Information Aggregation with Asymmetric Asset Payoffs[*]

Elias Albagli — Central Bank of Chile

Christian Hellwig — Toulouse School of Economics

Aleh Tsyvinski — Yale University

December 20, 2022

Abstract

We study noisy aggregation of dispersed information in financial markets beyond the usual parametric restrictions imposed on preferences, information, and return distributions. This allows a general characterization of asset returns by means of a risk-neutral probability measure that features excess weight on tail risks. Using this characterization, we show that noisy aggregation of dispersed information provides a unified explanation for several prominent cross-sectional return anomalies such as returns to skewness, returns to disagreement and interaction effects between the two. Moreover, this characterization can be linked to observable moments such as forecast dispersion and accuracy, and simple calibrations suggest the model can account for a significant fraction of empirical return anomalies.

[*]We thank Bruno Biais, John Geanakoplos, Narayana Kocherlakota, Felix Kubler, Stephen Morris, Guillaume Plantin, Jeremy Stein, Jean Tirole, Xavier Vives, Olivier Wang, Eric Young, and audiences at numerous conferences and seminars for helpful comments. This paper builds on and supersedes our earlier working paper "A theory of asset prices based on heterogeneous information" (Albagli, Hellwig, Tsyvinski 2011, 2015). Hellwig gratefully acknowledges financial support from the European Research Council (starting grant agreement 263790), the Institut Universitaire de France as well as funding from the French National Research Agency (ANR) under grant ANR-17-EURE-0010 (Investissements d'Avenir).Tsyvinski is grateful to NSF (grant 0551175) for support and EIEF for hospitality.

1 Introduction

Dispersed information and disagreement among investors is an ubiquitous feature of financial markets, and asset prices are often viewed as playing a central role in aggregating such information. We develop a flexible theory in which aggregation of dispersed information emerges as the core force determining asset prices and expected returns, and link these to the distribution of the underlying cash-flow risk and features of markets such as liquidity and investor disagreement. This theory provides a unified explanation for several prominent asset pricing anomalies, such as negative excess return to skewness, and the seemingly contradictory evidence on the impact of investor disagreement on returns in equity and bond markets.

We consider an asset market along the lines of Grossman and Stiglitz (1980), Hellwig (1980), and Diamond and Verrecchia (1981), populated by informed investors who observe a noisy private signal about asset payoffs, and noise traders whose random positions determine the net supply of the asset.[1] In such environment, the market-clearing price serves as an endogenous public signal of cash-flow. In contrast to the existing literature, we don't impose parametric restrictions on the distribution of cash flows, which allows us to derive return implications for a wide range of assets and compare the model-implied returns with their empirical counterparts.

The textbook no-arbitrage paradigm characterizes systematic return differences through a risk-neutral probability measure which summarizes investors expectations and attitudes towards risk. We build on no-arbitrage theory by constructing a risk-neutral probability measure for asset prices with dispersed information, and show that noisy information aggregation leads to *excess weight on tail risks*: the pricing kernel or change in probability measure is U-shaped, overweighting probabilities of both very high and low returns. This leads to price premia that can be interpreted as the value of a mean-preserving spread, whose magnitude scales up with the dispersion of investor expectations. Negative returns to skewness, negative (positive) returns to investor disagreement for positively (negatively) skewed securities, and positive interaction between skewness and investor disagreement emerge as direct corollaries. Importantly, these predictions distinguish noisy information aggregation from average risk premia, for which the pricing kernel is monotone and shifts probability mass from high to low returns. They also distinguish our theory from heterogeneous priors models with short-sales constraints in which disagreement unambiguously raises prices and lowers returns.

The main challenge in characterizing asset prices with noisy information aggregation comes

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

from the difficulty of tractably dealing with the endogeneity of information contained in the price. We address this challenge in three steps. First, we represent the equilibrium price by means of a *sufficient statistic* variable that summarizes the information aggregated through the price. This representation reveals the presence of an *updating wedge*: noisy information aggregation makes the asset price more sensitive to fundamental and noise trading shocks than the corresponding risk-adjusted dividend expectations. Hence the price is higher (lower) than expected dividends whenever the information aggregated through the price is sufficiently (un-)favorable.

Second, we represent the expected return of the asset by means of a risk-neutral probability measure and show that the updating wedge leads to *excess weight on tail risks*. Specifically, the risk-neutral differs from the objective distribution through a shift in the mean –akin to an average risk premium–, and a mean-preserving spread that captures the additional effect of noisy information aggregation. In consequence, securities characterized by upside (downside) risks are priced above (below) their fundamental value. Moreover, the over-pricing of upside or under-pricing of downside risks scales with excess weight on tail risks. Third, and to bring the model to the data, we represent excess weight on tail risks as a function of two sufficient statistics of investor beliefs: forecast dispersion and accuracy. The model attributes almost all the variation in excess weight on tail risks to forecast dispersion, which we can therefore interpret as a natural empirical proxy for excess weight on tail risks.

In section 2, we introduce our general model and illustrate these steps with three examples. First, we revisit the canonical CARA-Normal model and confirm the presence of the updating wedge and excess weight on tail risks property, but note it has no effects on average prices and returns due to the imposed symmetry of cash flows. Our second example replaces normally distributed with binary dividends and thus highlights the interaction between payoff asymmetry and excess weight of tail risks. While these examples illustrate our general insights, they still rely on strong parametric assumptions that limit their usefulness for comparative statics and empirical applications. Our third example assumes traders are risk-neutral but face position limits, which allows us to fully characterize the information content of prices for *arbitrary securities.* Moreover, by abstracting from risk aversion, the risk-neutral model focuses exclusively on the role of noisy information aggregation for asset returns. We show that the difference between the average price and cash flows, or the price premium, (i) is increasing in the degree of upside risk, (ii) is positive and increasing (negative and decreasing) in investor disagreement for securities characterized by upside (downside) risk, and (iii) displays positive interaction effects or increasing differences between payoff asymmetry and investor disagreement.

Section 3 shows how the model predictions and comparative statics are consistent with several empirical regularities. A substantial empirical literature recovers estimates of the pricing kernel by inverting no-arbitrage characterizations of option prices. They consistently find evidence of non-monotonic or U-shaped pricing kernels that seem at odds with standard risk-based no arbitrage conditions, i.e. the so-called *pricing kernel puzzle*, or negative variance risk premia (the so-called *variance premium puzzle*), which imply that the risk-neutral variance is a systematically upwards-biased predictor of the underlying return variance. These observations are all consistent with our core prediction that the risk-neutral measure displays excess weight on tail risks.[2]

Our other predictions then derive from the comparative statics that EWTR imply for the cross-sectional variation in returns.

1. Returns to skewness: A large empirical literature documents a negative relation between skewness of the return distribution and expected returns in equity markets.[3] In bond markets, returns to skewness are reflected in the credit spread puzzle, according to which markets appear to overweight default risks, especially for high-quality investment grade bonds.[4]

2. Returns to disagreement: The empirical evidence on returns to investor disagreement is divided. Several studies find negative returns to disagreement in equity markets, which are typically interpreted in support of heterogeneous priors models with short-sales constraints in which securities are over-priced due to an implicit re-sale option whose value is always increasing with forecast dispersion (Miller, 1977). Others find positive returns to disagreement in bonds markets, and interpret disagreement as a proxy for risk.[5]

3. Interaction effects: several studies find that returns to disagreement interact with returns to asymmetry, such as the value premium for equity, or leverage and default risk for bonds.[6]

Our theory is able to account for all three predictions. In particular, it explains why higher

[2]See Jackwerth (2000), Ait-Sahalia and Lo (2000), Bakshi et al. (2010), Christoffersen et al. (2013) and Audrino et al. (2022) for evidence of non-monotone or U-shaped pricing kernels, and Carr and Wu (2009) for evidence on the variance premium puzzle. These empirical findings are established using index options and market returns, while our predictions speak more naturally to cross-sectional return predictions for individual securities, especially in the risk-neutral variant of our model. While we are not aware of equally sharp empirical results for individual securities, the empirical results offer at least suggestive empirical support for the core forces at play in our model.

[3]See Conrad et al. (2013), Boyer et al. (2010) and Green and Hwang (2012).

[4]See Huang and Huang (2012), Feldhütter and Schaefer (2018) and Bai at al. (2020) for recent contributions.

[5]See Diether et al. (2002), and Gebhardt et al. (2001), and Yu (2011) for returns to disagreement in equity markets, Guntay and Hackbarth (2010) for bond markets and Carlin et al. (2014) for mortgage-backed securities.

[6]Yu (2011) documents that returns to disagreement are increasing with book-to-market ratios, and the value premium is increasing with forecast disagreement. Guntay and Hackbarth (2010) report that disagreement has larger impacts on bond spreads and returns for firms with high leverage and low credit ratings.

disagreement can lead to lower equity returns but higher bond returns by identifying upside vs. downside risk as the key determinant for signing the returns to disagreement.

We then use our model to conduct a hypothetical portfolio sorting exercise. We parametrize asset payoff distributions to match idiosyncratic skewness and volatility metrics reported in the literature,[7] and calibrate key informational parameters to match the observed variation in forecast dispersion in the I.B.E.S. data of analysts' earnings forecasts to impute excess weight on tail risks for a cross-section of equity returns. While not a perfect measure of investor private information (after all, analyst forecasts are publicly observed!), we argue below that the observed forecast dispersion is a reasonable proxy of dispersed information between different investors. The analyst forecast sample suggests excess weight on tail risks is highly skewed, i.e. small for most firms but very significant for those in the top skewness quintile, with excess weight on tail risks twice as large as at the average, and up to nine times as large as for the median firm. Our model can generate roughly 40% of the observed return differential between the highest and lowest skewness quintiles and about 70% of the observed return differential between the highest and lowest disagreement quintile. The model-implied returns also display strong positive interaction effects, with annualized excess monthly returns to skewness varying by about 6% between the highest and lowest disagreement quintiles, and excess returns to disagreement varying by a similar magnitude between the highest and lowest skewness quintiles–roughly 90% of the variation in returns to disagreement and the value premium reported in Yu (2011).

Section 4 generalizes the key steps of our theoretical arguments to generic distributions of asset payoffs, supply shocks and investor preferences. Excess weight on extreme tail risks materializes under a weak condition on the informativeness of the prior in the tails. The richer comparative statics results underlying our three main predictions require the somewhat stronger condition that the implied pricing kernel is log-convex. Log-convexity holds in all canonical examples. More generally, we show that the pricing kernel is log-convex whenever (i) posterior beliefs are "sufficiently well behaved", in the sense that agents update monotonically from signals (posterior beliefs satisfy a monotone likelihood ratio property w.r.t. the endogenous market signal), and the informativeness of signals does not vary too much across states, and (ii) the information contained in the price is well approximated by a noisy affine signal of fundamentals.

Section 5 studies extensions to multiple securities and coexistence of informed and uninformed traders. We highlight conditions under which the existence of a risk-neutral measure with an updating wedge generalizes to multi-asset environments, so that a unique pricing kernel can be

[7] See Boyer et al. (2010).

applied to simultaneously price multiple securities. Our extensions further highlight how dispersed information amplifies skewness premia in asset markets, which is key for generating quantitatively significant returns to skewness in the model.

Our paper contributes to the literature on noisy information aggregation in asset markets by offering a variant of the canonical noisy rational expectations model that dispenses with strong parametric assumptions about asset payoffs. In a similar spirit, Breon-Drish (2015) analyzes non-linear and non-normal variants of the noisy REE framework in the broad exponential family of distributions and CARA preferences.[8] Barlevy and Veronesi (2003), Peress (2004) and Yuan (2005) also study non-linear models of noisy information aggregation with a single asset market. Malamud (2015) and Chabakauri et al. (2021) study information aggregation in non-linear, multi-asset noisy REE models with a rich set of state-contingent securities, exploiting spanning properties of state prices with complete or incomplete markets. In contrast to our work, these papers all impose parametric assumptions on the underlying asset payoffs, probability, information and preference structure to fully characterize the information content of asset prices, rather than identifying properties of asset prices that apply beyond the specifics of their environment, and linking such properties to cross-sectional return anomalies.

Our equilibrium characterization with noisy information aggregation also shares similarities with common value auctions, which are especially pronounced in the case with risk-neutral agents and position bounds.[9] Yet whereas the auctions literature seeks to explore under what conditions prices converge to the true fundamental, we focus on the departures from this competitive limit under information frictions.

2 The general model and three examples

We begin by introducing our general model, including the information structure and the financial market. We then discuss three examples to illustrate the key theoretical ideas of our paper.

The financial market is a Bayesian trading game with a unit measure of informed traders and a single asset whose payoff is given by a strictly increasing function $\pi(\cdot)$ of a stochastic fundamental, θ. Nature draws $\theta \in \mathbb{R}$ according to a prior distribution with cdf. $H(\cdot)$. Each informed investor i then receives a private signal $x_i = \theta + \varepsilon_i$, where ε_i is i.i.d across agents, and distributed according

[8] He further derives powerful results on the incentives for information acquisition, whereas we take the information structure as given.

[9] See Wilson (1977), Milgrom (1979, 1981), Pesendorfer and Swinkels (1997), Kremer (2002) and Perry and Reny (2006) for important contributions to this literature.

to cdf. $F(\cdot)$ and smooth, symmetric density function $f(\cdot)$ with unbounded support. We assume $f'(\cdot)/f(\cdot)$ is strictly decreasing and unbounded above and below.[10] We denote the variance of fundamentals by $\sigma_\theta^2 \equiv Var(\theta)$ and the precision of private signals by $\beta \equiv 1/Var(\varepsilon)$.

