higher-order effects that are confounding the comparative statics of W_π with respect to z . Concretely, we will focus on risks that are symmetric or dominated by the upside or downside, and we will focus on $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right)$, $\mathbb{E}\left((V_\pi(z) - V_\pi(0))^2\right)$, and $\mathbb{E}\left((\pi(\theta) - \pi(0))^2\right)$ as our criterion for the variability of prices, expected dividends, and realized dividends, respectively, rather than the unconditional variances. The next theorem states our main result concerning excess variability:

Theorem 2 (Excess variability of prices) *For any payoff function $\pi(\cdot)$ that is symmetric, dominated by upside, or dominated by downside risk:*

(i) *The variability of expected dividends is always less than the variability of realized dividends and the variability of prices:*

$$\mathbb{E}\left((V_\pi(z) - V_\pi(0))^2\right) < \mathbb{E}\left((\pi(\theta) - \pi(0))^2\right) \quad \text{and} \quad \mathbb{E}\left((V_\pi(z) - V_\pi(0))^2\right) < \mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right)$$

(ii) *The excess variability of prices relative to expected dividends is increasing in γ_P and decreasing in γ_V .*

(iii) *For any γ_V , if γ_P is sufficiently high, then the variability of prices exceeds the variability of realized dividends. The same occurs if, for given γ_P , γ_V is sufficiently low.*

(iv) *If $\pi(\cdot)$ is unbounded on one side, then $\lim_{\gamma_V \rightarrow 0} \mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right) = \infty$.*

This theorem shows that the price is more variable than expected dividends, and if the market is sufficiently noisy, even more variable than realized dividends. The latter occurs in the limiting cases where supply shocks are unboundedly large ($\sigma_u^2 \rightarrow \infty$, $\gamma_V \rightarrow 0$), or the traders' private information is infinitely precise ($\beta \rightarrow \infty$, $\gamma_P \rightarrow 1$). In the former case, the variability of prices can be arbitrarily large, even as the variability of realized and expected dividends is bounded. The statement of the result relies on two restrictions which we used for analytical tractability. First, the focus on a variability measure which combines a variance with a bias between the average price and the price at the average fundamental. Second, we restrict ourselves to symmetric, upside or downside risks. With these restrictions, the results are the cleanest, and easiest to interpret.

To understand this result, and the source of excess price variability in our model, it is useful to think of a counter-factual third person who observes a signal z with distribution $z|\theta \sim$

$\mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$. Like the uninformed outsider, this third person is fully Bayesian, but has access to a more informative signal, whose precision matches that of the marginal trader's. Therefore in comparison to the marginal trader, the third person will form the same posterior beliefs, conditional on a realization of z , but z will be drawn from a distribution with a lower ex ante variance, and be consistent with Bayes' Rule derived from the objective signal precision. In comparison to the uninformed outsider, the third person is also fully Bayesian, but with simply a more precise signal. We break down the comparison between $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right)$ and the other terms into a comparison between $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right)$ and $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2 | z \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right)$, and the comparison of this latter term with the ex ante variability of expected and realized dividends. $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2 | z \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right)$ corresponds to the counter-factual variability of prices, if z had been drawn from a distribution $z|\theta \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$, such that $P_\pi(z)$ is consistent with a posterior expectation of π conditional on z .

For the comparison of the counter-factual variability of prices with the variability in expected and realized dividends, we first proceed to break down the variability measures into a variance and a bias term. The variance terms can then be compared using Blackwell's theorem on the comparison of experiments (Blackwell, 1951, 1953). Since $\pi(\theta)$, $P_\pi(z)$, and $V_\pi(z)$ correspond to the posterior expectation of $\pi(\theta)$ for respectively, an agent who observes the true θ , the counter-factual signal z , and the actual signal z , the unconditional variance of π exceeds the unconditional variance of P under the distribution $z \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$, which exceeds the unconditional variance of V under the distribution $z \sim \mathcal{N}(\theta, \sigma_u^2/\beta)$. For symmetric, upside and downside risks, the bias terms follow exactly the same ranking.¹¹

Therefore, if this second term was the only relevant component, the variability of prices in our model would satisfy the standard conditions resulting from the Blackwell comparison of experiments - namely, that a more informative price signal raises price volatility and expected dividend volatility, but both are bounded by the volatility of realized dividends. At best, the volatility gap can be brought close to zero when information in the market is sufficiently precise. For models in which asset prices are always equal to expected future dividends, this observation is made precise by West (1988).

The excess volatility then results from the first term, which measures the over-reaction of the

¹¹Our choice of variability measure (which is equivalent to the variance for symmetric risks) allows for the cleanest possible comparison between the actual and the counter-factual variance of prices. This variability measure then introduces the need to also rank the bias terms, which is done for symmetric, upside or downside risks. Since the bias terms are likely to be small compared to the variances, similar, but technically less clean results are likely to hold for arbitrary risks or other variability measures.

price compared to its true information content. This term measures the difference between the variability in prices under the objective signal distribution $z \sim \mathcal{N}(\theta, \sigma_u^2/\beta)$ with the variability in prices (derived from the same price function) for a counter-factual signal distribution $z \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$, under which the market's beliefs are consistent with Bayes' Rule. This over-reaction effect is always positive, and may be strong enough to cause the volatility of prices to exceed the volatility of realized dividends. This becomes possible in particular when information frictions in the market (as measured by the gap between γ_P and γ_V) are sufficiently severe.

To conclude, we point out that the same forces that lead to large excess volatility in prices also generate a low correlation of prices with realized dividends.

Proposition 2 (Low predictability of dividends) *Fix $\gamma_P > 0$. Then*

$$\lim_{\gamma_V \rightarrow 0} \text{corr}(P_\pi(z), \pi(\theta)) = 0 \text{ and } \lim_{\gamma_V \rightarrow 0} \frac{\text{cov}(P_\pi(z), \pi(\theta))}{\text{Var}(P_\pi(z))} = 0.$$

The key to this result is to note that the unconditional correlation of prices and realized dividends (in absolute value) is bounded above by the ratio between the unconditional variances of the expected and realized dividends $V_\pi(z)$ and $\pi(\theta)$. Likewise, the OLS regression coefficient for regressing realized dividends against prices is bounded by the ratio of the unconditional variances of $V_\pi(z)$ and $P_\pi(z)$. When γ_V is sufficiently low, i.e. when the market signal is very uninformative, then these ratios are close to zero (i.e. the posterior expectation remains much closer to the prior expectation than the actual dividend realization, and the posterior expectation is much less volatile than the price). In this case, the predictability of dividends from prices is very low. This turns out to be precisely the case in which the information aggregation wedge also has the potential to generate large excess price volatility.

These results offer a new perspective on the well documented “excess volatility puzzle” (Le Roy and Porter 1981; Shiller 1981), and the low predictability of future dividend growth. As reported by Shiller (1981) and Le Roy and Porter (1981), the volatility of realized dividends is much lower than the volatility in prices. In representative agent models with Bayesian updating, the volatility of expected dividends can never exceed realized dividends (West, 1988), whose importance in variance decomposition tests is very small.¹² Therefore, the literature has focused on variation in the stochastic discount factor coming from risk aversion as a source of excess price volatility in economies that allow a representative agent characterization (Campbell and Shiller 1988; Cochrane, 1992).

¹²For a recent digression, see Chen and Zhao (2009).

At the same time, a large body of empirical work in finance suggests that share prices are at best a very noisy predictor of future growth in dividends. As reviewed in Campbell (2003) (see references therein), quarterly real dividend growth and real stock returns for US post-war data have a correlation of only 0.03, which increases to 0.47 at 4-year horizons. This poses another challenge for risk-based explanations of asset price volatility, because it suggests that most of the price volatility results from factors that are largely orthogonal to expected future dividends.

Our theory suggests instead that high return volatility could result from volatile dividend expectations in a Bayesian environment despite low variation in observed dividends, as long as the informational frictions stressed above are severe enough. When noise trading is highly volatile, market information in prices is noisy and traders beliefs remain heterogeneous in equilibrium. With finite precision of private signals, large shifts in noisy demand are then absorbed by large shifts in the identity of the marginal trader, resulting in high price volatility. The ratio between price and realized dividend volatility can be made arbitrarily large by increasing the variance of noise trading shocks. At the same time, the correlation between prices and realized subsequent dividends can be arbitrarily close to zero, which is potentially consistent with the evidence summarized by Campbell (2003). If the price is sufficiently noisy so as to be a poor signal of fundamentals, yet individual traders have sufficiently precise private information, then our model can jointly account for large excess price volatility, and low predictability of future dividends.

Whether heterogeneous expectations can quantitatively account for observed excess price volatility and low predictability of future dividends is an empirical question we do not address here. Rather, the contribution of our model in this respect is to offer a theoretical framework, fully consistent with agent rationality, where this channel is not ruled out by the mere observation that the variability of actual dividends is modest, and not highly correlated with prices.

4 Applications

In this section, we study two applications of our theory. First, we reconsider the Modigliani-Miller Theorem. Second, in a dynamic extension of our model we show conditions under which bubbles arise.

4.1 Splitting Cash-flows to influence market value

The Modigliani-Miller theorem states that in perfect and complete financial markets, splitting a cash flow into two different securities, and selling these claims separately to investors does not influence

its total market value (Modigliani and Miller, 1958). Here we show that with noisy information aggregation, the Modigliani-Miller theorem remains valid only if the different claims are sold to investor pools with *identical* informational characteristics. When the investor pools for different claims have different characteristics, then the nature of the split influences the seller's revenue. The seller in turn can increase her revenues by tailoring the split to the different investor types.

Consider a seller who owns claims on a stochastic dividend $\pi(\cdot)$. This cash flow is divided into two parts, π_1 and π_2 , both monotone in θ , such that $\pi_1 + \pi_2 = \pi$, and then sold to traders in two separate markets. We assume (without loss of generality) that π_2 has more upside risk than π_1 . For each claim, there is a unit measure of informed traders who obtain a noisy private signal $x_i \sim \mathcal{N}(\theta, \beta_i^{-1})$, and a noise trader demand $\Phi(s_i)$, where

$$\begin{pmatrix} s_1 \\ s_2 \end{pmatrix} = \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{u,1}^2 & \rho\sigma_{u,1}\sigma_{u,2} \\ \rho\sigma_{u,1}\sigma_{u,2} & \sigma_{u,2}^2 \end{pmatrix} \right)$$

That is, each market is affected by a noise trader shock s_i with market-specific noise parameter $\sigma_{u,i}^2$. The environment is then characterized by the market-characteristics β_i and $\sigma_{u,i}^2$, and by the correlation of demand shocks across markets, ρ . Traders are active only in their respective market. However, we consider both the possibility that traders observe and condition on prices in the other market (informational linkages), and the possibility that they do not (informational segregation).

Under informational segregation, the analysis of the two markets can be completely separated; any correlation between the two in prices is the result of correlation in demand shocks, as well as the common underlying fundamental, but this doesn't influence expected revenues. The equilibrium characterization from proposition 1 applies separately in each market:

$$P_i(z_i) = \mathbb{E}(\pi_i(\theta) | x = z_i, z_i; \beta_i, \sigma_{u,i}^2) \quad \text{and} \quad V_i(z_i) = \mathbb{E}(\pi_i(\theta) | z_i; \beta_i, \sigma_{u,i}^2).$$

The seller's total expected revenue in excess of the cash flow's expected dividend value is then given by $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2})$, where $\sigma_{P,i}$ is determined as in lemma 2, and denotes the level of informational frictions in each market.

With informational linkages, the equilibrium analysis has to be adjusted to incorporate the information contained in price 1 for the traders in market 2, and vice versa. The characterization proceeds along the same lines as the previous model. Since expected dividends are monotone, informed traders in market i will buy a security if and only if their private signal exceeds a threshold $\hat{x}_i(\cdot)$, where $\hat{x}_i(\cdot)$ is conditioned on both prices. By market-clearing, it must be the case that $\hat{x}_i(\cdot) = z_i \equiv \theta + 1/\sqrt{\beta_i} \cdot s_i$. Observing P_i is then isomorphic to observing z_i , and observing both

prices is isomorphic to observing (z_1, z_2) . We let (z_1, z_2) denote the state, and consider equilibrium price functions $P_1(\cdot)$ and $P_2(\cdot)$ that are measurable w.r.t. (z_1, z_2) . It is then straight-forward to characterize posterior beliefs over θ using Bayes'rule, and to characterize the traders' indifference conditions and hence the market price functions, and expected dividends, conditional on (z_1, z_2) :

$$\begin{aligned} P_1(z_1, z_2) &= \mathbb{E}(\pi_1(\theta)|x = z_1; z_1, z_2) \text{ and } V_1(z_1, z_2) = \mathbb{E}(\pi_1(\theta)|z_1, z_2), \\ P_2(z_1, z_2) &= \mathbb{E}(\pi_2(\theta)|x = z_2; z_1, z_2) \text{ and } V_2(z_1, z_2) = \mathbb{E}(\pi_2(\theta)|z_1, z_2) \end{aligned}$$

In Appendix B we fully characterize expected prices and dividends for this two asset model. In particular, we show the following modified version of lemma 2.

Lemma 3 (Unconditional Wedge with two assets) *For each cash-flow π_i , the unconditional information aggregation wedge W_{π_i} is characterized by*

$$\begin{aligned} W_{\pi_i}(\sigma_{P,i}) &= \int_0^\infty (\pi'_i(\theta) - \pi'_i(-\theta)) \left(\Phi\left(\frac{\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\theta}{\sigma_{P,i}}\right) \right) d\theta, \text{ where} \\ \sigma_{P,i}^2 &= \sigma_\theta^2 + \frac{\beta_i}{(\beta_i + V)^2} (1 + \sigma_{u,i}^2) \text{ and } V = \sigma_\theta^{-2} + \frac{1}{1 - \rho^2} \left(\frac{\beta_1}{\sigma_{u,1}^2} + \frac{\beta_2}{\sigma_{u,2}^2} - 2\rho \frac{\sqrt{\beta_1\beta_2}}{\sigma_{u,1}\sigma_{u,2}} \right) \end{aligned}$$

Therefore, the cases of informational segregation and informational linkages only differ in terms of how our measure of informational frictions $\sigma_{P,i}$ depends on the underlying primitive parameters in each case, but for given values of π_i and $\sigma_{P,i}$, the seller's expected revenue net of expected dividends in both cases is $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2})$. We can now state a first version of the Modigliani-Miller theorem for expected revenues in our model.

Proposition 3 (Modigliani-Miller I) *(i) The cash-flow split does not affect the seller's expected revenue, if and only if the market characteristics are identical: $\sigma_{P,1} = \sigma_{P,2}$.*

(ii) If $\sigma_{P,1} > \sigma_{P,2}$, $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) > W_{\pi_1}(\sigma_{P,2}) + W_{\pi_2}(\sigma_{P,1})$, while if $\sigma_{P,1} < \sigma_{P,2}$, $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) < W_{\pi_1}(\sigma_{P,2}) + W_{\pi_2}(\sigma_{P,1})$.

The key to this proposition is that, for given values of σ_P , the expected information aggregation wedge is additive across cash flows: $W_{\pi_1}(\sigma_P) + W_{\pi_2}(\sigma_P) = W_{\pi_1 + \pi_2}(\sigma_P)$, for any σ_P , π_1 and π_2 . If the two markets have identical characteristics, i.e. $\sigma_{P,1} = \sigma_{P,2}$, only the combined cash flow matters for the total wedge - i.e. the Modigliani-Miller result applies. If on the other hand the two markets have different informational characteristics, then the increasing difference property of $W_{\pi_1}(\cdot)$ implies that the seller's revenue is influenced by how the two cash flows are matched to the two markets, and the revenue is higher when the upside risk is matched with the market that

has more severe information frictions (a higher value of σ_P). Intuitively, the seller exploits the information aggregation wedge to manipulate revenues, matching the pool of investors with high informational frictions with the upside risk, while selling the downside risks to an investor pool with lower informational frictions. This maximizes the gains from the positive wedge resulting on the upside, while it minimizes the losses from the negative wedge on the downside. This logic is pushed further by the next proposition, which considers how the seller can exploit the heterogeneity in investor pools if she gets to design the split of π into π_1 and π_2 .

Proposition 4 (Designing Cash flows) *The seller maximizes her expected revenues by splitting cash flows according to $\pi_1^*(\theta) = \min\{\pi(\theta), \pi(0)\}$ and $\pi_2^*(\theta) = \max\{\pi(\theta) - \pi(0), 0\}$, and then assigning π_1^* to the investor pool with the lower value of σ_P .*

Figure 4: Optimal Cash-flow Design

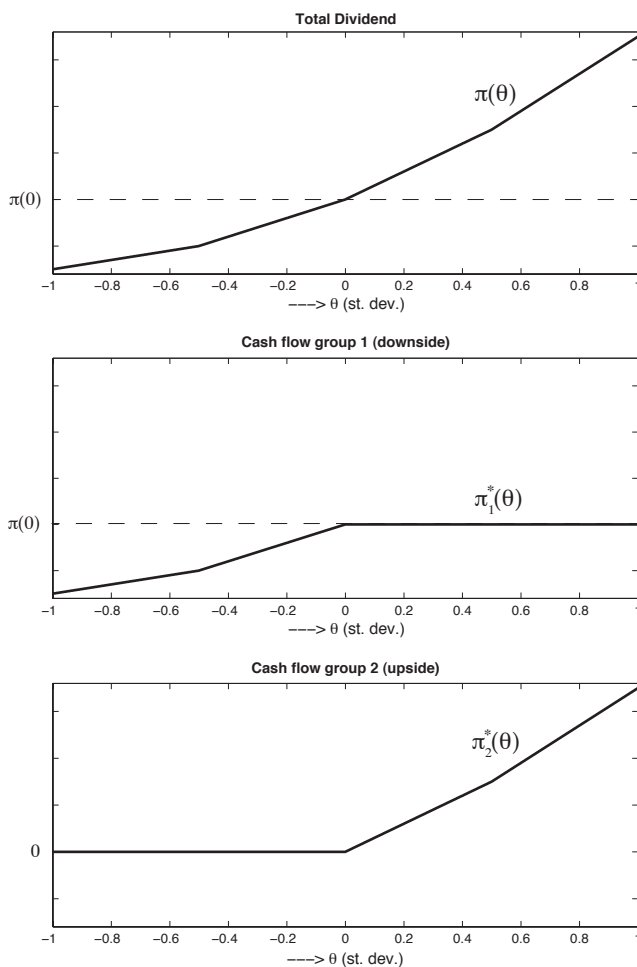


Figure 4 sketches the optimal dividend split for an arbitrary dividend function. The seller maximizes the total proceeds by assigning all the cash flow below the line defined by $\pi(\cdot) = \pi(0)$ to the investor group with the lowest information friction parameter; $\sigma_{P,1}$, and the complement to the investor group with the highest friction; $\sigma_{P,2}$. It is easy to show that any other arbitrary division of cash flows $\{\pi_1(\cdot), \pi_2(\cdot)\}$ implies that both $\pi_1^{*'}(\theta) - \pi_1^{*'}(-\theta) \leq \pi_1'(\theta) - \pi_1'(-\theta)$, and $\pi_2^{*'}(\theta) - \pi_2^{*'}(-\theta) \geq \pi_2'(\theta) - \pi_2'(-\theta)$. That is, π_1 has less downside risk than π_1^* , and π_2 has less upside risk than π_2^* . The increasing difference property of part iv) in Theorem 1 then implies that any transfer of cash flows between investor groups resulting from the alternative split reduces the total proceeds of the issuance. Intuitively, the optimal split loads the entire downside risk on the investor group that discounts the price of the claim the least with respect to its expected payoff (because of the low friction parameter; $\sigma_{P,1}$), while loading the entire upside risk to the group that overvalues the claim the most with respect to its expected dividend (due to the high information frictions; $\sigma_{P,2}$). When $\pi(\cdot) > 0$, this split has a straight-forward interpretation in terms of debt and equity, with a default point on debt that is set at the prior median $\pi(0)$.

An important limitation of the discussion in this section is that we take as given the differences in market characteristics. Moreover, we are implicitly assuming that, given these differences, the seller can freely assign the cash-flows to these two pools. In practice the situation is of course more complicated, because the investor's incentives to obtain information also depends on the asset risks they face. Analyzing this interplay between investor's information choices and the resulting market characteristics, along with the seller's security design question is clearly beyond the scope of this paper, but an important avenue for further work. The results here are simply intended to highlight the possibility of systematic departures from Modigliani and Miller's (1958) irrelevance result, and show that the information frictions give the owner of a cash flow distinct possibility to manipulate its market value through strategic security design.

We conclude this section by stating a second version of the Modigliani-Miller theorem for realized revenues, $P_{\pi_1}(z_1, z_2) + P_{\pi_2}(z_1, z_2)$. The original Modigliani-Miller theorem holds also at an interim stage conditional on new information, as long as the marginal traders in the two markets hold identical beliefs for each realization (z_1, z_2) .