Investors' preferences are characterized by a strictly increasing, concave utility function $U(\cdot)$ defined on realized gains or losses $d_i \cdot (\pi(\theta) - P)$, where demand $d_i \in [d_L, d_H]$ is restricted by position limits $d_L < 0 < d_H$.[11] Investors submit price-contingent demand schedules, defined as a mapping $d(x_i, P)$ from signal-price pairs into asset holdings. Aggregate demand is thus given by $D(\theta, P) = \int d(x, P)dF(x - \theta)$, where $F(x - \theta)$ is the cross-sectional cdf. of private signals x_i, conditional on θ.[12] The supply of securities s is stochastic with support $[d_L, d_H]$ and distributed according to cdf $G(\cdot)$. Once investors submit orders, a price $P(\theta, s)$ is selected, which clears the market if and only if, for all $(\theta, s) \in \mathbb{R} \times [d_L, d_H]$,

$$s = D(\theta, P) \equiv \int d(x, P)\, dF(x - \theta). \tag{1}$$

Let $H(\cdot|P)$ denote the posterior cdf of θ, conditional on observing P, and $H(\theta|x, P)$ the investors' posterior conditional on x and P. Given $H(\cdot|x, P)$, a demand function $d(x, P)$ is optimal, if it solves the investors' decision problem $\max_{d\in[d_L,d_H]} \int U(d(\pi(\theta) - P))\, dH(\theta|x, P)$. For $d(x, P) = d \in (d_L, d_H)$, this leads to the first-order condition

$$\int (\pi(\theta) - P) \cdot U'(d(\pi(\theta) - P))dH(\theta|x, P) = 0. \tag{2}$$

A *Perfect Bayesian Equilibrium* consists of a demand function $d(x, P)$, a price function $P(\theta, s)$, and posterior beliefs $H(\cdot|P)$ such that (i) $d(x, P)$ is optimal given $H(\cdot|x, P)$; (ii) $P(\theta, s)$ clears the market; and (iii) $H(\cdot|P)$ satisfies Bayes' rule whenever applicable, i.e., for all P such that $\{(\theta, s) : P(\theta, s) = P\}$ is non-empty. We focus on *price-monotone equilibria* $\{P(\theta, s); d(x, P); H(\cdot|P)\}$ in which $d(x, P)$ is decreasing in P whenever $d(x, P) \in (d_L, d_H)$.[13]

In our model, investors do not observe signals directly about asset payoffs, but rather about a fundamental θ, and the asset payoff is a monotone function of θ.[14] This formulation separates

[10] Monotonicity of $f'(\cdot)/f(\cdot)$ implies signals have log-concave density and satisfy the monotone likelihood ratio property. Unboundedness implies extreme signal realizations induce large updates in posterior beliefs, (almost) regardless of the information contained in other signals.

[11] Position limits may be infinite ($(d_L, d_H) = \mathbb{R}$) if the investors are strictly risk-averse (so that security demands are bounded by risk aversion), but must be finite when investors are risk-neutral. Depending on the context and application, both scenarios may be relevant.

[12] As is common in large anonymous games, we assume that a (Strong) Law of Large Numbers holds to equate aggregate demand to the expectation of individaul demand at a given aggregate state.

[13] Price monotonicity of demand arises automatically if trade takes place through a limit-order book.

[14] If $\pi(\theta) = \theta$, our model reduces to the canonical formulation in which investors observe noisy signals of dividends.

the distribution of asset payoffs from the investors' updating of beliefs, which strikes us as a reasonable approximation of many real world financial markets. For example, equity analysts gather information about a firm's earnings and investment opportunities, which affect dividend payouts to shareholders, while a bond analyst may assess the issuer's solvency which depends on revenues and leverage, among other variables often summarized in a single "distance to default" metric. An option trader will forecast where the underlying is heading. In all these cases, the fundamental about which information is gathered is distinct from the security's payoffs, and the mapping from fundamentals to asset payoffs is typically non-linear. Our model is flexible enough to accommodate any of these possibilities.

2.1 Example 1: the Canonical CARA-normal model

We begin with the textbook CARA-normal model of noisy information to introduce two key ideas that we will generalize throughout the rest of our paper: first, we show that the equilibrium price can be represented by a *sufficient statistic* which is defined as a linear combination of fundamental and supply shocks. Furthermore, the price displays an *information updating wedge*, which makes it more sensitive to this sufficient statistic than is warranted by the information it conveys about the fundamental. Second, we show that the updating wedge leads to a risk-neutral probability measure that displays *excess weight on tail risks* (EWTR, henceforth). Moreover, we illustrate that EWTR arises from the presence of supply shocks, and that it gets amplified through the presence of noisy private information. These properties are of limited interest when payoffs are symmetric, but they turn out to be quite consequential once the symmetry assumption is relaxed.

Our general model nests the canonical CARA-normal setup with the following assumptions: (i) normally distributed dividends, $\pi(\theta) = \theta$ with $\theta \sim \mathcal{N}\left(0, \sigma_\theta^2\right)$; (ii) normal distribution of supply: $s \sim \mathcal{N}\left(\overline{s}, \sigma_s^2\right)$; (iii) normally distributed private signals, $x_i|\theta \sim \mathcal{N}\left(\theta, \beta^{-1}\right)$; (iv) CARA preferences over terminal wealth, $U(w) = -\exp(-\chi w)$, and (v) no limits on portfolio holdings, $(d_L, d_H) = \mathbb{R}$.

Sufficient statistic representation: Recall that we can represent an informed trader's asset demand in the CARA-normal set-up as $d(x, P) = \left(\chi Var\left(\theta|x, P\right)\right)^{-1}\left(\mathbb{E}\left(\theta|x, P\right) - P\right)$. Now, define $z(P)$ as the private signal of the investor who, at a given price P, finds it optimal to hold exactly $\overline{s}$ units of the asset. If $z(P)$ is monotone in the price, we can think of the information that is contained in P as equivalent to the direct observation of z, i.e. z serves as a sufficient statistic for the information content of the price. Setting $d(z(P), P) = \overline{s}$, we can invert the demand relation to obtain the following *sufficient statistic representation* of the equilibrium price as a function of z:

$$P(z) = \mathbb{E}\left(\theta|x = z, z\right) - \chi Var\left(\theta|x = z, z\right) \cdot \overline{s}. \tag{3}$$

It is then straight-forward to check that $z|\theta$ is normally distributed with mean θ and variance $\tau^{-1} \equiv (\chi/\beta)^2 \cdot \sigma_s^2$, and therefore posterior beliefs are normal with $\mathbb{E}(\theta|x,z) = \frac{\beta x + \tau z}{1/\sigma_\theta^2 + \beta + \tau}$ and $Var(\theta|x,z) = \left(1/\sigma_\theta^2 + \beta + \tau\right)^{-1}$. Therefore, $\mathbb{E}(\theta|x=z,z) = \hat{\gamma} \cdot z$ and $Var(\theta|x=z,z) = (1-\hat{\gamma})\,\sigma_\theta^2$, where $\hat{\gamma} \equiv \frac{\beta+\tau}{1/\sigma_\theta^2+\beta+\tau}$ denotes the slope coefficient of $P(z)$ with respect to z.

In other words, the equilibrium price is represented as the risk-adjusted dividend expectation of a trader who, at the given price, chooses to hold exactly $\overline{s}$ units of the security. To understand the significance of this representation, we compare $P(z)$ to the "objective" Bayesian posterior of θ, given z, which is also normal, but with expectation $\mathbb{E}(\theta|z) = \gamma \cdot z$ and variance $Var(\theta|z) = (1-\gamma)\,\sigma_\theta^2$, where $\gamma = \frac{\tau}{1/\sigma_\theta^2+\tau}$ represents the slope coefficient of $\mathbb{E}(\theta|z)$ w.r.t. z. Since $\hat{\gamma} > \gamma$, the price differs from the expected dividend by responding more strongly to the market signal z than is justified purely by its information content: $\mathbb{E}(\theta|x=z,z)$ treats the signal as if it had precision $\beta+\tau$, when its true precision instead is only equal to τ. We term this excess price sensitivity the *information updating wedge*.[15] In addition, the price incorporates an expected risk premium $\chi(1-\hat{\gamma})\,\sigma_\theta^2 \cdot \overline{s}$, which scales with risk aversion, posterior uncertainty, and average supply.

This information updating wedge results from market clearing with dispersed information and is perfectly consistent with Bayesian rationality. An increase of the sufficient statistic from z to z' conveys positive news about θ through the information contained in the price and thus raises dividend expectations for all traders in the market. In addition, such a shift must come from either a reduction in supply s or an increase in θ, which shifts the distribution of private signals and thus increases asset demand. In both cases a further price adjustment is needed to clear the market. The expression for the equilibrium price incorporates the effect of the sufficient statistic z on posterior beliefs through the price signal weighted by τ, and the extra adjustment due to market-clearing with the additional weight β. In contrast, the Bayesian posterior of θ given z only includes the first effect. Theorem 1 in Section 4 shows that under weak regularity conditions, the equilibrium price with noisy information aggregation always admits a sufficient statistics representation with an information updating wedge of the form described by equation (3).

The information updating wedge arises from a combination of dispersed information and finitely elastic asset demand, due to risk aversion. The latter implies that demand (fundamental) and supply shocks have price impact, as reflected by the extra price adjustment required to clear markets. Dispersed information allows us to represent these price changes as shifts in the marginal investor's private information. Compare this with an otherwise identical market in which all investors share the same information $z \sim \mathcal{N}\left(\theta, \tau^{-1}\right)$. With CARA preferences and taking as given supply s, the

[15]The only prior discussion of the updating wedge we are aware of is by Vives (2008).

asset price with *common information* is $V(z,s) = \mathbb{E}(\theta|z) - \chi Var(\theta|z)\,s$.

Risk-neutral measure and excess weight on tail risks: The asset price can equivalently be represented as

$$P(z) = \mathbb{E}\left((\theta - R)\cdot m^I(\theta,z)\,|z\right)$$

where $m^I(\theta,z) \equiv \frac{h(\theta|x=z,z)}{h(\theta|z)}$ represents the change in probability measure implied by noisy information aggregation and $R \equiv \chi(1-\hat{\gamma})\sigma_\theta^2 \cdot \bar{s}$ the expected risk premium.[16] Taking prior expectations, we obtain

$$\mathbb{E}(P(z)) = \mathbb{E}((\theta - R)\cdot m(\theta)) = \int_{-\infty}^{\infty} \theta' d\hat{H}(\theta')$$

where $m(\theta) = \mathbb{E}\left(m^I(\theta,z)\,|\theta\right)$, and $\hat{H}(\theta) = \int_{-\infty}^{\theta} m(\theta'+R)dH(\theta'+R)$. It is straight-forward to check that $\mathbb{E}\left(m^I(\theta,z)\,|z\right) = \mathbb{E}(m(\theta)) = 1$, i.e. $m^I(\theta,z)$ and $m(\theta)$ represent the changes in probability measure and $\hat{H}(\theta)$ the risk neutral probability measure associated with the equilibrium price. The latter combines a risk adjustment $R = \chi(1-\hat{\gamma})\sigma_\theta^2\cdot\bar{s}$ that represents a parallel shift in the distribution of fundamentals with the informational adjustment $\hat{m}(\cdot)$ that captures the information updating wedge.

Equivalently, we can construct $\hat{H}(\cdot)$ by compounding the posterior $\theta|x=z,z\sim\mathcal{N}\left(\hat{\gamma}z,(1-\hat{\gamma})\sigma_\theta^2\right)$, with the prior distribution over z, $z\sim\mathcal{N}\left(0,\sigma_\theta^2+\tau^{-1}\right)$, and adjusting the mean for the risk-premium term $R = \chi(1-\hat{\gamma})\sigma_\theta^2\cdot\bar{s}$. Hence $\hat{H}(\cdot)$ is normal with mean $-R$ and variance $\hat{\sigma}_\theta^2 \equiv \left(1-\hat{\gamma}+\hat{\gamma}^2/\gamma\right)\sigma_\theta^2 > \sigma_\theta^2$. In other words, $\hat{H}(\cdot)$ departs from the prior $H(\cdot)$ through an adjustment of the mean and a mean-preserving spread, a property that we refer to as *excess weight on tail risks*: controlling for the mean $-R$, $\hat{H}(\cdot)$ places higher weight on realizations of θ in both upper and lower tails. EWTR distinguishes the risk-neutral measure under dispersed information from an average risk premium, which shifts probability mass from the upper to the lower tail realizations, analogous to the shift $-R$ in the mean of the distribution.

Combining the Law of Total Variance with the observation that $Var(\theta|x,z)$ is independent of the realization of x, we obtain that $\hat{\sigma}_\theta^2 - \sigma_\theta^2 = Var(\mathbb{E}(\theta|x=z,z)) - Var(\mathbb{E}(\theta|x,z))$. Hence EWTR ($\hat{\sigma}_\theta^2 > \sigma_\theta^2$) is equivalent to saying that the posterior expectations under the risk-neutral measure are strictly more variable than the posterior expectations of an arbitrary informed trader in the market. This property emerges because supply shocks introduce fluctuations in risk-neutral expectations that are orthogonal to the investors' private signals of fundamentals.

[16] When $m(\theta,z) = \frac{h(\theta|x=z,z)}{h(\theta|z)}$, $\mathbb{E}(\theta|x=z,z) = \int\theta\cdot\frac{h(\theta|x=z,z)}{h(\theta|z)}\cdot h(\theta|z)d\theta = \mathbb{E}(\theta\cdot m(\theta,z)\,|z)$.

Supply shocks vs. dispersed information: We can re-state $\hat{\sigma}_\theta^2$ as

$$\hat{\sigma}_\theta^2 = \sigma_\theta^2 + \frac{\beta + \chi^2 \sigma_s^2}{\left(1/\sigma_\theta^2 + \beta + \beta^2 / \left(\chi^2 \sigma_s^2\right)\right)^2}.$$

It follows that in the limit as $\beta \to 0$ (private information vanishes), $\hat{\sigma}_\theta^2$ converges to $\sigma_\theta^2 + \sigma_\theta^4 \cdot \chi^2 \sigma_s^2 > \sigma_\theta^2$. In other words, the presence of supply shocks on its own is already sufficient to generate EWTR, even when all traders have identical beliefs. The reason is that supply shocks vary the risk premium required for holding the asset, which generates price fluctuations that are independent of the underlying fundamentals. These orthogonal price fluctuations are captured by the EWTR in the risk-neutral measure.