Proposition 5 (Modigliani-Miller II) (i) *With informational segregation: $P_{\pi_1}(z_1, z_2) + P_{\pi_2}(z_1, z_2) = P_{\pi}(z_1, z_2)$ almost surely, if and only if the noise trading is perfectly correlated across markets ($\rho = 1$), and the two markets have identical informational characteristics ($\beta_1 = \beta_2$ and $\sigma_{u,1}^2 = \sigma_{u,2}^2$).*

(ii) *With informational linkages $P_{\pi_1}(z_1, z_2) + P_{\pi_2}(z_1, z_2) = P_{\pi}(z_1, z_2)$ almost surely, if and only if the noise trading is perfectly correlated across markets ($\rho = 1$), and either $\beta_1 = \beta_2$ and*

$$\sigma_{u,1}^2 = \sigma_{u,2}^2, \text{ or } \beta_1 \sigma_{u,1}^{-2} \neq \beta_2 \sigma_{u,2}^{-2}.$$

The perfect correlation reduces the noise to a single common shock. That this is necessary for the theorem to hold under segregation is immediate. It is also necessary for the case with informational linkages, because of the different weighting between the signals for the marginal traders in the two markets. In addition the signal distributions need to be the same, requiring that $\beta_1 = \beta_2$ and $\sigma_{u,1}^2 = \sigma_{u,2}^2$. Finally, the wedge needs to be the same in the two markets, or $\sigma_{P,1} = \sigma_{P,2}$. Together these conditions imply that the two markets have identical informational characteristics, and the marginal trader therefore holds identical beliefs. In the case with informational linkages, we need to consider the additional possibility that when $\beta_1 \sigma_{u,1}^{-2} \neq \beta_2 \sigma_{u,2}^{-2}$ and $\rho = 1$, the observations of two signals with different precision but perfectly correlated noise enables every trader to perfectly infer θ and u from the two prices regardless of the informational parameters, hence $P_{\pi_1}(z_1, z_2) + P_{\pi_2}(z_1, z_2) = \pi(\theta)$.

An interim version of the theorem therefore requires perfect correlation in the noise in different markets, on top of identical informational characteristics.

4.2 Dynamic Trading and Bubbles

As our second application, we consider a simple dynamic extension of our basic model, and show by means of an example how it can easily result in persistent (or even permanent) over-valuation of securities. Standard arguments imply that no arbitrage and common information rules out the possibility of rational bubbles for a general class of dynamic asset market economies (Tirole, 1982; Santos and Woodford, 1997). This is one of the classic no-arbitrage results: while a buy-and-hold strategy insures that a security can never be worth less than its fundamental value under no-arbitrage, a positive bubble component in the price is consistent with arbitrage by buy-and-sell strategies only if its date zero present value follows a Martingale process. But this is inconsistent with the implication of discounting and the transversality condition, that aggregate wealth and the present discounted value of aggregate consumption has finite present value - unless the bubble component is exactly zero.

Here, we show in a simple dynamic example how heterogeneous information and our limits to arbitrage break exactly this result. It is still true that the anticipation of a higher price in the future leads traders to bid up the price in the current period. In our environment, however, extending the insights of Theorem 1 to a dynamic environment, we show that a positive wedge (on average) is sustainable in the future, and leads to a higher willingness to pay in the current period. If this anticipation of a positive future wedge is sufficiently strong it can more than offset a negative

contribution of current payoffs to the wedge, implying that the security is priced above the present discounted value of future dividends in all periods and states.

We establish this result in a model in which per period cashflows $\pi(\cdot)$ are i.i.d. over time, and the security is infinitely lived. As conditions, we require that $\pi(\cdot)$ is dominated by upside risks, so that on average it is expected to trade at a premium, and bounded below (non-negative), so that there is a bound on the wedge on the downside. Inverting the conditions, we also obtain that a security that is dominated by downside risks and has a uniform bound on the upside may trade permanently at a discount.

Time is discrete and infinite, and in each period a new trading round takes place, with informed traders and noise traders. As before the total asset supply is 1. The asset pays dividends $\pi(\theta_t)$ after the current trading round has taken place (hence θ_t is publicly known before the start of period $t + 1$). For simplicity we assume that the fundamental θ_t and the stochastic demand shock u_t are distributed as specified as in section 2, and i.i.d. over time.¹³ Traders are long-lived and risk-neutral and discount the future at a rate $\delta \in (0, 1)$.

The lack of persistence in the dividend process implies that trading in round t only aggregates information about the current fundamental, but includes the anticipation of future prices. Formally, the payoff to a share bought in period t is $\pi(\theta_t) + \delta P(z_{t+1})$, where $P(z_{t+1})$ is the price in period $t + 1$, contingent on the period $t + 1$ state z_{t+1} . The price then satisfies the following recursive characterization:

$$P_\pi(z_t) = \mathbb{E}(\pi(\theta_t) + \delta P_\pi(z_{t+1}) | x = z_t, z_t), \quad (10)$$

and the expected dividend value of the asset satisfies:

$$V_\pi(z_t) = \mathbb{E}(\pi(\theta_t) + \delta V_\pi(z_{t+1}) | z_t). \quad (11)$$

Using the fact that in the i.i.d. case, the expected future prices and dividend values correspond to the unconditional ones, we have the following characterization of the information aggregation wedge in the dynamic model:

$$\begin{aligned} W_\pi(z_t) &= w_\pi(z_t) + \delta \mathbb{E}(W_\pi(z)) = w_\pi(z_t) + \frac{\delta}{1 - \delta} \mathbb{E}(w_\pi(z)), \text{ where} \\ w_\pi(z_t) &= \mathbb{E}(\pi(\theta_t) | x = z_t, z_t) - \mathbb{E}(\pi(\theta_t) | z_t) \end{aligned} \quad (12)$$

is the wedge resulting from the current period payoffs, and $\mathbb{E}(w_\pi(z))$ its corresponding unconditional expectation. Thus, the information aggregation wedge in the dynamic setting depends on

¹³It is possible but outside the scope of the current paper to extend the analysis to allow for persistent fundamental processes. The i.i.d. case is sufficient to convey the core insights that anticipated future wedges influence the current level of the wedge.

both the wedge resulting from current payoffs, and the expected discounted future wedge. Even when the current wedge is negative (at low realizations of z), the overall wedge may still be positive because traders anticipate higher share prices in the future. The following proposition formalizes this observation.

Proposition 6 (Sustainability of Bubbles) *Suppose that $\pi(\theta)$ is bounded below, increasing, and convex. Then, for any $\sigma_P > \sigma_\theta$, there exists $\hat{\delta} < 1$ such that for all $\delta > \hat{\delta}$, $W(z) > 0$, for all z .*

Proposition 6 shows how claims that have a lower bound on payoffs (for example, requiring them to be non-negative), and that generate a positive unconditional wedge, can be priced in the market at a value exceeding expected dividends at all times and in all states of the world. Symmetrically, a claim whose payoffs are bounded above may be undervalued in all future states. The positive (negative) exponential payoff function from example 1 exactly satisfies the required conditions for a permanent bubble (or discount).

The example illustrates the key forces that are at play to overturn the no-arbitrage argument against bubbles: First, with mean reversion in fundamentals and noise trading (captured by the i.i.d. assumption in shocks), the traders anticipation of future wedges are driven by the unconditional wedge. With upside risks, this is positive. Second, with bounded payoffs, there is a limit to how much the market's expectation of current dividends can be undervalued relative to the objective outsider's expectation. Third, the anticipation of a positive future wedge will dominate a negative current wedge, if traders are sufficiently patient.

This example is of course highly stylized, as a complete and exhaustive discussion of dynamic extensions of our model leads to additional difficulties on its own, which exceed the scope of this paper, and are left to future work. Nevertheless it is suggestive of the types of markets in which information-driven bubbles are likely to emerge, and when they are likely to occur, namely those that represent significant future upside opportunities, and/or markets in which investors face implicit protection against downside risks. Furthermore, such bubbles are more likely to occur in time periods where real interest rates are low.

5 Generalized Model

In this section, we explore the robustness of the information aggregation wedge to changes in the model's core assumptions. We first show that the information aggregation wedge arises as an

equilibrium property under much more general assumptions regarding trader’s risk preferences, alternative distributional assumptions, and importantly the bounds on the traders positions. Second, we modify our model from section 2 to show how the magnitude of the wedge is inversely related to the extent of arbitrage activity by risk-neutral, uninformed traders.

5.1 Risk aversion, distributional assumptions, and limits to arbitrage

In this section, we generalize the characterization of the information aggregation wedge to almost completely general specifications of preferences and general distributional assumptions, relaxing in particular the core assumptions of risk-neutrality and on the bounds of trader positions. This result confirms that the wedge is a very general property of asset pricing models with noisy information aggregation, and not due to the specific structure of preferences that we imposed. The formal result we prove is the following: for any noisy rational expectations equilibrium that satisfies a pair of regularity conditions on equilibrium posterior beliefs and demand functions by informed traders, there exists a sufficient statistic z , i.e. a function $z(\theta, u)$, such that the equilibrium price can be represented as a function only of z , or $P(\theta, u) = P(z(\theta, u))$, and it takes the form $P(z) = \mathbb{E}(\pi(\theta)|x = z, z)$, whereas the expected dividend conditional on P takes the form $V(z) = \mathbb{E}(\pi(\theta)|z)$. In other words, the core features of the equilibrium characterization that we were working with throughout this paper turns out to be an extremely general property of price formation with noisy information aggregation.

The asset market structure is the same as in our benchmark model, with the following modifications: (i) θ is distributed according to an arbitrary smooth prior $h(\cdot)$ on \mathbb{R} , (ii) private signals are distributed i.i.d. according to a distribution with cdf $F(\cdot|\theta)$, which satisfies the monotone likelihood ratio property, (iii) The traders’ preferences are characterized by some strictly increasing, concave utility function $U : \mathbb{R} \rightarrow \mathbb{R}$, which is defined on the traders’ realized gains or losses, given by $d(\pi(\theta) - P)$ when they purchase d units at a price P , (iv) traders submit price-contingent orders $d : \mathbb{R}^2 \rightarrow [0, 1]$ from signal-price pairs (x_i, P) into asset holdings $[d_L(P), d_H(P)]$, where $d_L(P) < 0 < d_H(P)$ represent arbitrary continuous price-continent bounds on the traders’ asset positions,¹⁴ and (v) the supply of the asset, net of noise trading is $S(u, P) \in [d_L(P), d_H(P)]$, which is continuous and increasing in P and the supply shock u which is on $[0, 1]$ according to some continuous cdf $G(\cdot)$, and independent of θ or the traders’ private signals.

¹⁴This formulation of the set of possible holdings also allows for situations where traders are still risk-neutral but their positions limited by, say, the wealth they can invest in the risky asset (so that each unit of wealth buys $1/P$ units of the asset).

A Perfect Bayesian Equilibrium is then defined by a price function $P(\theta, u)$, a demand function $d(x, P)$ for informed traders, and posterior beliefs for informed traders given by a cdf. $H(\cdot|x, P)$ and pdf $h(\cdot|x, P)$ such that traders' demand is optimal, given their posteriors conditional on observing x and $P(\theta, u) = P$, $H(\cdot|x, P)$ and $h(\cdot|x, P)$ are consistent with Bayes' rule, and the asset market clears: $S(u, P) = D(\theta, P) \equiv \int d(x, P) dF(x|\theta)$ for all (θ, u) . We now state the following characterization result:

Proposition 7 (Risk aversion) *Suppose $\{P(\theta, u); d(x, P); H(\cdot|x, P)\}$ is a Perfect Bayesian Equilibrium satisfying the following two conditions:*

- (i) $\lim_{x \rightarrow -\infty} \int \pi(\theta) dH(\theta|x, P) < P < \lim_{x \rightarrow +\infty} \int \pi(\theta) dH(\theta|x, P)$, and
- (ii) $\int \pi(\theta) dH(\theta|x, P)$ is Lipschitz continuous in P , with Lipschitz constant $\lambda < 1$.

Then there exists a sufficient statistic function $z(\theta, u)$, along with a conditional cdf $\Psi(z'|\theta) = \Pr(\{u \in [0, 1] : z(\theta, u) \leq z'\})$ and density $\psi(z|\theta)$ such that $P(\theta, u) = P(z(\theta, u))$, where

$$P(z) = \mathbb{E}(\pi(\theta)|x = z, z) = \frac{\int \pi(\theta) f(z|\theta) \psi(z|\theta) h(\theta) d\theta}{\int f(z|\theta) \psi(z|\theta) h(\theta) d\theta},$$

while expected dividends conditional on the price take the form

$$V(z) = \mathbb{E}(\pi(\theta)|z) = \frac{\int \pi(\theta) \psi(z|\theta) h(\theta) d\theta}{\int \psi(z|\theta) h(\theta) d\theta}.$$

Notice that we can construct an arbitrary number of sufficient statistic functions $z(\theta, u) = \hat{z}(P(\theta, u))$ for any strictly monotone function $\hat{z}(P)$. The contribution of proposition 7 therefore consists in constructing a sufficient statistic representation in terms of a random variable z , which exactly inherits and thus generalizes the equilibrium characterization of our baseline model with risk neutrality. The idea behind this representation result is to identify the sufficient statistic with the private signal of the trader who finds it exactly optimal to hold 0 assets. That is, define $z(P)$ implicitly by $d(z, P) = 0$. Condition (i) above implies that such a private signal threshold must exist and hence $z(P)$ is well-defined, while condition (ii) is sufficient to guarantee that it is strictly monotone, and hence observing z is informationally equivalent to observing P .¹⁵ The representation then follows from the observation that a risk averse trader will have an asset demand of zero, if and only if his expectation of dividends equals the price, implying $P = \mathbb{E}(\pi(\theta)|z(P), P)$, from which the representation follows immediately.

The conditions (i) and (ii) impose restrictions on the equilibrium beliefs $H(\cdot|x, P)$, for which our representation result applies. A drawback of our characterization is that these are conditions

¹⁵Our condition (ii) thus plays the same role as the invertibility property of the price function in lemma 1, except that there the invertibility w.r.t. z came immediately as an equilibrium property from the market-clearing condition.

imposed on the endogenous equilibrium beliefs and not on the exogenous primitives of the model. However, they have the advantage of lending themselves to easy interpretation: condition (i) is a restriction on the informativeness of the traders' private signals at extreme values. By log-concavity of the signal density $f(\cdot|\theta)$, the traders' conditional expectations of dividends are always increasing in x . This condition requires that the traders private signals have a sufficiently important impact on the posterior beliefs to ensure that there are always some traders who think the asset's expected return is positive, and some traders who think it's expected return is negative. Condition (ii) considers the impact of P on posterior expectations of dividends, and requires that dividend expectations always increase less than one-for-one in the price. This insures that for any given trader, the expected return from holding the asset is decreasing in its price.¹⁶

We conclude with several remarks. First, we emphasize that this representation extends the basic logic of our main characterization results to a much wider class of preference structures, and suggests that this is a general feature of noisy information aggregation in asset markets. The more delicate issue is to what extent the implications we draw from the wedge (such as the ex ante return and volatility results, and their applications) hold more generally. The real complication with such generalizations is two-fold: first, without any structure on the relevant distributions, such characterizations become a much harder task. They are complicated by the nature of the fixed point characterization. Second, in the initial characterization with risk neutrality, the representation of conditional beliefs and the sufficient statistic z was completely separate from the characterization of prices. As one moves away from these assumptions, the sufficient statistic representation is no longer separable from the characterization of prices and demand behavior, and therefore is no longer invariant to $\pi(\theta)$. This may affect some of the comparative statics results w.r.t. the shape of $\pi(\cdot)$, which took as given the distribution of z .

Second, notice that the only property of demand that we have exploited to arrive at the characterization is that a trader's asset demand is zero when the price equals the trader's dividend expectation. This implies among others that the above method of characterizing the equilibrium is even more general than what is suggested here: for example the same characterization still obtains when there is arbitrary heterogeneity in the shape of U across agents, since the point at which their demand is zero only depends on their expected return from holding the asset, and not on the shape of U . By the same argument the model can also be extended to allow for background risks, provided that there is no correlation between background risk and the asset return.¹⁷ This elements

¹⁶In fact, the formal argument only requires this property at prices that are observed on the equilibrium path.

¹⁷With correlation, there is an additional hedging component to demand which clouds the analysis.

of course alter the characterization of the conditional distribution of z , but not the representation result.

Third, while our proof derives the representation of the price function in terms of the sufficient statistic z by taking as given the traders' equilibrium demand behavior $d(x, P)$, we can also directly characterize the sufficient statistic z and the equilibrium price function and conditional densities as the solution to a fixed point problem. Specifically, a sufficient statistic function $z(\theta, u)$ with conditional cdf $\Psi(z'|\theta) = \Pr(\{u \in [0, 1] : z(\theta, u) \leq z'\})$, a demand function $d(x, P)$, and price function $P(z)$ form a Perfect Bayesian Equilibrium whenever

$$d(x, P(z)) \in \arg \max_{d \in [d_L(P(z)), d_H(P(z))]} \int U(d[\pi(\theta) - P(z)]) h(\theta|x, z) d\theta \text{ for all } (x, z), \quad (13)$$

$$P(z) = \int \pi(\theta) h(\theta|z, z) d\theta \text{ for all } z, \text{ where} \quad (14)$$

$$h(\theta'|x, z) = \frac{f(x|\theta) \psi(z|\theta) h(\theta)}{\int f(x|\theta') \psi(z|\theta') h(\theta') d\theta'}; \text{ and} \quad (15)$$

$$S(u, P(z)) = \int d(x, P(z)) dF(x|\theta) \text{ for all } (\theta, u) \text{ and } z = z(\theta, u). \quad (16)$$

In this characterization, condition (13) ensures that the demand function maximizes the traders' expected utility given any (x, z) . Condition (14) makes sure that the sufficient statistic variable z is consistent with its characterization above. Condition (16) guarantees that the sufficient statistic function $z(\theta, u)$ is consistent with market-clearing, and condition (15) characterizes the resulting conditional beliefs of θ , given (x, z) .¹⁸

Fourth, in the case with neutral preferences and position bounds, a threshold characterization of optimal trader demand applies more generally and provides significantly sharper equilibrium characterizations. By construction, the traders' demand is $d(x, P(z)) = d_L(P(z))$, when $x < z$, and $d(x, P(z)) = d_H(P(z))$, when $x > z$, and therefore total demand is $d_L(P(z)) + (d_H(P(z)) - d_L(P(z)))(1 - F(z|\theta))$. Solving the market-clearing condition then yields

$$F(z|\theta) = \frac{d_H(P(z)) - S(u, P(z))}{d_H(P(z)) - d_L(P(z))},$$

which has to hold for any (θ, u) . In this case, since the sufficient statistic z also corresponds to the signal threshold at which the informed traders switch their position, by the same argument as lemma 1 an equilibrium requires $P(z)$ to be invertible (i.e. the additional conditions we required

¹⁸Our solution approach here is related to Breon-Drish (2011), who relaxes dividend normality in a Grossman Stiglitz economy. He solves the fixed point of a (potentially non-linear) price function by imposing optimal trader behavior in the market clearing condition. Proposition 7 ensures an equilibrium characterization through the sufficient statistic z for more general preferences and distributions, as long as the equilibria satisfies our regularity conditions.

in proposition 7 to guarantee invertibility are not needed). What's more, when d_H , d_L and S do not depend on price, or even when the above ratio in the market-clearing condition doesn't depend on P , the sufficient statistic z is uniquely defined by the market-clearing condition, regardless of the dividend structure. In this case, the equilibrium is uniquely defined by the resulting cdf for the sufficient statistic: $\Psi(z|\theta) = G(\hat{u}(z, \theta))$, where $\hat{u}(z, \theta)$ is defined by $S(\hat{u}(z, \theta)) = d_H - (d_H - d_L)F(z|\theta)$. In addition, we observe that setting the position bounds to $[d_L, d_H] = [0, 1]$ amount to nothing more than a re-normalization of the noise-trading distribution, such that a wider band in positions (as measured by a higher value of $d_H - d_L$) is equivalent to a reduction of the variance in supply shocks. In other words, a model with position bounds $[d_L(P), d_H(P)]$ and a supply of assets $S(u, P)$ is thus equivalent in terms of prices to a model with position bounds normalized to $[0, 1]$ along with a normalized supply function given by

$$\hat{S}(u, P) = \frac{S(u, P) - d_L(P)}{d_H(P) - d_L(P)}.$$

The closed-form solutions and unique equilibrium characterization that we obtained in section 2 thus generalize to any model in which traders are risk-neutral and position limits are constant.

Finally, notice that risk aversion does not appear directly in this representation. In fact, risk premia are accounted for in the threshold level of the private signal that is required to make it optimal for a trader to hold zero assets. If the asset is in positive net supply, and informed traders on average have to hold a positive position in equilibrium, then the trader with position zero will have a private signal z that is lower than the average private signal in the market, and thus a lower expectation $\mathbb{E}(\pi(\theta)|x = z, z)$. This difference between the average private signal and the signal of the person whose dividend expectation equals the price thus just corresponds to the risk premium.