On the other hand, for small positive β, $\hat{\sigma}_\theta^2 \approx \sigma_\theta^2 + \sigma_\theta^4 \cdot \left(\beta + \chi^2 \sigma_s^2\right)$ is increasing in β. Dispersed information thus amplifies EWTR when private signals are sufficiently noisy. Furthermore, consider the limit when β and σ_s^2 become small but the precision of the price signal τ is held constant. Since $\beta = \sqrt{\tau} \cdot \chi \sigma_s$, the EWTR $\hat{\sigma}_\theta^2 - \sigma_\theta^2$ is of order $\chi \sigma_s$ with dispersed information, but of order $\chi^2 \sigma_s^2$ when $\beta = 0$. Hence when supply shocks are small and private signals sufficiently noisy, amplification can become arbitrarily large. On the other hand, private information obviously reduces excess weight on tail risks, if signals are sufficiently informative: in the limit as $\beta \to \infty$, the price becomes perfectly revealing and must therefore converge to the true fundamental, in which case the risk-neutral distribution coincides with the prior distribution over θ. The amplification arises because noisy information aggregation generates negative co-movement between investor dividend expectations and the stochastic risk premia required to compensate investors for holding the asset: a positive supply shock that increases exposures leads to a higher risk premium and a lower equilibrium price. But since traders can't distinguish whether price movements are driven by fundamentals or supply shocks, they view the price reduction as possibly negative news about fundamentals thus adjusting their expectations downwards, which in turn amplifies price fluctuations and EWTR.[17]

Asset-pricing predictions: Figure 1 compares the price (thick solid line) with expected dividends (thin solid line), conditional on z (assuming average supply $\overline{s} = 0$). Since $\hat{\gamma} > \gamma$, the price responds more to z than the underlying dividend expectations. We can then derive predictions for return premia by taking expectations w.r.t. z, or equivalently by comparing expected cash-flows under the risk-neutral measure $\hat{H}(\cdot)$ and the physical distribution $H(\cdot)$. However, such a comparison is of limited interest in the CARA-normal model, since the average price only depends on the expected risk premium R and not on the EWTR property. This can be seen by comparing the solid horizontal line (average price with $\overline{s} = 0$) with the *dashed* horizontal line (average price

[17] We generalize these observations below in section 5.

Figure 1: Linear payoff model –CARA-normal example

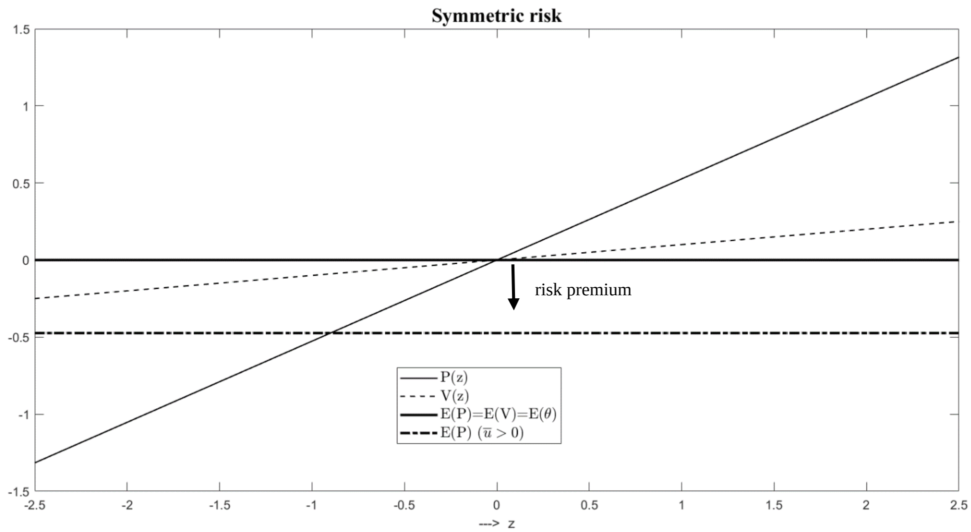

with $\overline{s} > 0$). Indeed, when $\overline{s} = 0$, the unconditional expectations of prices and dividends coincide, as payoff symmetry implies that overpricing when z is positive is exactly offset by underpricing when z is negative. We must therefore look beyond the symmetry assumptions embedded in the CARA-normal example to analyze how noisy information aggregation and EWTR affect average prices and returns.

2.2 Example 2: the CARA-binary model

Our second example illustrates how noisy information aggregation leads to a non-zero premium in expected prices when asset payoffs are asymmetric, even when the expected asset supply is zero.

We consider a model with CARA preferences but assume binary payoffs: $\pi(\theta) = \theta$, where $\theta \in \{0, 1\}$, with ex-ante probability $\Pr(\theta = 1) = \lambda > 0$. The parameter λ measures the degree of *upside* versus *downside risk*: if $\lambda > 1/2$, the security is a downside risk; if $\lambda = 1/2$, the risk is symmetric; if $\lambda < 1/2$, the security is an upside risk. All other elements and notation are kept as in section 2.1.[18]

With CARA preferences, an investor with private signal x demands

$$d(x, P) = \frac{1}{\chi}\left(\log\left(\frac{\mu(x, P)}{1 - \mu(x, P)}\right) - \log\left(\frac{P}{1 - P}\right)\right) \tag{4}$$

units of the security, where $\mu(x, P)$ represents the informed investor's posterior belief that $\theta = 1$, conditional on observing private signal x and price P. As before, we construct a sufficient statistic

[18] This example is a special case of example 1 in Breon-Drish (2015).

$z(P)$ as the private signal of the investor who, at a given price P, finds it optimal to hold exactly $\bar{s}$ units of the asset. We conjecture that $z|\theta \sim \mathcal{N}\left(\theta, \tau^{-1}\right)$ is normally distributed with mean $\theta \in \{0,1\}$ and precision τ, in which case $\mu(x, P(z))$ takes the form

$$\log\left(\frac{\mu(x, P(z))}{1-\mu(x, P(z))}\right) = \log\left(\frac{\lambda}{1-\lambda}\right) + \beta\left(x - \frac{1}{2}\right) + \tau\left(z - \frac{1}{2}\right). \tag{5}$$

Substituting (5) into (4) and inverting the condition $d(z, P) = \bar{s}$ then leads to the following sufficient statistic representation of the equilibrium price:

$$P(z) = \frac{\lambda e^{(\beta+\tau)\left(z-\frac{1}{2}\right)-\chi\bar{s}}}{\lambda e^{(\beta+\tau)\left(z-\frac{1}{2}\right)-\chi\bar{s}} + 1 - \lambda}. \tag{6}$$

From the market-clearing condition $s = \int d(x, P)\, d\Phi\left(\sqrt{\beta}(x-\theta)\right)$, for $\theta \in \{0,1\}$, it is then straight-forward to verify that $z = \theta - \chi/\beta \cdot (s - \bar{s})$ and hence $z|\theta \sim \mathcal{N}\left(\theta, \tau^{-1}\right)$ with $\tau^{-1} = (\chi/\beta)^2 \cdot \sigma_s^2$, which confirms our initial conjecture. As in the CARA-normal model, equation (6) represents the price as a function of the sufficient statistics z, with an updating wedge: the log-odds ratio implied by the price attributes a weight $\tau + \beta$ to the market signal z, rather than just τ, and also includes a risk adjustment $-\chi\bar{s}$ to compensate investors for their expected exposure $\bar{s}$.

With binary payoffs, the expected price is equal to the market-implied or risk-neutral prior that $\theta = 1$, which we can compute by taking expectations over the price function in (6). Proposition 1 shows that when the asset is in zero expected supply, the expected price or risk-neutral prior attaches a higher probability to "tail risks":

Proposition 1 *: If $\bar{s} = 0$, there exists $\Delta \in (0,1)$ such that the expected price takes the form*

$$\mathbb{E}(P(z)) = \lambda + \left(\frac{1}{2} - \lambda\right)\Delta. \tag{7}$$

Moreover, $\lim_{\beta\to 0} \Delta = \lambda(1-\lambda)(\chi\sigma_s)^2 + o\left((\chi\sigma_s)^4\right)$ *and* $\lim_{\beta\to 0, \chi\sigma_s = \frac{\beta}{\sqrt{\tau}}} \frac{\Delta}{\beta} = \lambda(1-\lambda) + o(\tau)$.

Hence when the expected asset supply is 0, the expected price includes an adjustment $(1/2 - \lambda)\Delta$ that increases the expected asset price whenever the asset is characterized by upside risk ($\lambda < 1/2$), and decreases the expected asset price whenever the asset is characterized by downside risk ($\lambda > 1/2$). This adjustment introduces a positive relation between the skewness of the security and its expected price, we therefore refer to $(1/2 - \lambda)\Delta$ as the skewness premium.

To understand the intuition of this result, Panels a) and b) in Figure 2 plot the price and dividend expectation for different realizations of z, for the cases of downside ($\lambda = 0.9$), and upside ($\lambda = 0.1$) risks, all under $\bar{s} = 0$. As before, the price responds more strongly to z than the dividend expectation and thus crosses the latter from below for some realization of z. For symmetric payoffs (not plotted), and as was the case for the CARA-normal model, the average price coincides with the average payoff.

For downside risks, Panel a) of the Figure shows that the difference between price and expected dividend becomes highly asymmetric and is much more pronounced for low realizations of z than for high draws. Intuitively, for a given dispersion of private signals, disagreement among investors about payoffs is rather limited for high values of z, as cash-flow variability is much more bounded on the upside relative to the prior. For low draws of z, on the other hand, the risk shifts towards a situation where the security payoff may fall considerably below its prior, heightening the dispersion of payoff expectations among investors, thus amplifying the difference between the price and the expected dividend as function of z. Taking the average across realizations, this asymmetry naturally leads to expected prices (thick solid horizontal line)below average payoffs (thin solid horizontal line), even when the risk-premium compensation is by construction set to zero ($\bar{s} = 0$). The opposite situation arises with upside risks.

In addition, we obtain, as in the CARA-normal model, that the skewness premium is of order $\chi^2\sigma_s^2$ when private information vanishes, but of order $\beta = \sqrt{\tau} \cdot \chi\sigma_s$, when the price remains informative in the limit. In other words, once again the skewness premium already results from the presence of supply shocks, but is amplified by noisy information aggregation.

Recall that in the CARA-normal model, we linked the presence of EWTR without private signals to stochastic risk premia that generate price fluctuations that are orthogonal to fundamentals, and we linked the amplification from private information to co-movement between average risk premia and dividend expectations. The same forces are at play in the CARA-binary model, but in addition the asymmetry in payoffs makes both the stochastic risk premia and their co-movement with dividend expectations asymmetric: with upside (downside) risk, the risk premium required to compensate investors for a given size positive asset exposure is strictly smaller (larger) in absolute value than the risk premium for taking on a negative exposure of equal size. This is due to *downside risk aversion*: investors require extra compensation for downside exposures, and when the security is characterized by upside risk ($\lambda < 1/2$), the downside exposure is larger for negative than for positive positions, and the opposite is true for downside risks ($\lambda > 1/2$). The magnitude of the asymmetry between upside and downside risks is then scaled by the variance of supply shocks.

Likewise, with noisy private signals, a positive supply shock increases required compensation for risk and lowers the asset price, which in turn leads investors to lower their dividend expectations. But this also reduces their posterior uncertainty for upside risks, since the unlikely event of a high payoff becomes even more remote, but increases posterior uncertainty for downside risks since investors become more worried about the unlikely scenario that the asset may actually fail to pay off. For upside risks, this means that the co-movement of risk premia with dividend expectations is stronger on the upside than the downside, since uncertainty and exposures are negatively correlated (and risk premia are the product of the two), while for downside risks, the co-movement of risk premia with dividend expectations is stronger on the downside than on the upside. This asymmetric amplification then generates a positive price premium for upside risks and a negative price premium for downside risk, which is captured by the value of Δ in proposition 1. This information-based skewness premium scales with the standard deviation of supply shocks σ_s and thus becomes the dominant force behind price premia for skewness when private information is noisy yet the price remains informative.[19]

The characterization in proposition 1 can be extended to the case with $\overline{s} \neq 0$ as follows:[20] Let $\tilde{\lambda} = \lambda e^{-\chi \overline{s}} / \left(\lambda e^{-\chi \overline{s}} + 1 - \lambda\right)$ denote a risk-adjusted prior that $\pi = 1$. With $\overline{s} \neq 0$, we can replace λ with $\tilde{\lambda}$ and go through the same steps using the risk-adjusted prior $\tilde{\lambda}$. However, when computing $\mathbb{E}\left(P\left(z\right)\right) - \lambda$ one must correct for the gap between λ and $\tilde{\lambda}$, which yields the additional risk premium term $\left(\lambda - \tilde{\lambda}\right) R$, with $R < 1$. The risk premium scales with the difference between the objective and the risk-adjusted probability that $\theta = 1$, which depends on the expected asset supply $\overline{s}$. The contribution of this term is illustrated in Figure 2 through the dashed horizontal line, corresponding to the average price when $\overline{s} = 0$ (thick solid horizontal line) minus the risk premium that arises when $\overline{s} > 0$ (a negative average supply would lead to the opposite result).

Summing up, the CARA-binary example illustrates that noisy information aggregation generates a skewness premium in asset prices that is separate from the risk premium that emerges from non-zero average supply. However, this example relies heavily on the specific features introduced by CARA preferences and binary payoffs, limiting its applicability for broader, more realistic security classes and quantitative explorations. Moreover, the binary distribution of payoffs links expected payoff, payoff uncertainty and payoff asymmetry all to the same parameter λ. Hence our interpretation of Δ as a price premium linked to skewness, along with the amplification result from noisy information aggregation, is at best suggestive, since the formal analysis doesn't clearly separate the respective roles of payoff uncertainty and asymmetry. Our next example relaxes the present payoff assumptions and formally establishes a tight link between the expected price premium and

[19] In section 5, we generalize this decomposition into preference- and information-based skewness premia, along with the amplification result, to general preferences and securities, in a manner that clearly highlights the respective roles of downside risk aversion and posterior uncertainty.

[20] Details are presented in the online appendix.

Figure 2: Non-linear payoff models –CARA-binary and risk neutral, normal model

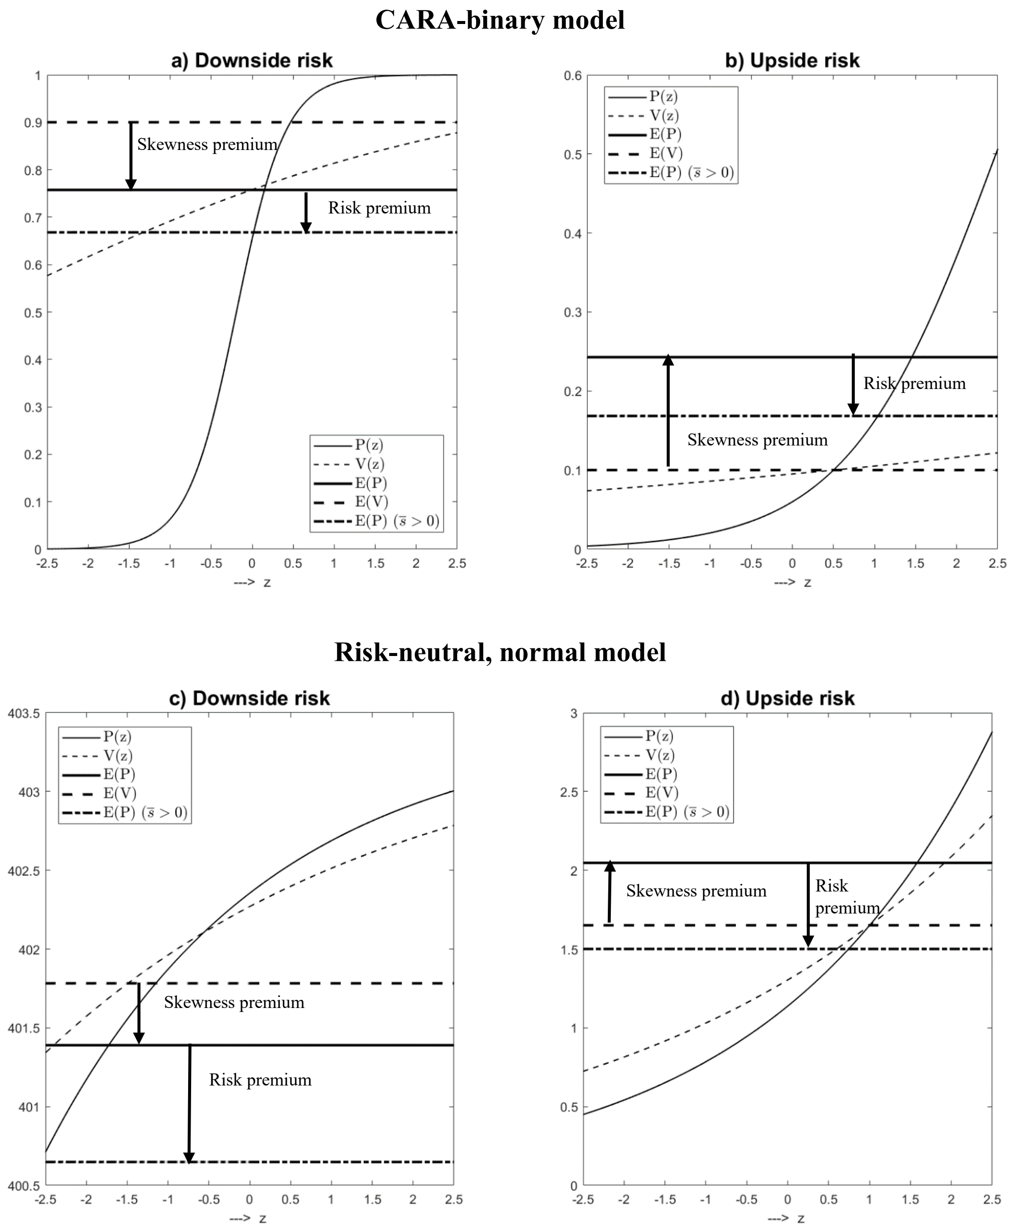

the asymmetry of asset payoffs.

2.3 The risk-neutral, normal model

We now introduce a variant of the general model that allows us to generalize the key ideas of the previous examples without imposing strong restrictions on security payoffs. Specifically, we assume that informed investors are risk-neutral and face binding position limits. This model is rich enough to convey our main theoretical results, yet tractable enough to allow closed-form solutions that facilitate the derivation of rich empirical predictions for returns to skewness and disagreement. It also helps to tie model parameters more closely to observables, which we explore in the next section.

We view the risk-neutral model with position limits as depicting the activity of one among many parallel securities markets, and interpret comparative statics results as cross-sectional predictions about asset returns. This model is of special interest because risk preferences disappear from the equilibrium characterization. It therefore strikes us as a natural laboratory for studying cross-sectional return predictions with noisy information aggregation, since investors should be able to diversify asset-specific risks by investing across a wide range of assets, i.e. risk preferences disappear from the characterization of returns. Such diversification can be achieved by limiting the wealth that is invested in any given security, akin to position limits in our model. In section 5 we return to the discussion of multi-asset extensions of our model, and also address conditions under which one may consider asset markets one-by-one, in isolation from each other.

We consider the following specialization of the general model: (i) normally distributed fundamentals $\theta \sim \mathcal{N}\left(0, \sigma_\theta^2\right)$ and private signals $x_i|\theta \sim \mathcal{N}\left(\theta, \beta^{-1}\right)$, (ii) risk-neutral preferences $U(w) = w$ with positions limited by $[d_L, d_H] = [0, 1]$,[21] and (iii) stochastic asset supply $s = \Phi(u)$, where $\Phi(\cdot)$ is the cdf of a standard normal distribution, and $u \sim \mathcal{N}\left(0, \sigma_u^2\right)$. Importantly, our analysis imposes no restriction on the payoff function $\pi(\cdot)$, allowing us to offer asset pricing predictions for arbitrary payoff distributions. The functional form assumption about asset supply is adapted from Hellwig et al. (2006) and appears in similar form in Goldstein et al. (2013).[22] It keeps the updating problem tractable by preserving normality of the investors' posterior beliefs.

As before, we first derive a price function in terms of a sufficient statistic z, and then use it to obtain a risk-neutral representation of the expected price that displays EWTR. Finally we derive

[21] Our main predictions are not dependent on the specific bounds assumed, such as the short-sale constraint ($d_L = 0$) which we assume only for simplicity.

[22] Goldstein et al.(2013) similarly assume risk-neutral investors and position limits, but focus on specific return assumptions that allow them to analyze informational feedback from financial markets to investment decisions.

comparative statics of expected prices from the interaction between EWTR and payoff asymmetries.

Characterization of equilibrium price:

Standard arguments imply that $H(\cdot|x,P)$ must be first-order stochastically increasing in the investor's signal x. There then exists a unique signal threshold $\bar{x}(P)$ such that

$$\int \pi(\theta)\, dH(\theta|x,P) \gtreqless P \quad \text{if and only if} \quad x \gtreqless \bar{x}(P), \tag{8}$$

and investors find it optimal to purchase one unit of the security if and only if their private signal x exceeds the threshold $\bar{x}(P)$, otherwise they do not buy. This leads to an aggregate demand by informed investors that is equal to $D(\theta,P) = Pr(x \geq \bar{x}(P)|\theta) = 1 - \Phi(\sqrt{\beta}(\bar{x}(P) - \theta))$.

Setting $z \equiv \bar{x}(P)$, market-clearing then implies $1 - \Phi(\sqrt{\beta}(z-\theta)) = s = \Phi(u)$, or equivalently $z = \theta - 1/\sqrt{\beta} \cdot u$, and $z|\theta \sim \mathcal{N}(\theta, \tau^{-1})$, where $\tau \equiv \beta/\sigma_u^2$. Substituting this characterization of z into the marginal investor's indifference condition (8) leads to the following proposition:

Proposition 2 *: The unique price-monotone equilibrium is characterized by the equilibrium price function*

$$\begin{aligned} P_\pi(z) &= \mathbb{E}(\pi(\theta)|x=z,z) = \int \pi(\theta) d\Phi\left(\frac{\theta - \hat{\gamma}z}{\sqrt{1-\hat{\gamma}}\sigma_\theta}\right) \\ \textit{where } \hat{\gamma} &\equiv \frac{\beta+\tau}{1/\sigma_\theta^2 + \beta + \tau}. \end{aligned} \tag{9}$$

As before, Proposition 2 represents the equilibrium price in terms of a sufficient statistic z for the information conveyed through the asset price. This sufficient statistic corresponds to the private signal of the marginal investor who is just indifferent between buying and not buying the asset. Once again we can equate the price to the dividend expectation of a hypothetical investor who treats the market signal z as if it had precision $\beta + \tau$.

The expected dividend conditional on z instead takes the form

$$\mathbb{E}(\pi(\theta)|z) = \int \pi(\theta) d\Phi\left(\frac{\theta - \gamma \cdot z}{\sqrt{1-\gamma}\sigma_\theta}\right) \quad \text{where} \quad \gamma \equiv \frac{\tau}{1/\sigma_\theta^2 + \tau}. \tag{10}$$

The comparison between equations (9) and (10) shows exactly the same information updating wedge as in the CARA-normal and CARA-binary examples. The intuition behind these expressions is also the same: in order to clear the market, the expectation of the marginal investor must respond more strongly to changes in the sufficient statistic z than is warranted purely by the information conveyed through it, which occurs through the *systematic shift in the identity* of this marginal investor in order to accommodate fluctuations in supply.

Asset pricing implications of noisy information aggregation:

Proposition 2 allows us to represent the expected price $\mathbb{E}(P(z)) = \int \pi(\theta)\, d\hat{H}(\theta)$ as an expectation of dividends under a risk-neutral measure $\hat{H}$. As in the CARA-normal example, we compound the risk-neutral posterior $\theta|x = z, z \sim \mathcal{N}\left(\hat{\gamma} \cdot z, \sigma_\theta^2 (1 - \hat{\gamma})\right)$ with the prior $z \sim \mathcal{N}\left(0, \sigma_\theta^2/\gamma\right)$ to show that the risk-neutral measure $\hat{H}(\cdot)$ is normal with mean 0 and variance

$$\hat{\sigma}_\theta^2 = \left(1 - \hat{\gamma} + \hat{\gamma}^2/\gamma\right) \sigma_\theta^2 > \sigma_\theta^2. \tag{11}$$

This representation is identical to the one found in the CARA-normal example, aside from the fact that the precision τ of the sufficient statistic is calculated differently. The difference between $\hat{H}(\cdot)$ and $H(\cdot)$ thus corresponds to a mean-preserving spread characterized by the difference between the risk-neutral prior and the objective prior variance, $\hat{\sigma}_\theta^2$ as compared to σ_θ^2. Taking expectations we represent the difference between the expected price and expected payoff as

$$W(\pi, \hat{\sigma}_\theta) \equiv \mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta)) = \int_{-\infty}^{\infty} \left(\pi\left(\frac{\hat{\sigma}_\theta}{\sigma_\theta}\theta\right) - \pi(\theta)\right) d\Phi\left(\frac{\theta}{\sigma_\theta}\right). \tag{12}$$

The term $W(\pi; \hat{\sigma}_\theta)$ summarizes the impact of the mean-preserving spread on the expected price premium. Our next definition provides a partial order on payoff functions that we use for the comparative statics of $W(\pi, \hat{\sigma}_\theta)$.

Definition 1 ***(Cash flow risks)*** *:*
(i) Payoff function π has symmetric risk if $\pi(\theta_1) - \pi(\theta_2) = \pi(-\theta_2) - \pi(-\theta_1)$ for all $\theta_1 > \theta_2 \geq 0$.
(ii) Payoff function π is dominated by upside risk if $\pi(\theta_1) - \pi(\theta_2) \geq \pi(-\theta_2) - \pi(-\theta_1)$, and dominated by downside risk if $\pi(\theta_1) - \pi(\theta_2) \leq \pi(-\theta_2) - \pi(-\theta_1)$, for all $\theta_1 > \theta_2 \geq 0$.
(iii) Payoff function π_1 has more upside risk than π_2 if $\pi_1 - \pi_2$ is dominated by upside risk.

This definition classifies payoff functions by comparing marginal gains and losses at fixed distances from the prior mean to determine whether payoff fluctuations are larger on the upside or on the downside. Any linear payoff function has symmetric risks, any convex function is dominated by upside risks, and any concave function by downside risks, but the classification extends to more general non-linear functions with symmetric gains and losses, non-convex functions with upside risk or non-concave functions with downside risk.

Securities are easy to classify according to this definition when the fundamental and the return are both observable (for example in the case of defaultable bonds or options). But even without direct information about fundamentals, for any symmetric prior $H(\cdot)$, upside and downside risk directly translate into the distribution of returns being more spread out above or below the median of the return distribution. Intuitively, this means that a security that is dominated by upside (downside) risk has positive (negative) skewness. The reverse may not be true, however: a security may be positively skewed overall, but have local violations of the upside risk condition for some realizations of θ. As an expectation of third moments, skewness only offers a summary measure of all the local asymmetries between upside and downside risks that are captured by definition 1.[23]

The following proposition generalizes the insights of proposition 1 from the CARA-binary model to arbitrary returns in the risk-neutral model with position limits.