We can illustrate this observation about risk premia in the context of the CARA-normal model, which is a special case of the general structure defined above. Specifically, suppose that (i) traders start with initial share-holdings of zero, (ii) their preferences over terminal wealth are $U(w) = -\exp(-\gamma w)$, (iii) the dividend is normally distributed, $\pi(\theta) = \theta$, with $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$; and (iv) traders do not face limits on their portfolio holdings, and (v) the supply of shares is stochastic and normally distributed, according to $u \sim \mathcal{N}(\bar{u}, \sigma_u^2)$. We conjecture and verify that the equilibrium is characterized by a sufficient statistic $z(P)$ which is distributed according to $z(P) \sim \mathcal{N}(\theta - R, \tau_P^{-1})$, where τ_P denotes the informativeness of the price signal, and R will measure the risk premium. The informed traders' posterior conditional on $z(P)$ and their private signal x is normal, and their

optimal demand satisfies

$$d(x, P) = \frac{\mathbb{E}(\theta|x, z(P)) - P}{\gamma \mathbb{V}(\theta|x, z(P))},$$

where $\mathbb{E}(\theta|x, z(P)) = \frac{\beta x + \tau_P(z(P) + R)}{\sigma_\theta^{-2} + \beta + \tau_P}$ and $\mathbb{V}(\theta|x, z(P)) = (\sigma_\theta^{-2} + \beta + \tau_P)^{-1}$.

Aggregating demand across traders and using the market-clearing condition $\int d(x, P) d\Phi(\sqrt{\beta}(x - \theta)) = u$, the equilibrium price satisfies:

$$P = \frac{\beta\theta + \tau_P(z(P) + R)}{\sigma_\theta^{-2} + \beta + \tau_P} - \frac{\gamma}{\sigma_\theta^{-2} + \beta + \tau_P} u.$$

The equilibrium price as a function of the sufficient statistic $z(P)$ thus satisfies the representation $P(z) = \mathbb{E}(\theta|x = z, z)$, when $z(P)$ is defined as $z(P) = \theta - \gamma/\beta \cdot u$. This confirms the conjecture with $R = \gamma/\beta \cdot \bar{u}$ and $\tau_P = (\beta/\gamma)^2 \cdot \sigma_u^{-2}$. On the other hand, the expected dividend conditional on z is $\mathbb{E}(\theta|z) = \tau_P/(\sigma_\theta^{-2} + \tau_P) \cdot (z + R)$. Therefore, the wedge is

$$W(z) = \left(\frac{\beta + \tau_P}{\sigma_\theta^{-2} + \beta + \tau_P} - \frac{\tau_P}{\sigma_\theta^{-2} + \tau_P} \right) (z + \gamma/\beta \cdot \bar{u}) - \frac{\gamma}{\sigma_\theta^{-2} + \beta + \tau_P} \cdot \bar{u}$$

and thus consists of a constant term that measures the risk premium, and a linear term that corresponds to the over-reaction of the price to the information it contains. The linear equilibrium of the CARA-normal model thus embeds both a risk-premium and the over-reaction to the market-information. Remarkably, the only prior discussion of this observation that we are aware of is by Vives (2008), which briefly mentions the comparison between prices and expected dividends for the CARA case. The restrictions of the CARA model to normally distributed asset payoffs on the other hand limit the analysis to the symmetric case in which any ex ante wedge is attributed to risk premia. This precludes the results linking the shape of asset payoffs to their expected prices and dividends.

5.2 Price impact of information

We now generalize our previous formulation to allow for a response of uninformed traders to perceived excess returns on the asset, as well as stochastic trading motives which are unrelated to dividend expectations (for example, liquidity or hedging needs). We keep the same model as in section 2, but consider the following formulation for asset demand:

$$D(u, P) = \Phi(u + \eta(\mathbb{E}(\pi(\theta)|P) - P)), \quad (17)$$

with $u \sim \mathcal{N}(0, \sigma_u^2)$. Uninformed traders' demand is increasing in the expected return conditional on the price, $\mathbb{E}(\pi(\theta) | P) - P$, with an elasticity given by η .¹⁹ The parameter η captures the responsiveness of uninformed traders to the expectation of dividends in excess of prices, or in other words, the extent to which they are willing or able to arbitrage away the difference between expected price and dividend value. Equivalently, η measures the price impact of private information which relates naturally to the concept of market liquidity.

We follow our previous equilibrium characterization and asset prices with minor changes to account for the endogeneity of demand to asset prices. Market-clearing implies $\Phi(\sqrt{\beta}(\hat{x}(P) - \theta)) = \Phi(u + \eta(\mathbb{E}(\pi(\theta) | P) - P))$, or

$$z = \theta + \frac{1}{\sqrt{\beta}}u = \hat{x}(P) - \eta/\sqrt{\beta} \cdot (\mathbb{E}(\pi(\theta) | P) - P). \quad (18)$$

Observing P is thus isomorphic to observing $z \sim \mathcal{N}(\theta, \sigma_u^{-2}/\beta)$, and Lemma 1 continues to hold without any changes. Using the fact that the expected dividend is $\mathbb{E}(\pi(\theta) | P) = \mathbb{E}(\pi(\theta) | z) = V(z)$, the equilibrium price function is implicitly defined by the marginal trader's indifference condition

$$P(z) = \mathbb{E}(\pi(\theta) | \hat{x}(P), z) = \mathbb{E}\left(\pi(\theta) | z + \eta/\sqrt{\beta} \cdot (V(z) - P(z)), z\right). \quad (19)$$

This condition implicitly defines the equilibrium price. Let $P(z; \eta)$ denote the equilibrium price as a function of the elasticity parameter η , and $P(z; 0) = P(z)$ the price function with inelastic supply. The next proposition shows that the magnitude of the information aggregation wedge is inversely related to the uninformed trader's demand elasticity.

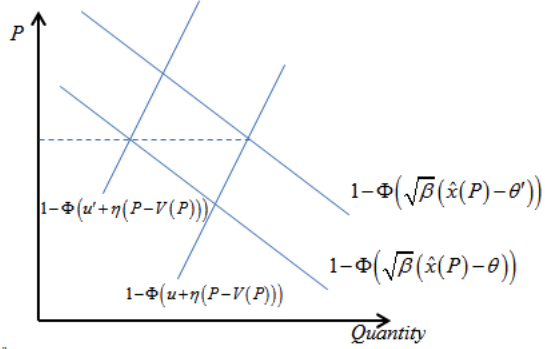
Proposition 8 (Price-elastic demand) *If $P(z; 0) = V(z)$, then $P(z; \eta) = V(z)$, for all η . If $P(z; 0) \neq V(z)$, then $|P(z; \eta) - V(z)|$ is strictly decreasing in η and $\lim_{\eta \rightarrow \infty} |P(z; \eta) - V(z)| = 0$.*

Therefore, the more elastically the uninformed and noise traders respond to the wedge between prices and expected dividends, the more they arbitrage away this difference, and the smaller the information aggregation wedge becomes. This is illustrated by figure 5, which illustrates (in comparison to figure 1) how the price impact of a shift in θ is muted by the price elasticity of uninformed traders.

The wedge results from the informed traders' impact on equilibrium prices. The more traders move prices by acting on their private information, the larger is the wedge. In the inelastic case, the wedge was maximized as the informed traders fully determined prices. In the other extreme,

¹⁹Exactly the same analysis can be conducted if uninformed traders responded to the expected return $\mathbb{E}(\pi(\theta) | P) / P$ (provided the latter is well-defined, i.e. $\pi(\theta)$ is always non-negative), instead of the payoff difference $\mathbb{E}(\pi(\theta) | P) - P$.

Figure 5: Identity Shift with Elastic Supply



their price impact vanishes in the limit as $\eta \rightarrow \infty$ and the uninformed traders completely arbitrage the wedge. The parameter η can thus also be intuitively interpreted as a measure of the limits to arbitrage by uninformed, risk neutral outsiders.

We can illustrate these effects simply in the example with linear dividends: $\pi(\theta) = \theta$. In this case, the expected dividend value is $V(z) = \gamma_V \cdot z$, as before. The price however solves

$$P(z) = \gamma_P z + \frac{\eta\sqrt{\beta}}{\sigma_\theta^2(1 - \gamma_P)} (V(z) - P(z)) = \left(\gamma_P - (\gamma_P - \gamma_V) \frac{\eta\sqrt{\beta}}{\sigma_\theta^2(1 - \gamma_P) + \eta\sqrt{\beta}} \right) z.$$

Compared to the case with inelastic demand, the over-reaction is smaller, i.e. the coefficient in front of z is decreasing in η , and converges to 1 as $\eta \rightarrow \infty$. The wedge

$$W(z) = (\gamma_P - \gamma_V) \frac{\sigma_\theta^2(1 - \gamma_P)}{\sigma_\theta^2(1 - \gamma_P) + \eta\sqrt{\beta}} z$$

is also decreasing in η and vanishes as $\eta \rightarrow \infty$. The information aggregation wedge is therefore largest when uninformed traders are not actively arbitraging the expected return difference coming from the information aggregation wedge.

6 Concluding Remarks

In this paper we have presented a theory of asset price formation based on heterogeneous information and limits to arbitrage. This theory ties expected asset returns to properties of their risk profile, and generates a channel for excess price volatility. The theory is parsimonious, in the sense that all its results follow directly from the interplay between heterogeneous information and limits to arbitrage. The theory is general, in the sense that we do not impose any strong restrictions on the distribution of asset payoffs for the purpose of tractability (although we do impose such restrictions on information, risk preferences and noise-trading assumptions), but rather aim to identify

the relevant underlying features of cash flows at a general level. And the theory is tractable and lends itself easily to applications, as suggested by our discussion of the Modigliani-Miller theorem and the sustainability of bubbles.

We conclude with short remarks on future potential research directions and our related research. An important avenue for future work is to extend our analysis from section 5 to offer richer insights into the interaction between risk aversion and dispersed information for asset prices. A second direction, which we explore in one-going related research (Albagli, Hellwig, and Tsyvinski, 2011a), is to incorporate the release of public news and disclosures into our asset pricing model, and explore both positive and normative implications of public information and disclosure rules for asset prices. A third important extension is to extend the analysis of a multi-period, and multi-asset extensions of our market model, both of which have already been touched upon in this paper in the context of specific examples. A final important direction lies in the integration of financial market frictions with real decisions that endogenize the dividend payoff function we considered here. In a companion paper (Albagli, Hellwig, and Tsyvinski, 2011b), we consider one such model in which there is interplay between information aggregation, firm decisions and managerial incentives in a simple model of informational feedback.