Proposition 3 ***(Sign and comparative statics of*** $W(\pi, \hat{\sigma}_\theta)$***):***

(i) If π *has symmetric risk, then* $W(\pi; \hat{\sigma}_\theta) = 0$*. If* π *is dominated by upside risk, then* $W(\pi; \hat{\sigma}_\theta)$ *is positive and increasing in* $\hat{\sigma}_\theta$*. If* π *is dominated by downside risk, then* $W(\pi; \hat{\sigma}_\theta)$ *is negative and decreasing in* $\hat{\sigma}_\theta$*. Moreover,* $lim_{\hat{\sigma}_\theta \to \sigma_\theta} |W(\pi; \sigma_\theta)| = 0$*; and* $lim_{\hat{\sigma}_\theta \to \infty} |W(\pi; \hat{\sigma}_\theta)| = \infty$*, whenever* $lim_{\theta \to \infty} |\pi(\theta) + \pi(-\theta)| = \infty$*.*

(ii) If π_1 *has more upside risk than* π_2*, then* $W(\pi_1; \hat{\sigma}_\theta) - W(\pi_2; \hat{\sigma}_\theta)$ *is non-negative and increasing in* $\hat{\sigma}_\theta$*.*

Proposition 3 shows how the expected price premium arises as a combination of asymmetry in the payoff function π and noisy information aggregation through its effect on EWTR ($\hat{\sigma}_\theta > \sigma_\theta$). The price premium is positive for upside risks, negative for downside risks and larger in absolute value for assets with larger return asymmetries. Furthermore, the price premium increases and can grow arbitrarily large as EWTR becomes more important, and it vanishes as $\hat{\sigma}_\theta \to \sigma_\theta$. The latter case arises, for example, when private signals become infinitely noisy ($\beta \to 0$), or the price signal infinitely precise ($\sigma_u^2 \to 0$). Finally, price impact of information aggregation frictions and return asymmetry are mutually reinforcing, since the price premium has increasing differences in upside risk and EWTR. These results follow directly from our interpretation of the skewness premium as a mean-preserving spread, which becomes more valuable when the payoff function shifts towards more upside risk.

Panels c) and d) in Figure 2 help with the intuition of Proposition 2. With symmetric $\pi(\cdot)$ (not plotted), EWTR does not affect the average price: over-pricing when z is high is just offset by under-pricing when z is low. When instead $\pi(\cdot)$ is dominated by downside risk (panel c), the lower tail risk of dividends is more important than the upper tail, and divergence of opinion between investors becomes more consequential and widens the gap between expectations of the marginal

[23] Indeed the relation between the partial order defined by upside and downside risk according to Definition 1 and the summary measure provided by skewness is akin to the comparison between ranking distributions by second-order stochastic dominance and the summary ranking by the variance of the distribution.

investor (eg., the price) and the expected dividend conditional only on the price signal, z. Hence under-pricing on the downside is larger than over-pricing on the upside, which results in a positive skewness premium. When instead $\pi(\cdot)$ is dominated by upside risk (panel d), reverse arguments apply and the skewness premium is positive.

So far, the risk-neutral model abstracted from average supply effects like the risk premium in the CARA-normal and CARA-binary models. It is possible to reintroduce them by assuming that supply takes the form $s = \Phi(u)$, where $u \sim \mathcal{N}\left(\bar{u}, \sigma_u^2\right)$, with $\bar{u} \neq 0$. In this case, the risk-neutral measure $\hat{H}(\cdot)$ is normal with mean $\bar{\theta} \equiv -\frac{\sqrt{\beta}}{1/\sigma_\theta^2+\beta+\tau}\bar{u}$ and variance $\hat{\sigma}_\theta^2$. Hence, a higher supply ($\bar{u} > 0$) results in a downwards shift of the risk-neutral distribution, similar to the CARA-normal case. The expected price premium then decomposes into a component reflecting the shift in means and a mean-preserving spread that inherits the same properties as described above in Proposition 3. The downwards shift in means leads to a lower average price and is plotted in panels c) and d) of Figure 2 (dashed horizontal lines). The opposite case (upward shift) arises for $\bar{u} < 0$ (not plotted).

EWTR and Forecast Dispersion:

Equation (11) defines EWTR $\hat{\sigma}_\theta/\sigma_\theta$ in terms of two parameters: the precision of private information and the variance of noise trading. We conclude this section by showing how EWTR can equivalently be represented in terms of two statistics that can, in principle, be estimated using data on investors' forecasts of fundamentals. These statistics allow us to translate the comparative statics of Proposition 3 into testable predictions.

Define $\hat{\gamma} = 1 - \frac{1/\sigma_\theta^2}{1/\sigma_\theta^2+\beta+\tau} = 1 - \frac{Var(\theta|x,z)}{Var(\theta)}$ as a measure of the *accuracy* of investors' forecasts of fundamentals, and

$$\tilde{D} \equiv \sqrt{\frac{Var\left(\mathbb{E}\left(\theta|x,z\right)|\theta,z\right)}{Var\left(\theta\right)}} = \frac{\beta}{1/\sigma_\theta^2+\beta+\tau}\frac{\beta^{-1/2}}{\sigma_\theta} = \frac{\sqrt{\beta}}{\left(1/\sigma_\theta^2+\beta+\tau\right)}\frac{1}{\sigma_\theta} \tag{13}$$

as a measure of forecast *dispersion*, defined as the cross-sectional standard deviation of investor expectations, normalized by the standard deviation of fundamentals. We then represent EWTR $\hat{\sigma}_\theta/\sigma_\theta$ in terms of these two statistics:

$$\frac{\hat{\sigma}_\theta}{\sigma_\theta} = \sqrt{1+\tilde{D}^2\frac{\hat{\gamma}\left(1-\hat{\gamma}\right)}{\hat{\gamma}\left(1-\hat{\gamma}\right)-\tilde{D}^2}} \tag{14}$$

Representation (14) shows that EWTR increases with forecast dispersion and is U-shaped in forecast accuracy, reaching a minimum at $\hat{\gamma} = 1/2$, and diverging to infinity when $\hat{\gamma}\left(1-\hat{\gamma}\right) \rightarrow \tilde{D}^2$. However, for most of its range, the effect of forecast accuracy $\hat{\gamma}$ on $\hat{\sigma}_\theta/\sigma_\theta$ is very mild as we move away from $\hat{\gamma} = 0.5$.[24] In other words, forecast dispersion $\tilde{D}$ emerges as a natural empirical proxy variable for EWTR. We can then re-cast the comparative statics with respect to $\hat{\sigma}_\theta$ in proposition 3 in terms of forecast dispersion. Our model thus predicts a positive (negative) relation between forecast dispersion and returns for securities with downside (upside) risks. This distinguishes our theory from models of heterogeneous priors with short-sales constraints in which higher disagreement generates a positive option value of resale regardless of a security's return structure.

3 Returns to skewness and disagreement: empirical evidence and numerical calibration

3.1 Empirical evidence

Proposition 3 offers three qualitative predictions: returns to skewness, returns to disagreement, and positive interaction effects. They all follow from the core prediction that the risk-neutral measure displays excess weight on tail risks. We now briefly summarize the empirical evidence that offers qualitative support for these predictions.

Core prediction (Excess Weight on Tail Risks): Starting with Jackwerth (2000) and Ait-Sahalia and Lo (2000), a substantial empirical literature documents the *pricing kernel puzzle*, i.e. the empirical observation that pricing kernels recovered from option prices on stock indices appear to be non-monotone, and in certain cases U-shaped, confirming the core implication of our theoretical model. Christoffersen et al. (2013) use a GARCH-based model of option pricing with stochastic volatility, Audrino et al. (2022) apply the Ross (2015) recovery theorem, and Bakshi et al. (2010) derive pricing implications of securities characterized by pure upside risks. All three focus on index options and market returns, and find strong evidence in support of pricing kernels that are U-shaped or upwards-sloping on the upside. While our theory (in particular, the risk-neutral, normal model) may more naturally relate to cross-sectional variation in asset returns, the evidence presented in these papers offers at least suggestive support for our theory. Moreover, Carr and Wu (2008) document the variance premium puzzle: i.e. the observation that risk-neutral variance of returns is an upwards-biased predictor of the true variance of returns. Christoffersen et al. (2013) discuss the relation between the pricing kernel and variance premium puzzles. In their estimates, the implied variance of returns overestimates the true variance of returns by about 12 to 18%, which is consistent with an estimate of EWTR in the range of $\frac{\hat{\sigma}_\theta}{\sigma_\theta} \approx 1.06$ to 1.09 in our model. As

[24] For example, a value of $\tilde{D} = 0.2$, close to the average dispersion of earnings forecasts in GH, implies a minimum value of $\hat{\sigma}_\theta/\sigma_\theta = 1.0235$, but does not exceed 1.0244 for values of forecast accuracy in the range of $[0.25, 0.75]$.

we wil see below, these estimates are very similar to the ones that we obtain by calibrating $\frac{\hat{\sigma}_\theta}{\sigma_\theta}$ to match forecast accuracy and disagreement.

Our qualitative predictions then concern returns to skewness, disagreement and interaction effects. For a given EWTR $\hat{\sigma}_\theta$, we define a price premium for disagreement $W\left(\pi^{Q_H};\hat{\sigma}_\theta\right) - W\left(\pi^{Q_L};\hat{\sigma}_\theta\right)$, where π^Q represents a return distribution at a given quantile Q in the distribution of upside or downside risks (or skewness), with $Q_H > Q_L$. By sorting stocks according to their implied level of upside or downside risk, or skewness, we can obtain the empirical analogue of the skewness premium in returns. Likewise, for a given security π, we define a price premium for disagreement $W\left(\pi;\hat{\sigma}_\theta^{q_H}\right) - W\left(\pi;\hat{\sigma}_\theta^{q_L}\right)$, where $\hat{\sigma}_\theta^q$ represents the EWTR associated with a given disagreement quantile q, with $q_H > q_L$. By sorting stocks according to their level of disagreement, we can obtain the empirical analogue of the disagreement premium in returns.

Prediction 1 (Returns to skewness): *Price premia are positive (negative) and return premia negative (positive) for securities dominated by upside (downside) risk. Price premia are increasing and returns decreasing with skewness or upside risk.*

A sizable empirical literature documents a negative relationship between expected skewness and equity returns. For example, Conrad et al. (2013) estimate skewness of equity returns from option prices, Boyer et al. (2010; BMV henceforth) from forecasting regressions. Both studies then sort stocks by expected skewness and find that securities with higher skewness earn about 0.7% lower average returns per month, equivalent to more than 8% of yearly excess returns for the strategy of going long/short on low/high skewness stocks. Green and Hwang (2012) find that IPOs with high expected skewness earn significantly more negative abnormal returns in the following one to five years. Zhang (2013) shows that skewness correlates positively with the book-to-market factor and thus helps account for the value premium.

Returns to skewness also manifest themselves in bond markets through the *credit spread puzzle*, i.e. the difficulty to reconcile high levels of corporate bonds spreads with historical default data in standard models pricing credit risk. This shortfall is most severe for short maturity, high investment grade securities, which are almost as safe as treasuries of similar maturity, yet pay significantly larger return premia. [25]

[25] Huang and Huang (2012) calibrate a number of structural models to historical default data and show that they all produce spreads relative to treasuries that fall significantly short of their empirical counterparts. Chen (2010), He and Milbradt (2014) and Chen et al. (2018) develop dynamic models of credit risk with endogenous default, long-run risks and market liquidity. While they come closer to matching empirical counterparts, most purely risk- and liquidity-based models account for at most a small fraction of the level and volatility of spreads that are observed in practice, especially for short-horizon investment grade bonds. A separate literature has linked credit spreads and

Existing explanations for these empirical findings rely on investor preferences for positive skewness.[26] Proposition 3 offers an alternative explanation for a negative returns to skewness and high spreads on investment-grade corporate bonds as the result of EWTR. In contrast to preference-based theories, this explanation also links returns to skewness and disagreement. In section 5 we return to the comparison between preference- and information-based returns to skewness and show that dispersed information amplifies returns to skewness.

Prediction 2 (Returns to disagreement): *Price premia for disagreement are positive (negative) and return premia negative (positive) for securities dominated by upside (downside) risk.*

A growing literature uses forecast dispersion as an empirical proxy for disagreement. Dietheret al. (2002; DMS henceforth) sort stocks by the dispersion of earnings forecasts across analysts covering each security. They find that stocks in the highest dispersion quintile have monthly returns which are about 0.62% lower than those in the lowest dispersion quintile, amounting to a yearly excess return over 7% for the strategy of going long/short on low/high dispersion stocks. They interpret this result as evidence consistent with the hypothesis of Miller (1977) of investor disagreement interacting with short-selling constraints.[27] Yu (2011) reports similar results and Gebhardt et al. (2001) document that an alternative measure of stock risk premia, the cost of capital, is also negatively related to analyst forecast dispersion.

Güntay and Hackbarth (2010; GH henceforth) perform a similar analysis for bond yields but reach the opposite conclusion as DMS: yield spreads and bond returns are increasing with forecast dispersion, and spreads are 0.14% higher and returns 0.08% higher in the top dispersion quintile, which amounts to a yearly excess return of about 1% for the strategy long/short on high/low dispersion bonds. GH replicate DMS' result of negative returns to disagreement in equity returns in their sample (though the measured excess returns are smaller), which suggests a systematic difference in returns to disagreement for equity and bond markets. GH interpret returns to disagreement as a proxy for risk premia. Carlin et al. (2014) confirm GH's results for mortgage-backed securities.

Proposition 3 reconciles the seemingly contradictory empirical results about return premia for disagreement by noting that studies that find negative returns to disagreement focus on securities

equity returns through capital structure models with time-varying default risk (See Bhamra et al., 2010 and citations therein, as well as McQuade, 2018). Our analysis instead links credit spreads and (levered) equity returns through dispersed investor information and EWTR.

[26] In Brunnermeier and Parker (2005) and Brunnermeier et al. (2007), overinvestment in highly skewed securities, along with under-diversification, results from a model of optimal expectations. Barberis and Huang (2008) show that cumulative prospect theory results in a demand for skewness or a preference for stocks with lottery-like features. Mitton and Vorkink (2007) develop a model in which investors have heterogeneous preference for skewness.

[27] They rule out a risk-based explanation for the anomaly by controlling for stocks exposure to standard risk factors.

with upside risk, while studies that find positive returns to disagreement focus on securities where downside risk is dominant.

Our third prediction focuses on interaction effects: returns to disagreement increase with asset skewness, and returns to skewness increase with investor disagreement.

Prediction 3 (Interaction effects): *There is positive interaction between returns to disagreement and returns to skewness.*

Evidence on interaction effects between skewness and disagreement is more limited. Ideally, one would like to match measures of returns, expected skewness, and forecast disagreement and accuracy at the level of individual stocks, and then conduct a dual sorting exercise which could speak to all three predictions simultaneously. Unfortunately we are not aware of studies that take exactly this approach.[28] Yu (2011) comes closest to what we need by sorting stocks by book-to-market ratio and disagreement. He reports that the value premium increases from 4.3% annual return with the lowest tercile disagreement to 11.3% with the highest tercile, and the returns to disagreement range from -0.26% annual for the highest quintile of book-to-market ratios to -7.2% for the lowest quintile. Following Zhang (2013) who interprets book-to-market ratios as a proxy for skewness, these results suggest substantial interaction between returns to skewness and disagreement in equity returns.[29]

For bond markets, GH report that the effect of disagreement on spreads and yields doubles in high leverage or low-rated rated firms, two plausible proxies for downside risks. In a regression of credit spreads on leverage, disagreement and their interaction, the interaction term turns out to be highly significant, but disagreement and leverage are insignificant on their own. These empirical results suggest that returns to skewness and disagreement interact in the data along the lines suggested by our theoretical results.[30]

[28] Short of this ideal, the empirical studies on returns to skewness and disagreement that we cite are either based on different samples data sources, or focus on imperfect proxies for the skewness and disagreement or EWTR measures implied by our theory.

[29] More specifically, Zhang (2013) first documents strong positive correlation between book-to-market ratios and skewness of returns, and then shows that the explanatory power of book-to-market ratios for returns is significantly lower when controling for skewness.

[30] Hou et al. (2020) question the statistical robustness of various return anomalies in equity markets including the studies on returns to skewness and disagreement. They replicate existing studies on a uniform sample and emphasize the importance of small capitalizations and equal vs. value weighting in estimating return anomalies. They replicate DMS and show that value-weighting leads to much lower and statistically insignificant returns to disagreement. This suggests that returns to disagreement are concentrated in markets with small capitalizations, which is consistent with the replication of DMS by GH for a subsample of firms that are active also in bond markets. They also find returns to skewness that are small and insignificant, but by using the realized skewness of past returns rather than predicted

The ability to account for returns to skewness and disagreement in both equity and bond markets sharply distinguishes our theory from heterogeneous prior models with short-sales constraints following Miller (1977): in those models prices incorporate a resale option value that lowers future returns irrespective of the asset characteristics. They are therefore inconsistent with Predictions 2 and 3 and thus unable to explain why these comparative statics would be different for different security classes such as stocks and bonds, as suggested by the empirical evidence discussed in the preceding paragraphs.

Indeed, the ambiguous empirical relationship between disagreement and asset returns remains one of the major unresolved puzzles in asset pricing. Perhaps Carlin et al. (2014) put it best: *"Understanding how disagreement affects security prices in financial markets is one of the most important issues in finance. ...Despite the fundamental nature of this issue, though, there still remains significant controversy in the literature about how disagreement risk affects expected returns and asset prices."* To our knowledge, ours is the first explanation that can reconcile the seemingly contradictory empirical results as direct predictions of a unified theory, tractable enough to encompass assets with different underlying cash-flow risks.

3.2 A simple calibration

We offer a first attempt to quantify the role of noisy information aggregation for asset returns by informing key model parameters from data onasset payoffs and investor disagreement. The objective of the exercise is two-fold: first, to illustrate more formally how our model can be applied to shed light on specific empirical asset pricing puzzles. Second, to provide a first assessment of the quantitative potential of information frictions in explaining such puzzles. In doing so, it is important to note that since our contribution is essentially theoretical, we regard these results by no means as a conclusive test of our model, but we do hope that they will open the door for more sophisticated empirical work in the future.

Forecast dispersion and EWTR in the data:

We infer EWTR from measures of forecast dispersion $\tilde{D}$ and forecast accuracy $\hat{\gamma}$ (see online appendix for a more detailed description). In line with the empirical literature reviewed above, we interpret the asset fundamental as a firm's earnings, and derive measures of forecast dispersion and accuracy from the I.B.E.S. data of analyst earnings forecasts for a cross-section of listed firms. The

future skewness, they do not directly replicate Conrad et al. (2013) or BMV. Their estimates of the value premium are similar to the ones reported in Yu (2011). Our predictions are broadly consistent with Hou et al. (2020), if one assumes that larger markets are more liquid and less subject to noisy information aggregation frictions.

data reports a measure of forecast dispersion, along with a consensus or average earnings forecast and realized earnings per share for each firm-year in the sample. We base our estimates of forecast dispersion, forecast accuracy and EWTR on a sample of 5,320 firms used in the empirical study by Guntay and Hackbarth (GH), which uses forecasts over relatively short horizons (within quarter) from 1987-1998.[31] We use the time series of realized forecast errors to compute forecast accuracy, and substitute our estimates of forecast dispersion and accuracy into equation (14) to construct a firm-level estimate of EWTR.

This approach ties our model parameter most directly to empirical counterparts, but it implicitly assumes that analysts' forecast dispersion is representative of dispersion in investors' private beliefs. However, analyst forecasts are in the public domain and thus not part of private information sets. This distinction doesn't play a major role in heterogeneous priors models where investors' beliefs may disagree about public information, but it does for models of noisy information aggregation under a common prior, which precludes public information as a source of dispersion in beliefs.[32] One possible resolution is to argue that public disagreement in analyst forecasts is broadly representative of dispersion in private investor forecasts: as long as the former is positively correlated with the latter, the qualitative predictions discussed above will remain valid.

Using analyst's earnings forecasts to quantify EWTR requires in addition that the quantitative magnitudes are comparable. If one is willing to accept that analysts forecasts are more precise and less dispersed than investor forecasts, then the measures drawn from analyst forecast dispersion represents a lower bound on the overall magnitude of EWTR. Alternatively, we may assume that analysts are representative of the wider investor pool and treat survey forecasts as a noisy finite sample of private investor expectations that are publicly disclosed to the market. We formally develop this interpretation in the online appendix and show that we can use equation (14) with minor adaptations for noisy public signals along with the sample estimates of forecast dispersion and accuracy to infer EWTR in equity markets.[33]

Table 1 reports the mean, the mid-point of each quintile bin of the distributions of forecast dispersion, and forecast accuracy and EWTR across firms from the GH sample. These objects are highly skewed: for most firms, forecast dispersion is low and accuracy is fairly high, so the implied

[31] We thank these authors, as well as Ludwig Straub and Robert Ulbricht, for sharing the data used in this section.

[32] This concern arises whenever public survey expectations are used to estimate dispersion in private beliefs, as virtually all papers cited here do.

[33] Kovbasyuk and Pagano (2022) discuss anecdotal and empirical evidence suggesting that analysts have an incentive to disclose their information after taking certain positions to realize gains from informed trading, consistent with our view that analyst forecast disagreement as a proxy measure for forecast dispersion among a wider pool of investors.

	Mean	10%	30%	Median	70%	90%
Dispersion ($\tilde{D}$)	**0.19**	0.05	0.11	**0.17**	0.24	0.36
Accuracy ($\hat{\gamma}$)	**0.75**	0.395	0.673	**0.817**	0.917	0.976
EWTR ($\frac{\hat{\sigma}_\theta}{\sigma_\theta}$)	**1.10**	1.002	1.008	**1.022**	1.053	1.192

Table 1: Forecast dispersion, accuracy, and implied EWTR

EWTR is very small. However, both forecast dispersion and EWTR is fairly significant in the top quintile of the distribution –while the 90th percentile of dispersion is about twice the mean, EWTR is up to nine times as high. As a ballpark estimate, the data suggest an average EWTR of about 10%, but most of this average is driven by the top quintile where EWTR is close to 20%. For robustness, we replicate our estimates on a second sample from Straub and Ulbricht (2015), who use the entire I.B.E.S. sample (1976-2016) and forecasts over a longer 8 month horizon. We further restrict our sample to a subset of 2,103 firms which have at least 10 years of forecast data, which substantially reduces the noise in measuring forecast dispersion and accuracy. The distribution of forecast dispersion, accuracy and EWTR obtained is qualitatively similar to the GH sample, with a highly skewed distribution where most variation is concentrated in the top quintile. The fact that two substantially different data sets deliver qualitatively similar distributions of forecast dispersion and EWTR gives us some confidence in the robustness of our numerical examples.

Model-implied returns to skewness and disagreement:

We now compute model-implied returns to skewness and disagreement for EWTR in the range reported in Table 1. We define a parametric asset return function $\pi(\theta) = e^{kx(\theta)}$ such that $x(\cdot)$ follows a beta distribution, setting the key parameters to match target values for expected skewness and volatility at the firm level. We then vary informational parameters to generate levels of forecast dispersion and EWTR in line with the distributions from GH (Table 1). Concretely, we set $\hat{\gamma}$ to the sample mean of 0.75, and we then vary $\tilde{D}$ to match the mean friction, as well as the 10th, 30th, 50th, 70th and 90th percentile of the distribution of forecast dispersion.[34]

Table 2 compares the empirical and model-implied returns for securities sorted into quintiles by skewness (in rows) and forecast dispersion (in columns). The row and column labeled "Targets" reports forecast dispersion (in the row) from Table 1 and skewness (in the column) and from BMV (Table 3, column 4). The row and the column labeled "Returns" report the corresponding average

[34]The challenge is to identify targets for expected return skewness and volatility for firms matching the distribution of earnings forecasts. Since DMS or GH do not report these moments, we calibrate our returns to match the mean realized firm-level return skewness and volatility reported in BMV (Table 3, Columns 4 and 5).

			Mean	Q1	Q2	Q3	Q4	Q5	Q5-Q1
	Targets	*Disp (GH)*	*0.315*	*0.05*	*0.11*	*0.17*	*0.24*	*0.36*	
	Skew (BMV)	**Returns**	**GH**	**0**	**-10**	**-18**	**-27**	**-26**	**-26**
Mean	*0.851*	**BMV**	*-9.5*	*-0.1*	*-0.6*	*-1.5*	*-3.8*	*-18.7*	*-18.6*
Q1	*0.167*	**0**	*-0.8*	*-0.0*	*-0.0*	*-0.1*	*-0.3*	*-1.5*	*-1.5*
Q2	*0.375*	**-7**	*-2.3*	*-0.0*	*-0.1*	*-0.4*	*-0.9*	*-4.6*	*-4.6*
Q3	*0.565*	**-9**	*-4.7*	*-0.0*	*-0.3*	*-0.8*	*-1.8*	*-9.2*	*-9.2*
Q4	*0.809*	**-13**	*-8.6*	*-0.1*	*-0.5*	*-1.4*	*-3.4*	*-16.9*	*-16.8*
Q5	*1.629*	**-67**	*-25.8*	*-0.3*	*-1.6*	*-4.2*	*-10.3*	*-50.7*	*-50.4*
Q5-Q1		**-67**	*-25*	*-0.3*	*-1.6*	*-4.1*	*-10.0*	*-49.2*	*-48.9*

Table 2: Returns to skewness and disagreement (model vs data)

returns (in basis points,bp) from the two empirical studies from BMV (Table 3, column 1) in the column and GH (Table 12, Panel A) in the row; in both cases we report excess returns of stocks in Q2-Q5 over stocks in Q1 to focus on the variation in returns and improve the comparability between model and data. The subsequent rows and columns then report model-implied returns for different levels of skewness and forecast dispersion, while the final row and column report the difference in returns between top and bottom quintiles.

To quantify the implications of varying skewness, the column labeled "mean" sets the level of forecast dispersion to target a mean level of information frictions equal to $\frac{\hat{\sigma}_\theta}{\sigma_\theta} = 1.1$, and varies the level of skewness across rows. Our model generates returns to skewness of about -25 bp per month between the top and bottom quintiles, corresponding to 37% of the skewness premium of -67 bp reported in BMV.[35] The subsequent columns repeat the same exercise varying the level of forecast dispersion across quintiles, which yield skewness premia that range from -0.3 up to -50 bp in the top quintile of the distribution. Analogously, to quantify the effects of varying dispersion, the row labeled "mean" sets the level of idiosyncratic skewness and volatility to the sample mean reported by BMV (Table 1) and varies forecast dispersion across quintiles. Our model generates returns to disagreement of about -19 bp, or 72% of the returns to disagreement of reported for equity in GH (26 bp). The subsequent rows repeat the same exercise varying the level of skewness across quintiles, which yield returns to disagreement from less than 2 bp in the lowest quintile all the way up to 50 bp for the highest.

[35] The model accounts for between 30% and 60% of excess returns for the other skewness quintiles.

The table also allows to illustrate interaction effects. Variation in returns to disagreement from the bottom to top skewness quintile (-49 bp), as well as the variation in returns to skewness from the bottom to the top disagreement quintiles (-50 bp) are both significant and correspond to about 6.2% annual returns. We do not have direct empirical counterparts for the joint variation of returns with skewness and disagreement, but Yu (2011) reports that annualized returns to disagreement vary by about 7% between the highest and lowest quintiles of book-to-market value, and the value premium varies by a similar amount between high and low disagreement terciles. DMS (Table IV) report similar magnitudes of variation in returns to disagreement with changes in book-to-market value and market capitalization.

Overall, and with the caveats that apply to a simple calibration, our results suggest model-implied returns to skewness and forecast dispersion may explain a relevant fraction of those documented empirically, hopefully inviting further studies on these important issues in asset pricing.