References

- [1] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2011a). “Public Disclosures and Asset Price Fluctuations,” Work in progress.
- [2] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2011b). “Information Aggregation, Investments and Managerial Incentives,” NBER Working paper No. 17330.
- [3] Allen, Franklin, Stephen Morris, and Andrew Postlewaite (1993). “Finite Bubbles with Short Sale Constraints and Asymmetric Information,” *Journal of Economic Theory*, 61, 206-229.
- [4] Allen, Franklin, Stephen Morris, and Hyun Song Shin (2006). “Beauty Contests and Iterated Expectations in Asset Markets,” *Review of Financial Studies*, 19, 161-177.
- [5] Bacchetta, Philippe and Eric van Wincoop (2006). “Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?,” *American Economic Review*, 96 (3), 552-576.
- [6] Blackwell, David (1951). “Comparison of Experiments,” in: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, 93-102.

- [7] Blackwell, David (1953). "Equivalent Comparison of Experiments," *Annals of Mathematical Statistics*, 24, 265-272.
- [8] Breon-Drish, Bradyn (2011). "Asymmetric Information in Financial Markets: Anything Goes," working paper, Stanford University.
- [9] Brunnermeier, Markus. (2001). *Asset Pricing under Asymmetric Information - Bubbles, Crashes, Technical Analysis and Herding*. Oxford University Press.
- [10] Campbell, John and Robert J. Shiller (1988). "The Dividend-price Ratio and Expectations of Future Dividends and Discount Factors," *The Review of Financial Studies*, 1 (3), 195-228.
- [11] Campbell, John (2003). "Consumption-based asset pricing," in: Constantinides, G., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance* (Chapter 13), vol. IB. North-Holland, Amsterdam, pp. 803-887.
- [12] Chen, Long and Xinlei Zhao (2009). "Return Decomposition," *Review of Financial Studies*, 22 (12), 5213-5249.
- [13] Chen, Joseph, Harrison Hong and Jeremy Stein (2002). "Breadth of Ownership and Stock Returns," *Journal of Financial Economics*, 66 (2-3), 171-205.
- [14] Cochrane, John (1992). "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 2, 463-85.
- [15] Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998), "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 53 (6), 1839-1885.
- [16] Diamond, Douglas and Robert Verrecchia (1981). "Information Aggregation in a Noisy Rational Expectations Economy," *Journal of Financial Economics*, 9 (3), 221-235.
- [17] Gromb, Denis, and Dimitri Vayanos (2010), "Limits of Arbitrage," *Annual Review of Financial Economics*, 2, 251-275.
- [18] Grossman, Sanford J. and Joseph E. Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70 (3), 393-408.
- [19] Harrison, Michael and David Kreps (1978). "Speculative Investors in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92 (2), 323-336.

- [20] Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski (2006). "Self-Fulfilling Currency Crises: The Role of Interest Rates," *American Economic Review*, 96 (5), 1769-1787.
- [21] Hellwig, Martin (1980). "On the aggregation of Information in Competitive Markets," *Journal of Economic Theory*, 22. 477-498.
- [22] Hong, Harrison, and David Sraer, (2011). "Quiet Bubbles," working paper, Princeton University.
- [23] Hong, Harrison and Jeremy Stein, (2007). "Disagreement and the Stock Market," *Journal of Economic Perspectives*, 21 (2), 109-128.
- [24] Kyle, Albert (1985). "Continuous Auctions and Insider Trading," *Econometrica*, 53 (6), 1315-1335.
- [25] LeRoy, Stephen and Richard Porter (1981). "The Present-Value Relation: Tests Based on Implied Variance Bounds," *Econometrica*, 49 (3), 555-574.
- [26] Milgrom, Paul (1981). "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49 (4), 921-943.
- [27] Modigliani, Franco and Merton H. Miller (1958). "The Cost of Capital, Corporate Finance and the Theory of Investment," *American Economic Review*, 48 (3), 261-297.
- [28] Odean, Terrance (1998). "Volatility, Price and Profit when all Traders are above Average," *Journal of Finance*, 53 (6), 1887-1934.
- [29] Santos, Manuel, and Michael Woodford (1997). "Rational Asset Pricing Bubbles," *Econometrica* 65 (1), 19-57.
- [30] Scheinkman, Jose and Wei Xiong (2003). "Overconfidence and Speculative Bubbles" *Journal of Political Economy*, 111 (6), 1183-1220.
- [31] Shiller, Robert (1981). "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71 (3), 421- 436.
- [32] Spiegel, Matthew (1998). "Stock Price Volatility in a Multiple Security Overlapping Generations Model," *Review of Financial Studies*, 11 (2), 419-447.
- [33] Tirole, Jean (1982). "On the Possibility of Speculation Under Rational Expectations," *Econometrica*, 50 (5), 1163-1181.

- [34] Veldkamp, Laura (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.
- [35] Vives, Xavier (2008). *Information and Learning in Markets*. Princeton University Press.
- [36] Watanabe, Masahiro (2008). “Price Volatility and Investor Behavior in an Overlapping Generations Model with Information Asymmetry,” *Journal of Finance*, 63 (1), 229-272.
- [37] West, Kenneth (1988). “Dividend Innovations and Stock Price Volatility,” *Econometrica*, 56 (1), 37-61.

7 Appendix: Proofs

Proof of Lemma 1. Part (i): By market-clearing, $z = \hat{x}(P(z))$ and $\hat{x}(P(z')) = z'$, and therefore $z = z'$ if and only if $P(z) = P(z')$.

Part (ii): Since $P(z)$ is invertible, observing P is equivalent to observing $z = \hat{x}(P(z))$ in equilibrium. But $z|\theta \sim \mathcal{N}(\theta, \sigma_u^2/\beta)$, from which the characterization of $H(\cdot|x, P)$ follows immediately from Bayes' Law. ■

Proof of Proposition 1. Substituting $\hat{x}(P) = z$, a price function $P(z)$ is part of an equilibrium if and only if it satisfies (6) and is invertible. $\pi(\cdot)$ is strictly increasing, and an increase in z represents a first-order stochastic shift in the posterior over θ , so the price function $P_\pi(z)$ is continuous and monotone over its domain and spans its entire range, hence invertible. Moreover, all prices are observed in equilibrium (and hence out-of-equilibrium beliefs play no role). Thus, $P_\pi(z)$ defines the unique equilibrium in which prices are conditioned only on z . ■

Proof of Lemma 2. By the law of iterated expectations, $\mathbb{E}(V(z)) = \mathbb{E}(\pi(\theta)) = \int_{-\infty}^{\infty} \pi(\theta) d\Phi(\theta/\sigma_\theta)$. To find $\mathbb{E}(P(z))$, define $\sigma_P^2 = \sigma_\theta^2(1 + (\gamma_P/\gamma_V - 1)\gamma_P)$. Simple algebra shows that

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{1 - \gamma_P\sigma_\theta}} \phi\left(\frac{\theta - \gamma_P z}{\sqrt{1 - \gamma_P\sigma_\theta}}\right) d\Phi\left(\frac{\sqrt{\gamma_V}z}{\sigma_\theta}\right) = \frac{1}{\sigma_P} \phi\left(\frac{\theta}{\sigma_P}\right).$$

With this, we compute $\mathbb{E}(P(z))$:

$$\begin{aligned} \mathbb{E}(P(z)) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \pi(\theta) d\Phi\left(\frac{\theta - \gamma_P z}{\sqrt{1 - \gamma_P\sigma_\theta}}\right) d\Phi\left(\frac{\sqrt{\gamma_V}z}{\sigma_\theta}\right) \\ &= \int_{-\infty}^{\infty} \pi(\theta) \int_{-\infty}^{\infty} \frac{1}{\sqrt{1 - \gamma_P\sigma_\theta}} \phi\left(\frac{\theta - \gamma_P z}{\sqrt{1 - \gamma_P\sigma_\theta}}\right) d\Phi\left(\frac{\sqrt{\gamma_V}z}{\sigma_\theta}\right) d\theta \\ &= \int_{-\infty}^{\infty} \pi(\theta) \frac{1}{\sigma_P} \phi\left(\frac{\theta}{\sigma_P}\right) d\theta. \end{aligned}$$

Therefore, W_π is

$$\begin{aligned}
W_\pi &= \int_{-\infty}^{\infty} \pi(\theta) \left(\frac{1}{\sigma_P} \phi\left(\frac{\theta}{\sigma_P}\right) - \frac{1}{\sigma_\theta} \phi\left(\frac{\theta}{\sigma_\theta}\right) \right) d\theta \\
&= \int_{-\infty}^{\infty} \pi'(\theta) \left(\Phi\left(\frac{\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\theta}{\sigma_P}\right) \right) d\theta \\
&= \int_0^{\infty} (\pi'(\theta) - \pi'(-\theta)) \left(\Phi\left(\frac{\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\theta}{\sigma_P}\right) \right) d\theta,
\end{aligned}$$

where the first equality proceeds by integration by parts, the second by a change in variables, and the third step uses the symmetry of the normal distribution ($\Phi(-x) = 1 - \Phi(x)$). ■

Proof of Theorem 1. Parts (i) – (iii) follow immediately from lemma 2, the definition of upside and downside risk, and the fact that $\Phi(\theta/\sigma_\theta) > \Phi(\theta/\sigma_P)$ for all θ (since $\sigma_P > \sigma_\theta$). For part (iv) notice that

$$\begin{aligned}
W_{\pi_1}(\sigma_P) - W_{\pi_2}(\sigma_P) &= \int_0^{\infty} \Delta(\theta) \left(\Phi\left(\frac{\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\theta}{\sigma_P}\right) \right) d\theta, \\
\text{where } \Delta(\theta) &= \pi_1'(\theta) - \pi_1'(-\theta) - (\pi_2'(\theta) - \pi_2'(-\theta)).
\end{aligned}$$

Since π_1 has more upside risk than π_2 , $\Delta(\theta) \geq 0$ for all θ , which implies that $W_{\pi_1}(\sigma_P) - W_{\pi_2}(\sigma_P)$ is increasing in σ_P . ■

Proof of Theorem 2. Part (i): To compare the volatility of prices with that of expected dividends, we write $\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2\right)$ as

$$\begin{aligned}
&\mathbb{E}\left((P_\pi(z) - P_\pi(0))^2 \mid z \sim \mathcal{N}(\theta, \sigma_u^2/\beta)\right) - \mathbb{E}\left((P_\pi(\hat{z}) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right) \\
&+ \mathbb{E}\left((P_\pi(z) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right)
\end{aligned}$$

where we have just made explicit the distribution of z conditional on θ , and we have added and subtracted the term $\mathbb{E}\left((P_\pi(\hat{z}) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right)$. This term evaluates the variability of prices under a counter-factual distribution of the signal, such that $P_\pi(\hat{z})$ can be interpreted as a posterior expectation of π conditional on \hat{z} . The various comparisons are now based on (i) evaluating the difference in the first line, which we will label the amplification term, and (ii) comparing $\mathbb{E}\left((P_\pi(\hat{z}) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})\right)$ with $\mathbb{E}\left((V_\pi(z) - V_\pi(0))^2\right)$ and $\mathbb{E}\left((\pi(\theta) - \pi(0))^2\right)$.