4 Generalizing the risk-neutral model

We now generalize the equilibrium characterization and comparative statics from section 3 to the general model set-up of section 2. Suppose that there exists a *price-monotone equilibrium* $\{P(\theta,s); d(x,P); H(\cdot|P)\}$ in which $d(x,P)$ is strictly decreasing in P for $d(x,P) \in (d_L, d_H)$.[36] Fix any $\bar{D} \in (d_L, d_H)$ and define $z \equiv z(P)$ as the private signal of an informed investor who finds it optimal to hold $\bar{D}$ units at price P. $z(P)$ is implicitly defined by $d(z,P) = \bar{D}$. Since $d(x,P)$ is strictly increasing in x, $z(P)$ is strictly increasing in P, and is therefore a sufficient statistic for the information conveyed in the price. By inverting $z(P)$, we can represent the price as a function of z only. In addition, we can construct posterior beliefs directly from the market-clearing condition. Since aggregate demand $D(\theta, P)$ is decreasing in P, we have $Pr(P \le P'|\theta) = Pr(D(\theta,P) \ge D(\theta,P')) = Pr(s \ge D(\theta,P')) = 1 - G(D(\theta,P'))$. Therefore conditional on θ, z is distributed according to

$$\Psi(z|\theta) = 1 - G(D(\theta, P_\pi(z))). \tag{15}$$

Together with the prior $H(\cdot)$, equation (15) defines the joint distribution of P and θ, from which we derive the posterior $H(\cdot|P)$ using Bayes' Rule whenever applicable. These observations lead to the following theorem:

Theorem 1 *: For any price-monotone equilibrium* $\{P(\theta,s); d(x,P); H(\theta|P)\}$*, and any* $\bar{D} \in (d_L, d_H)$*, there exists a sufficient statistic* $z = z(\theta, s)$*, with cdf given by (15), such that the price function*

[36] To our knowledge, no equilibrium existence results are available for this general class of models.

takes the form $P(\theta, s) = P_\pi(z(\theta, s))$*, where* $P_\pi(\cdot)$ *satisfies*

$$P_\pi(z) = \frac{\mathbb{E}\left(U'\left(\bar{D}\left(\pi(\theta) - P_\pi(z)\right)\right) \cdot \pi(\theta) \,|x = z, z\right)}{\mathbb{E}\left(U'\left(\bar{D}\left(\pi(\theta) - P_\pi(z)\right)\right) |x = z, z\right)}. \tag{16}$$

Theorem 1 generalizes the *sufficient statistic representation* of Proposition 2 for any price-monotone equilibrium. For each $\bar{D} \in (d_L, d_H)$, there exists a state variable z, function of θ and s only, such that the price is represented as the risk-adjusted expectation of dividends of an investor who finds it optimal to hold exactly $\bar{D}$ units of the asset when the state is z.[37]

Equation (16) generalizes the updating wedge discussed in the context of the CARA-normal, CARA-binary and risk-neutral examples. The risk-neutral, conditional expectation of dividends processes the price signal twice, once as a public price signal, and once as the private signal of the threshold investor who purchases $\bar{D}$ units of the asset. The intuition for this characterization is as before: shifts in fundamentals or noise trading result in price adjustment, due to market-clearing, over and above the mere information content of the price. In the price expression, these effects are represented by the sufficient statistic z appearing twice in the conditioning set, once through the price signal, and once through the marginal investor's private information. This wedge between the market expectation of dividends and the Bayesian posterior is thus a necessary characteristic of *any* model with noisy information aggregation through asset prices.

Theorem 1 only offers a partial equilibrium characterization: to fully characterize asset valuations, we still need to compute, for some $\bar{D}$, the distribution of the associated sufficient statistic z. This distribution, however, derives from the market clearing condition $D(\theta, P) = s$, which still requires information about the entire demand schedule. Nevertheless, Theorem 1 allows us to develop implications for asset prices and returns through a risk-neutral representation of the price. Specifically, equation (16) can be rewritten as $P_\pi(z) = \mathbb{E}(\pi(\theta) m(\theta, z)|z)$, where

$$m(\theta, z) = \frac{U'\left(\bar{D}\left(\pi(\theta) - P_\pi(z)\right)\right)}{\mathbb{E}\left(U'\left(\bar{D}\left(\pi(\theta) - P_\pi(z)\right)\right) m^I(\theta, z)\,|z\right)} m^I(\theta, z),$$

and $m^I(\theta, z) = \frac{h(\theta|x=z,z)}{h(\theta|z)} = \frac{f(z-\theta)}{\mathbb{E}(f(z-\theta')|z)}$. Since $\mathbb{E}(m(\theta, z)|z) = 1$, the asset price admits a risk-neutral representation, where the pricing kernel $m(\theta, z)$ can be decomposed into a risk adjustment $\frac{U'\left(\bar{D}(\pi(\theta)-P_\pi(z))\right)}{\mathbb{E}\left(U'\left(\bar{D}(\pi(\theta)-P_\pi(z))\right)m^I(\theta,z)|z\right)}$ that weighs states according to the investor's marginal utility of

[37] The representation in theorem 1 depends on the choice of $\bar{D}$, but the representations for different values of $\bar{D}$ are all monotonic transformations of each other. They correspond to different decompositions of the price into expected dividend and risk premium: the higher is $\bar{D}$, the higher is the required risk premium, and hence also the dividend expectation of the investor who holds $\bar{D}$ in equilibrium. It may be natural to set $\bar{D}$ equal to $\mathbb{E}(s)$, so that the risk adjustment accounts for the risk preferences of an investor who holds the unconditional average exposure.

consumption at given exposure $\bar{D}$, and an informational adjustment $m^I(\theta, z)$ that weighs states according to the ratio between the marginal trader's and the objective posterior density. Notice that the first factor vanishes when traders are either approximately risk-neutral ($U'(\cdot)$ is constant) or $\bar{D} = 0$, while the second factor vanishes if private information becomes infinitely noisy (no private information). Equation (16) therefore provides an analogous representation to the "usual" no-arbitrage representation of prices that weighs states according to the marginal investors' attitudes towards risk (the first component in $m(\theta, z)$), and an additional adjustment factor that is new and specific to models with noisy information aggregation.

We can therefore represent the asset price as the conditional dividend expectation under $\hat{H}(\theta|z) \equiv \int_{-\infty}^{\theta} m(\theta, z)\, dH(\cdot|z)$: $P_\pi(z) = \hat{\mathbb{E}}(\pi(\theta)|z) = \int \pi(\theta)\, d\hat{H}(\theta|z)$. As before, we compound $\hat{H}(\cdot|z)$ with the prior over z to define the *risk-neutral probability measure* $\hat{H}(\theta) = \int \hat{H}(\theta|z) d\Psi(z)$, where $\Psi(z) \equiv \int (1 - G(D(\theta, P_\pi(z))))\, dH(\theta)$ denotes the prior cdf of z. Hence, the expected price is represented as the expectation of dividends under the risk-neutral measure $\hat{H}$: $\mathbb{E}(P_\pi(z)) = \int \hat{\mathbb{E}}(\pi(\theta)|z)\, d\Psi(z) = \hat{\mathbb{E}}(\pi(\theta)) = \mathbb{E}(\pi(\theta) \cdot m(\theta))$, where

$$m(\theta) = \frac{\hat{h}(\theta)}{h(\theta)} = \mathbb{E}(m(\theta, z)|\theta),$$

with $\mathbb{E}(m(\theta)) = 1$.[38]

Excess Weight on Tail Risks: We now provide conditions under which $\hat{H}(\cdot)$ overweighs tail realizations of the fundamental. We first show that the risk-neutral measure $\hat{H}(\cdot)$ overweights extreme tail probabilities under a simple regularity condition on the market-implied signal:

Proposition 4 *: Suppose that* $\lim_{z\to-\infty} H(z+k|z) = 0$ *and* $\lim_{z\to\infty} H(z+k|z) = 1$*, for any finite* k*. Then* $\lim_{\theta\to\infty} m(\theta) = \lim_{\theta\to-\infty} m(\theta) = \infty$.

Proposition 4 shows that the upper and lower tail densities of the risk-neutral measure are infinitely thicker than the corresponding prior densities, whenever the updating conditional on the price remains bounded, even in the face of extreme realizations of the sufficient statistic z, or in other words, if z is arbitrarily high (low), the posterior belief assigns probability close to 1 to the

[38] With CARA preferences and absolute risk aversion χ, it is furthermore possible to rewrite equation (16) as $P_\pi(z) = \mathbb{E}^R\left(\pi(\theta)\, m^I(\theta, z)|z\right)$, where expectations $\mathbb{E}^R(\cdot)$ are formed given a risk-adjusted prior with pdf $h^R(\theta) \sim h(\theta)\, e^{-\chi\bar{D}\pi(\theta)}$. In other words, with CARA preferences, it is straight-forward to incorporate the risk adjustment into the prior distribution over θ. The expected price then satisfies $\mathbb{E}(P_\pi(z)) = \mathbb{E}^R(\pi(\theta)) + Cov^R\left(\pi(\theta), m^I(\theta, z)\right)$, and the expected price premium decomposes into a risk adjustment $\mathbb{E}^R(\pi(\theta)) - \mathbb{E}(\pi(\theta))$ and an informational adjustment $Cov^R\left(\pi(\theta), m^I(\theta, z)\right)$. Without the CARA assumption, the informational and risk adjustments are not independent, since z may affect the risk premium through a wealth effect of P on asset demand.

event that the fundamental θ is lower (higher) than any fixed difference k from the realized value of z. This property emerges naturally from Bayesian updating, provided that the prior remains informative (relative to the sufficient statistic) in the tails.

Under the conditions of Proposition 4, the risk neutral measure displays fatter tails than the prior distribution, or $\hat{H}(\theta) > H(\theta)$ for sufficiently low θ and $\hat{H}(\theta) < H(\theta)$ for sufficiently high θ. These conditions are sufficient to establish comparative statics and return predictions based on excess weight on tail risks for securities where risks are concentrated in the extreme tails, like high grade corporate bonds or deeply out of the money options.

However, we need a stronger characterization of excess weight on tail risks to establish comparative statics globally. Specifically, we say that $\hat{H}(\cdot)$ displays *excess weight on tail risks*, if $m(\cdot)$ is log-convex with $\lim_{\theta\to\infty} m(\theta) = \lim_{\theta\to-\infty} m(\theta) = \infty$. We further say that $\hat{H}_1(\cdot)$ has more excess weight on tail risk than $\hat{H}_2(\cdot)$ if $m_1(\theta)/m_2(\theta)$ is log-convex. Log-convexity of $m(\cdot)$ implies that $m(\cdot)$ is U-shaped (rather than, say, W-shaped), and that $\hat{H}(\cdot)$ intersects $H(\cdot)$ exactly once. Log-convexity also identifies the key distinction between the risk-neutral measure under noisy information aggregation and the "usual" risk adjustment for a security in positive net supply, since the latter typically leads to a strictly downwards-sloping pricing kernel (rather than a U-shaped one) to shift probability mass from higher towards lower states.

Suppose that the prior $h(\cdot)$ and the signal density $f(\cdot)$ are strictly log-concave with

$$\bar{\tau}_h \geq -\left(\frac{h''}{h} - \left(\frac{h'}{h}\right)^2\right) \geq \underline{\tau}_h > 0 \text{ and } \bar{\tau}_f \geq -\left(\frac{f''}{f} - \left(\frac{f'}{f}\right)^2\right) \geq \underline{\tau}_f > 0.$$

This assumption imposes that variation in the log-curvature in the two distributions, or in the informativeness of the prior and the private signals, is bounded on both sides. When f and h are normal, then $\frac{h''}{h} - \left(\frac{h'}{h}\right)^2$ and $\frac{f''}{f} - \left(\frac{f'}{f}\right)^2$ are constant with $\bar{\tau}_h = \sigma_\theta^{-2} = \underline{\tau}_h$ and $\bar{\tau}_f = \beta = \underline{\tau}_f$.[39]

Proposition 5 *: Suppose that $\psi(z|\theta) \equiv \psi(z-\theta)$, where $\psi(\cdot)$ is strictly log-concave with $\bar{\tau}_\psi \geq -\left(\frac{\psi''}{\psi} - \left(\frac{\psi'}{\psi}\right)^2\right) \geq \underline{\tau}_\psi > 0$. Define $\underline{\hat{\gamma}} = \frac{\underline{\tau}_f + \underline{\tau}_\psi}{\underline{\tau}_f + \underline{\tau}_\psi + \underline{\tau}_h}$ and $\bar{\gamma} = \frac{\bar{\tau}_\psi}{\bar{\tau}_\psi + \bar{\tau}_h}$, and suppose further that $\underline{\hat{\gamma}}\underline{\tau}_h - \bar{\tau}_h\bar{\gamma} > 0$. Then, $m(\cdot)$ is strictly log-convex and*

$$\frac{d}{d\theta}\frac{m'(\theta)}{m(\theta)} = \frac{d}{d\theta}\left(\frac{\hat{h}'(\theta)}{\hat{h}(\theta)} - \frac{h'(\theta)}{h(\theta)}\right) \geq \frac{\underline{\hat{\gamma}}\left(\underline{\hat{\gamma}} - \frac{\bar{\tau}_h}{\underline{\tau}_h}\bar{\gamma}\right)}{\underline{\hat{\gamma}} - \left(1-\underline{\hat{\gamma}}\right)\left(\underline{\hat{\gamma}} - \frac{\bar{\tau}_h}{\underline{\tau}_h}\bar{\gamma}\right)}\underline{\tau}_h > 0.$$

Moreover, whenever f, ψ, and h converge to normal densities, then $\frac{d}{d\theta}\frac{m'(\theta)}{m(\theta)}$ converges to $\frac{1}{\sigma_\theta^2} - \frac{1}{\hat{\sigma}_\theta^2} > 0$, where $\hat{\sigma}_\theta^2$ is given by equation (11).

[39] See Saumard and Wellner (2014) for a primer to log-concave distributions. Proposition 5 makes use of their proposition 10.1.