For the amplification term, we have

$$\begin{aligned}
& \mathbb{E} \left((P_\pi(z) - P_\pi(0))^2 \mid z \sim \mathcal{N}(\theta, \sigma_u^2/\beta) \right) - \mathbb{E} \left((P_\pi(\hat{z}) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1}) \right) \\
&= \int (P_\pi(z) - P_\pi(0))^2 \left[\frac{\sqrt{\gamma_V}}{\sigma_\theta} \phi \left(\frac{\sqrt{\gamma_V} z}{\sigma_\theta} \right) - \frac{\sqrt{\gamma_P}}{\sigma_\theta} \phi \left(\frac{\sqrt{\gamma_P} z}{\sigma_\theta} \right) \right] dz \\
&= \int 2P'_\pi(z) (P_\pi(z) - P_\pi(0)) \left[\Phi \left(\frac{\sqrt{\gamma_P} z}{\sigma_\theta} \right) - \Phi \left(\frac{\sqrt{\gamma_V} z}{\sigma_\theta} \right) \right] dz
\end{aligned}$$

For $z > 0$, $P_\pi(z) > P_\pi(0)$ and $\Phi(\sqrt{\gamma_P}z/\sigma_\theta) > \Phi(\sqrt{\gamma_V}z/\sigma_\theta)$, while for $z < 0$, both the inequalities are reversed. It follows immediately that this integral is always positive.

For the second part, we first break down the three terms into a variance and a bias component:

$$\begin{aligned}
\mathbb{E} \left((\pi(\theta) - \pi(0))^2 \right) &= \text{Var}(\pi(\theta)) + (\mathbb{E}(\pi(\theta)) - \pi(0))^2 \\
\mathbb{E} \left((V_\pi(z) - V_\pi(0))^2 \right) &= \text{Var}(V_\pi(z)) + (\mathbb{E}(V_\pi(z)) - V_\pi(0))^2 \\
\mathbb{E} \left((P_\pi(\hat{z}) - P_\pi(0))^2 \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1}) \right) &= \text{Var}(P_\pi(\hat{z}) \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})) \\
&\quad + (\mathbb{E}(P_\pi(\hat{z}) \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})) - V_\pi(0))^2
\end{aligned}$$

Now, the functions $\pi(\cdot)$, $V_\pi(\cdot)$, $P_\pi(\cdot)$ are all equal to $\mathbb{E}(\pi(\theta) \mid s)$, for different specifications of the conditioning information s : $s \mid \theta = \theta$ for $\pi(\cdot)$, $s \mid \theta \sim \mathcal{N}(\theta, \sigma_u^2/\beta)$ for $V_\pi(\cdot)$, and $s \mid \theta \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$, for $P_\pi(\cdot)$. We can rank these three signal distributions in that $\mathcal{N}(\theta, \sigma_u^2/\beta)$ is a mean-preserving spread over $\mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})$, which is a mean-preserving spread over $s \mid \theta = \theta$. It follows immediately from Blackwell (1951) that

$$\text{Var}(\pi(\theta)) > \text{Var}(P_\pi(\hat{z}) \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1})) > \text{Var}(V_\pi(z)).$$

Moreover, $\mathbb{E}(\pi(\theta)) = \mathbb{E}(V_\pi(z)) = \mathbb{E}(P_\pi(\hat{z}) \mid \hat{z} \sim \mathcal{N}(\theta, (\beta + \beta\sigma_u^{-2})^{-1}))$, by the Law of Iterated Expectations. To compare the bias terms, observe that

$$\begin{aligned}
\int_{-\infty}^{+\infty} \pi(\sqrt{1-\gamma}\sigma_\theta u) \phi(u) du &= \int_0^{+\infty} \left(\pi(\sqrt{1-\gamma}\sigma_\theta u) + \pi(-\sqrt{1-\gamma}\sigma_\theta u) \right) \phi(u) du \\
&= \pi(0) + \int_0^{+\infty} (\pi'(\theta) - \pi'(-\theta)) \left(1 - \Phi \left(\frac{\theta}{\sqrt{1-\gamma}\sigma_\theta} \right) \right) d\theta.
\end{aligned}$$

Applying this formula to $P(0)$ with $\gamma = \gamma_P$, to $V(0)$ with $\gamma = \gamma_V$, to $\mathbb{E}(\pi(\theta))$ with $\gamma = 0$, we find that for upside risks, $\mathbb{E}(\pi(\theta)) > V(0) > P(0) > \pi(0)$, while for downside risks, $\mathbb{E}(\pi(\theta)) < V(0) < P(0) < \pi(0)$. In both cases, $(\mathbb{E}(\pi(\theta)) - \pi(0))^2 > (\mathbb{E}(\pi(\theta)) - P(0))^2 > (\mathbb{E}(\pi(\theta)) - V(0))^2$, which completes the proof for part (i).

Part (ii) Follows immediately from observing that these comparative statics apply separately to each of the terms used in the decomposition in part (i).

Part (iii) Fixing $\gamma_V < 1$, if $\gamma_P \rightarrow 1$ the variance and bias terms for P approach those for $\pi(\cdot)$. We thus wish to show merely that the amplification term doesn't vanish. But this term converges to

$$\int 2P'_\pi(z) (P_\pi(z) - P_\pi(0)) \left[\Phi\left(\frac{z}{\sigma_\theta}\right) - \Phi\left(\frac{\sqrt{\gamma_V}z}{\sigma_\theta}\right) \right] dz > 0.$$

Likewise, fixing $\gamma_P < 1$, as $\gamma_V \rightarrow 0$, the bias and variance terms are fixed, and we therefore consider the amplification term, which converges to

$$\begin{aligned} & \int 2P'_\pi(z) (P_\pi(z) - P_\pi(0)) \left[\Phi\left(\frac{\sqrt{\gamma_P}z}{\sigma_\theta}\right) - \frac{1}{2} \right] dz \\ &= \lim_{z \rightarrow \infty} \frac{1}{2} (P_\pi(z) - P_\pi(0))^2 + \lim_{z \rightarrow -\infty} \frac{1}{2} (P_\pi(-z) - P_\pi(0))^2 \\ & \quad - \mathbb{E} \left((P_\pi(z) - P_\pi(0))^2 \mid z \sim \mathcal{N}\left(0, \frac{\sigma_\theta^2}{\gamma_P}\right) \right), \end{aligned}$$

which is strictly positive, and infinite whenever $P_\pi(\cdot)$ (or equivalently $\pi(\cdot)$) is unbounded on at least one side (part iv). ■

Proof of Proposition 2. The covariance of $P_\pi(z)$ with $\pi(\theta)$ satisfies

$$\begin{aligned} |\text{cov}(P_\pi(z), \pi(\theta))| &= |\mathbb{E}((P_\pi(z) - \mathbb{E}(P_\pi(z))) (\pi(\theta) - \mathbb{E}(\pi(\theta))))| \\ &= |\mathbb{E}((P_\pi(z) - \mathbb{E}(P_\pi(z))) (V_\pi(z) - \mathbb{E}(V_\pi(z))))| \leq \sqrt{\text{Var}(P_\pi(z)) \text{Var}(V_\pi(z))}. \end{aligned}$$

Therefore the correlation of $P_\pi(z)$ with $\pi(\theta)$ satisfies

$$|\text{corr}(P_\pi(z), \pi(\theta))| = \frac{|\text{cov}(P_\pi(z), \pi(\theta))|}{\sqrt{\text{Var}(P_\pi(z)) \text{Var}(\pi(\theta))}} \leq \sqrt{\frac{\text{Var}(V_\pi(z))}{\text{Var}(\pi(\theta))}},$$

and the regression beta of $\pi(\theta)$ against $P_\pi(z)$ satisfies

$$\frac{|\text{cov}(P_\pi(z), \pi(\theta))|}{\text{Var}(P_\pi(z))} \leq \sqrt{\frac{\text{Var}(V_\pi(z))}{\text{Var}(P_\pi(z))}}.$$

The result then follows from observing that, for given σ_θ^2 , $\lim_{\gamma_V \rightarrow 0} \text{Var}(V_\pi(z)) = 0$, while $\text{Var}(\pi(\theta))$ remains constant and $\text{Var}(P_\pi(z))$ is bounded away from 0. ■

Proof of Lemma 3. We focus on market 1; the characterization is identical for market 2. The two market signals (z_1, z_2) and θ are jointly normally distributed according to

$$\begin{pmatrix} \theta \\ z_1 \\ z_2 \end{pmatrix} = \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & & \\ \sigma_\theta^2 & \sigma_\theta^2 + \tau_1^{-1} & \\ \sigma_\theta^2 & \sigma_\theta^2 + \rho\sqrt{\tau_1^{-1}\tau_2^{-1}} & \sigma_\theta^2 + \tau_2^{-1} \end{pmatrix} \right),$$

where $\tau_1 = \beta_1/\sigma_{u,1}^2$ and $\tau_2 = \beta_2/\sigma_2^2$ denote the signal's precision levels and ρ the correlation in their errors. Define

$$\Sigma = \begin{pmatrix} \sigma_\theta^2 + \tau_1^{-1} & \sigma_\theta^2 + \rho\sqrt{\tau_1^{-1}\tau_2^{-1}} \\ \sigma_\theta^2 + \rho\sqrt{\tau_1^{-1}\tau_2^{-1}} & \sigma_\theta^2 + \tau_2^{-1} \end{pmatrix} \text{ and } \mathbf{1} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

By Bayes' Rule, $\theta|z_1, z_2 \sim \mathcal{N}(\mu(z_1, z_2), V^{-1})$, where

$$\mu(z_1, z_2) = \sigma_\theta^2 \mathbf{1}' \Sigma^{-1} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \text{ and } V = \{\sigma_\theta^2 - \sigma_\theta^2 \mathbf{1}' \Sigma^{-1} \mathbf{1} \sigma_\theta^2\}^{-1} = \sigma_\theta^{-2} + \frac{\tau_1 + \tau_2 - 2\rho\sqrt{\tau_1\tau_2}}{1 - \rho^2}.$$

If $x \sim \mathcal{N}(\theta, \beta_1^{-1})$ is the traders' private signal distribution, then $\theta|x, z_1, z_2 \sim \mathcal{N}(\hat{\mu}(x, z_1, z_2), (\beta + V)^{-1})$, where $\hat{\mu}(x, z_1, z_2) = (\beta_1 x + V\mu(z_1, z_2)) / (\beta_1 + V)$. Therefore,

$$\hat{\mu}(x = z_1, z_1, z_2) = (\beta_1 + V)^{-1} (\beta' + V\sigma_\theta^2 \mathbf{1}' \Sigma^{-1}) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix},$$

where $\beta' = (\beta_1, 0)$. This fully characterizes the price and expected dividend functions $P_\pi(z_1, z_2) = \mathbb{E}(\pi(\theta)|x = z_1, z_1, z_2)$ and $V_\pi(z_1, z_2) = \mathbb{E}(\pi(\theta)|z_1, z_2)$. From an ex ante perspective, $\hat{\mu}(x = z_1, z_1, z_2) \sim \mathcal{N}(0, \hat{\sigma}_1^2)$, where

$$\begin{aligned} \hat{\sigma}_1^2 &= (\beta_1 + V)^{-1} (\beta' + V\sigma_\theta^2 \mathbf{1}' \Sigma^{-1}) \Sigma (\Sigma^{-1} \mathbf{1} \sigma_\theta^2 V + \beta) (\beta_1 + V)^{-1} \\ &= (\beta_1 + V)^{-2} (\beta_1^2 (\sigma_\theta^2 + \tau_z^{-1}) + 2\beta_1 \sigma_\theta^2 V + V\sigma_\theta^2 \mathbf{1}' \Sigma^{-1} \mathbf{1} \sigma_\theta^2 V) \\ &= (\beta_1 + V)^{-2} (\beta_1^2 (\sigma_\theta^2 + \tau_z^{-1}) + 2\beta_1 \sigma_\theta^2 V + V^2 \sigma_\theta^2 - V) \\ &= \sigma_\theta^2 + \left(\frac{\beta_1}{\beta_1 + V}\right)^2 \tau_z^{-1} - \frac{V}{(\beta_1 + V)^2}, \end{aligned}$$

Therefore, we compute $\sigma_{P,1}^2$ as $\sigma_{P,1}^2 = \hat{\sigma}^2 + (\beta + V)^{-1}$ or

$$\sigma_{P,1}^2 = \sigma_\theta^2 + \left(\frac{\beta_1}{\beta_1 + V}\right)^2 \tau_z^{-1} - \frac{V}{(\beta_1 + V)^2} + (\beta_1 + V)^{-1} = \sigma_\theta^2 \left(1 + \frac{\beta_1 \sigma_\theta^{-2}}{(\beta_1 + V)^2} \frac{\tau_z + \beta_1}{\tau_z}\right).$$

The proof is completed by substituting for τ_1 and τ_2 in the definition of V . ■

Proof of Proposition 3. If $\sigma_{P,1} = \sigma_{P,2} = \sigma_P$, then $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) = W_\pi(\sigma_P)$, and hence the total expected revenue is not affected by the split. If instead $\sigma_{P,1} \neq \sigma_{P,2}$, then by Theorem 1, $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) > W_{\pi_1}(\sigma_{P,2}) + W_{\pi_2}(\sigma_{P,1})$, whenever $\sigma_{P,2} > \sigma_{P,1}$ (since π_2 has more upside risk than π_1). ■

Proof of Proposition 4. For any alternative split (π_1, π_2) , the monotonicity requirements imply that $0 \leq \pi_1'(\theta) = \pi'(\theta) - \pi_2'(\theta) \leq \pi'(\theta)$. This in turn implies that for all $\theta \geq 0$, $\pi_1^{*'}(\theta) - \pi_1^{*'}(-\theta) =$

$-\pi'(-\theta) \leq \pi'_1(\theta) - \pi'_1(-\theta)$ and $\pi_2^{*'}(\theta) - \pi_2^{*'}(-\theta) = \pi'(\theta) \geq \pi'_2(\theta) - \pi'_2(-\theta)$, i.e. π_1 has less downside risk and more upside risk than π_1^* , and π_2 has more downside risk and less upside risk than π_2^* . Moreover,

$$(\pi'_1(\theta) - \pi'_1(-\theta)) + (\pi'_2(\theta) - \pi'_2(-\theta)) = \pi'(\theta) - \pi'(-\theta) = (\pi_1^{*'}(\theta) - \pi_1^{*'}(-\theta)) + (\pi_2^{*'}(\theta) - \pi_2^{*'}(-\theta))$$

But then, the expected revenue of selling π_1 to the investor pool with $\sigma_{P,1}$ and π_2 to the investor pool with $\sigma_{P,2}$ is $W_{\pi_1}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) = W_{\pi}(\sigma_{P,1}) + W_{\pi_2}(\sigma_{P,2}) - W_{\pi_2}(\sigma_{P,1})$, while the expected revenue from selling π_1^* to the investor pool with $\sigma_{P,1}$ and π_2^* to the investor pool with $\sigma_{P,2}$ is $W_{\pi_1^*}(\sigma_{P,1}) + W_{\pi_2^*}(\sigma_{P,2}) = W_{\pi}(\sigma_{P,1}) + W_{\pi_2^*}(\sigma_{P,2}) - W_{\pi_2^*}(\sigma_{P,1})$. The difference in revenues is therefore $W_{\pi_2^*}(\sigma_{P,2}) - W_{\pi_2^*}(\sigma_{P,1}) - (W_{\pi_2}(\sigma_{P,2}) - W_{\pi_2}(\sigma_{P,1}))$, which is positive, since π_2^* contains more upside and less downside risk than π_2 , and $\sigma_{P,2} \geq \sigma_{P,1}$ (Theorem 1, part (iv)). ■

Proof of Proposition 5.

Clearly, $z_1 = z_2 = z$ almost surely if and only if $\rho = 1$. If $\beta_1 = \beta_2$ and $\sigma_{u,1}^2 = \sigma_{u,2}^2$, it then follows that $P_{\pi_1}(z_1, z_2) + P_{\pi_2}(z_1, z_2) = P_{\pi}(z)$, almost surely, if and only if $\rho = 1$. Moreover, it follows from the characterizations of P_{π_1} and P_{π_2} that the price function is no longer additive (even if $\rho = 1$), whenever $\beta_1 \neq \beta_2$ or $\sigma_{u,1}^2 \neq \sigma_{u,2}^2$, unless the markets are informationally linked, and $\beta_1 \sigma_{u,1}^{-2} \neq \beta_2 \sigma_{u,2}^2$. In this last case, we find that signals have different precision, but perfectly correlated errors, so θ and the correlated error can be perfectly inferred from the two signals, i.e. $V \rightarrow \infty$ in the characterization in lemma 3, and the wedge disappears. ■

Proof of Proposition 6. If $\pi(\cdot)$ is convex, then by Theorem 1, for any finite \underline{w} , there exists $\hat{\delta} < 1$, s.t. $\delta > \hat{\delta}$, $\delta \mathbb{E}(w(z)) > -(1 - \delta) \underline{w}$. We therefore need to establish a lower bound for $w(z)$. But if $\pi(\cdot)$ is bounded below, then $\lim_{z \rightarrow -\infty} w(z) = 0$, and $w(z)$ is positive for sufficiently high z , so it is necessarily bounded. ■

Proof of Proposition 7. Define $z(P)$ implicitly by $d(z, P) = 0$. By condition (i) and the continuity of $H(\cdot|x, P)$ w.r.t. x , $z(P)$ is well defined. Next, we show that $z(P)$ is strictly increasing: Notice that $d(x, P)$ satisfies the first-order condition

$$\int [\pi(\theta) - P] U'(d(x, P) [\pi(\theta) - P]) dH(\theta|x, P) = 0$$

or equivalently

$$\text{cov}(\pi(\theta) - P; U'(d(x, P) [\pi(\theta) - P]) | x, P) = -\mathbb{E}(\pi(\theta) - P | x, P) \mathbb{E}(U'(d(x, P) [\pi(\theta) - P]) | x, P).$$

The LHS of this condition is positive/zero/negative, if and only if $d(x, P)$ is negative/zero/positive. The RHS on the other hand is positive/zero/negative if and only if $\int [\pi(\theta) - P] dH(\theta|x, P)$ is negative/zero/positive. It follows that $d(x, P) \stackrel{\geq}{\leq} 0$, if and only if $\int \pi(\theta) dH(\theta|x, P) \stackrel{\geq}{\leq} P$. But then, since $\int \pi(\theta) dH(\theta|x, P) - P$ is strictly increasing in x (by log-concavity) and strictly decreasing in P (by condition (ii)) when $x = z(P)$, it follows immediately that $z(P)$ is strictly increasing in P . Therefore, observing P is informationally equivalent to observing $z(P)$. Moreover, since $d(z(P), P) = 0$ implies $P = \int \pi(\theta) dH(\theta|z(P), P)$, inverting this expression yields $P(z) = \mathbb{E}(\pi(\theta)|x = z, z)$. ■

Proof of Proposition 8.

If $P(z; 0) = V(z)$, then $P(z; \eta) = V(z)$ solves the pricing equation for any $\eta > 0$. If $P(z; 0) \neq V(z)$, then define the function $\mathcal{T}^\eta(P, z)$ as

$$\mathcal{T}^\eta(P, z) = \mathbb{E}\left(\pi(\theta) | z + \eta/\sqrt{\beta} \cdot (V(z) - P), z\right).$$

$\mathcal{T}^\eta(P, z)$ is continuous and decreasing in P , and $\mathcal{T}^\eta(V(z), z) = P(z; 0)$. Moreover, if $V(z) > P(z; 0)$, then $\mathcal{T}^\eta(P(z; 0), z) > P(z; 0)$, $\mathcal{T}^\eta(V(z), z) < V(z)$, and therefore there exists a unique $P(z; \eta) \in (P(z; 0), V(z))$, such that $\mathcal{T}^\eta(P(z; \eta), z) = P(z; \eta)$. If instead $V(z) < P(z; 0)$, then $\mathcal{T}^\eta(P(z; 0), z) < P(z; 0)$, $\mathcal{T}^\eta(V(z), z) > V(z)$, and $\mathcal{T}^\eta(P(z; \eta), z) = P(z; \eta)$ for a unique $P(z; \eta) \in (V(z), P(z; 0))$. Moreover, replacing $P(z; 0)$ with $P(z; \eta')$ for $\eta' < \eta$ in the steps above shows that $|P(z; \eta) - V(z)|$ is strictly decreasing in η . For the limit, notice that since $P(z; \eta)$ is monotone in η and bounded, it must converge to a limit $P(z; \infty) = \lim_{\eta \rightarrow \infty} P(z; \eta)$. If $P(z; \infty) > V(z)$, then $P(z; \infty) = \lim_{\eta \rightarrow \infty} \mathcal{T}^\eta(P(z; \eta), z) = -\infty$, whereas if $P(z; \infty) < V(z)$, then $P(z; \infty) = \lim_{\eta \rightarrow \infty} \mathcal{T}^\eta(P(z; \eta), z) = \infty$, both of which are contradictions. Therefore, we are left with $P(z; \infty) = V(z)$ at the limit. ■