Proposition 5 identifies sufficient conditions for log-convexity of the risk-neutral measure, which generalize equation (11) in the linear-normal setting. When z is affine in θ with $\frac{\partial z(\theta,s)}{\partial \theta} = 1$, $\hat{m}$ is log-convex whenever $\varphi(z) \equiv \mathbb{E}(f(z-\theta')|z)$ is log-concave, and the latter is shown to be true whenever $\underline{\hat{\gamma}}\underline{\tau}_h - \bar{\tau}_h\bar{\gamma} > 0$. This expression can be interpreted as a lower bound on the gap between the posterior uncertainty under the objective and the risk-neutral probability measure, just as $\underline{\hat{\gamma}}$ is a lower bound on the combined informativeness of the private and market signal relative to the prior, and $\bar{\gamma}$ an upper bound on the informativeness of the market-signal relative to the prior. The condition that $\underline{\hat{\gamma}}\underline{\tau}_h > \bar{\tau}_h\bar{\gamma}$ thus imposes that private signals are sufficiently informative so that the risk-neutral measure has uniformly lower posterior uncertainty than the objective posterior.

Moreover, equation (11) for the linear-normal model can be restated as $\frac{1}{\sigma_\theta^2} - \frac{1}{\hat{\sigma}_\theta^2} = \frac{\hat{\gamma}(\hat{\gamma}-\gamma)}{\hat{\gamma}-(1-\hat{\gamma})(\hat{\gamma}-\gamma)}\frac{1}{\sigma_\theta^2}$. When all three densities converge to log-quadratic (normal) densities, then $\underline{\hat{\gamma}} \to \hat{\gamma}$, $\bar{\gamma} \to \gamma$, and the above bound converges to $\frac{1}{\sigma_\theta^2} - \frac{1}{\hat{\sigma}_\theta^2}$. Hence if the densities are approximately Gaussian, then equation (11) provides a good approximation to the extent of excess weight on tail risks for noisy information aggregation models that are "close" to the linear-normal model used in the preceding examples, and more generally the same equation can be used to construct a lower bound on excess weight on tail risks by using upper and lower bounds on the informativeness of signals in each state.

The proposition relies on three conditions: first, log-concavity of the densities h, f, and ψ insures that agents update monotonically from both the private signal and the price signal, and therefore posterior beliefs are monotone in the signal realization.

Second, we require a bound on the variation in log-curvature of the signal densities, or equivalently a lower bound on the precision of trader's private signals. Variation in log-curvature implies that posterior uncertainty may vary across states and signal realizations. The additional bound insures that such variation in posterior uncertainty cannot become too important (or equivalently, private signals are sufficiently informative) so that the risk-neutral posterior displays a uniformly stronger response to variation in z than the objective posterior throughout the state space.

Third, the condition that $\frac{\partial z(\theta,s)}{\partial \theta} = 1$ for all (θ, s), or equivalently, that $\psi(z|\theta) = \psi(z-\theta)$ for all (z,θ), imposes that the sufficient statistic variable takes the canonical form of "fundamental plus noise". To interpret this condition, notice that differentiating the market-clearing condition $D(\theta, P) = s$ with respect to θ, we obtain

$$\frac{D_\theta(\theta, P_\pi(z))}{D_P(\theta, P_\pi(z))} = -\frac{\partial P_\pi(z)}{\partial z}\frac{\partial z(\theta, s)}{\partial \theta}.$$

In general, $\frac{\partial P_\pi(z)}{\partial z} = -d_x(z,P)/d_P(z,P)$ measures the rate at which the marginal investor trades off between higher price and higher dividend expectation, while $-D_\theta/D_P$ represents the same marginal rate of substitution for aggregate demand, or investors on average. Additive separability then obtains whenever the marginal and average investors' marginal rates of substitution coincide. Departures from this benchmark require that $-d_x(x, P)/d_P(x, P)$ varies with x, and that this variation does not wash out through aggregation.

Alternatively, we may replace the assumption that $\psi(z|\theta) = \psi(z-\theta)$ for all (z, θ) with stronger assumptions on the bounds to variation in log-curvature of the signal distributions. Intuitively speaking, these bounds impose that departures from the canonical benchmark are not too large so as to change the sign of the bound constructed in proposition 5.[40]

To summarize, proposition 5 shows that EWTR emerges naturally if agents' posteriors satisfy a monotone likelihood ratio property with regards to private signal realizations and prices, the informativeness of the prior and the signals doesn't vary too much over the state space, and the sufficient statistic is reasonably close to the canonical "fundamental plus noise" structure. While the former amounts to regularity conditions on the prior and the private signal densities, the latter imposes restrictions on the endogenous distribution of the sufficient statistic which we unfortunately have not been able to translate into conditions on exogenous primitives. Nevertheless, they clarify in what sense the results from the risk-neutral and normal updating models can be expected to generalize. Alternatively, we can prove log-convexity of $\hat{m}(\cdot)$ and EWTR by invoking other restrictions on primitive parameters, for example in limiting cases with either very large or very small supply noise and private signal precisions. The online appendix provides further details.

Generalizing Proposition* 3*: For a given change in probability measure $\hat{m}$, we then write the expected price premium as

$$W(\pi, \hat{m}) \equiv \mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta)) = cov(\pi(\theta), m(\theta)).$$

If $H(\cdot)$ and $\hat{m}(\cdot)$ are symmetric and centered around the same mean (say, around 0), we can then directly generalize the comparative statics predictions of Proposition 3, applying the partial order on returns given by Definition 1:

Theorem 2 *Suppose that $m_1(\cdot)$ and $m_2(\cdot)/m_1(\cdot)$ are symmetric around 0 and log-convex, and $H(\cdot)$ is symmetric around 0.*

(i) ***Comparative Statics w.r.t.*** m***:*** *If π has symmetric risk, then $W(\pi; m) = 0$. If π is dominated by upside (downside) risk, then $W(\pi_1, m_2) \geq W(\pi_1, m_1) \geq 0$ ($W(\pi_1, m_2) \leq W(\pi_1, m_1) \leq$*

[40] Such departures introduce additional terms in the characterization of $\frac{d}{d\theta}\frac{m'(\theta)}{m(\theta)}$, which vanish when $\psi(z|\theta) \equiv \psi(z-\theta)$.

0). *Moreover,* $\lim_{m\to 1} W(\pi, m) = 0$, *and* $|W(\pi, m)|$ *grows arbitrarily large if* m *has arbitrarily large excess weight on tail risk and* $lim_{\theta\to\infty}|\pi(\theta) + \pi(-\theta)| = \infty$.

(*ii*) ***Comparative Statics w.r.t.*** π ***and Increasing differences:*** *If* π_2 *has more upside risk than* π_1, *then* $W(\pi_2, m_2) - W(\pi_1, m_2) \geq W(\pi_2, m_1) - W(\pi_1, m_1) \geq 0$.

This generalization relies on the symmetry and equal means assumption that were already embedded in the risk-neutral normal model. We can still obtain variants of the comparative statics in theorem 2 upon relaxing these assumptions. With equal means but asymmetric distributions, the theorem continues to hold if the partial order on upside and downside risks is restricted to payoff functions that are strictly concave or strictly convex.

If means are not equal, then the expected price premium responds to both a shift in means and a mean-preserving spread in the risk-neutral distribution (Section 3). The comparative statics of theorem 2 then continue to hold if these two shifts are mutually reinforcing, or equivalently (in the notation used above), if $\int \theta d\hat{H}_2(\theta) \geq \int \theta d\hat{H}_1(\theta) \geq \int \theta dH(\theta)$ for convex return functions (upside risks) or $\int \theta d\hat{H}_2(\theta) \leq \int \theta d\hat{H}_1(\theta) \leq \int \theta dH(\theta)$ for convex return functions (downside risks).

Alternatively, we can decompose the expected premium $\mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta))$ into a shift in means and a mean-preserving spread:

$$\mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta)) = \int_{-\infty}^{\infty} (\pi(\theta) - \pi(\theta - \delta))\, d\hat{H}(\theta) + \int_{-\infty}^{\infty} \left(H(\theta) - \hat{H}(\theta + \delta)\right) d\pi(\theta)$$

where $\delta \equiv \int \theta d\hat{H}(\theta) - \int \theta dH(\theta) = cov(\theta, \hat{m}(\theta))$. The shift in means $\int_{-\infty}^{\infty} (\pi(\theta) - \pi(\theta - \delta))\, d\hat{H}(\theta)$ varies with the expected asset supply: for a given distribution of dividends, a first-order stochastic increase in the supply distribution $G(\cdot)$ requires that informed investors buy more shares in equilibrium, which lowers the marginal investor's z. This downwards shift in the price distribution is captured by a decrease in δ. The second-order shift in the distribution $\int_{-\infty}^{\infty} \left(H(\theta) - \hat{H}(\theta + \delta)\right) d\pi(\theta)$ instead captures the excess weight on tail risks implied by the risk-neutral distribution, controlling for the difference in means. We can then apply the above comparative statics results to this second term, under the assumption that $\hat{H}(\theta + \delta)$ has EWTR over $H(\theta)$.

Finally, without log-convexity, the condition of proposition 4 implies that there exists θ_L and $\theta_H > \theta_L$ such that comparative statics from theorem 2 continue to apply to securities for which $\pi(\theta_L) = \pi(\theta_H)$, i.e. all the variation in returns is concentrated in the tails $\theta \leq \theta_L$ and $\theta \geq \theta_H$.

Numerical solution methods: The generalized equilibrium characterization described in this section can also be used to develop a new procedure to solve noisy REE equilibrium with general preferences numerically, by iterating over the information content of prices. This is of interest for two reason. First, it provides a method for verifying, at least numerically, whether the model-implied risk-neutral measure satisfies the sufficient conditions for excess weight on tail risks described in proposition 4; hence short of having explicit sufficient conditions for EWTR, these conditions can at least be verified numerically for a given set of primitives. Second, the lack of sharp equilibrium characterizations, except in special cases that are analytically solvable, has been a long-standing challenge to bring information aggregation models closer to standard preferences used in finance and to derive asset pricing predictions.[41]

Fix a support of the fundamental θ and prior $H(\cdot)$. We start conjecturing a distribution of prices conditional on a given value of θ: $\Psi^{(0)}(P'|\theta) \equiv Pr(P \leq P'|\theta)$, along with a conditional density $\psi^{(0)}(P|\theta)$. From $\psi^{(0)}(P|\theta)$, we calculate the posterior distribution for each investor using Bayes rule: $Pr^{(0)}(\theta|x_i, P) = \psi^{(0)}(P|\theta) \cdot Pr(\theta|x_i) / \sum_{\theta'} \psi^{(0)}(P|\theta') \cdot Pr(\theta'|x_i)$, where $Pr(\theta|x_i)$ corresponds to the posterior conditional on observing x_i only. Using the posterior distribution, we find optimal demand functions $d^{(0)}(x, P)$, and then determine aggregate demand $D^{(0)}(\theta, P)$ numerically by integrating over x. Using the market-clearing condition, we then characterize the resulting informational content of prices: $\Psi^{(1)}(P'|\theta) \equiv 1 - G(D^{(0)}(\theta, P'))$. This new conditional price distribution $\Psi^{(1)}$ is used then as the starting guess in place of $\Psi^{(0)}$, and the exercise is iterated until convergence. Finally, we calculate the price function $P(\theta, s)$ by inverting the function $D(\theta, P) = s$ to obtain $P = P(\theta, s = D)$.

In the online appendix, we apply this method to a model with binary payoffs and CRRA preferences. The numerical results confirm the analytical results for the CARA case in section 2 and the risk-neutral model in section 3: downside risk is on average under-priced, upside risk is over-priced.[42]

[41] Bernardo and Judd (2000) and Peress (2004) numerically solve a RE equilibrium under asymmetric information and CRRA preferences by "guessing" price and demand functions using hermite polynomials under the structural moment conditions implied by demand optimality and market clearing. The central difference of our approach is that here we explicitly solve for the price likelihood function, which allows a clean characterization of the informational content of prices, for different price realizations. To our knowledge, this methodology is new in the REE literature.

[42] We make our matlab code available for the CRRA, binary payoff case solved in the online appendix. Under generic preferences and payoff structures, aggregate demand monotonicity w.r.t. prices is not guaranteed (see for example, Barlevy and Veronesi (2003)). Without strict monotonicity, it is no longer true that $Pr(P \leq P'|\theta) = Pr(u \geq D(\theta, P))$ for any price level, and the solution method proposed here would not work. The example presented in this section uses parameters which satisfy monotonicity of demand as a function of the price.

5 Multiple assets and uninformed traders

We consider extensions of our model to multiple assets and uninformed, as well as informed traders. Within this context, we discuss several results.

First, we discuss conditions under which our construction of the risk-neutral probability measure based on noisy information aggregation generalizes to multi-asset models. We illustrate this characterization with two polar cases, one in which securities are linked to the same fundamental (i.e. stocks and bonds issued by the same firm, or different options on the same underlying security), and one in which securities are completely independent, and each market can be analyzed in isolation.

Second, we use the latter representation to show that stochastic risk premia due to noisy supply shocks are sufficient to generate excess weight on tail risks, but this is then amplified by the presence of informed traders and noisy information aggregation. In the process, we also generalize the decomposition into preference-based and information-based skewness premia that we discussed in the context of the CARA-normal and CARA-binary examples.

General set-up: Suppose that there are two securities with stochastic payoff $\pi_n(\cdot)$, with $n \in \{1,2\}$.[43] Their supply is stochastic with $s \equiv (s_1, s_2) \in [d_{L,1}, d_{H,1}] \times [d_{L,2}, d_{H,2}]$ distributed according to a smooth cdf. G.[44] Security n is conditioned on a fundamental θ_n, and the vector of fundamentals $\theta = (\theta_1, \theta_2) \in \mathbb{R}^2$ is distributed according to a common prior $H(\cdot)$. We suppose that a measure $\kappa_I > 0$ of traders are informed, and a measure $\kappa_U > 0$ are uninformed. The informed traders receive two private signals $x_n^i = \theta_n + \varepsilon_n^i$, where ε_1^i and ε_2^i are iid across agents i, but may be correlated across securities.

We assume throughout that traders' preferences are as in Section 4. Let P_n denote the price of asset n, $d_n^i \in [d_{L,n}, d_{H,n}]$ the position of trader i in security n (as before, we impose that positions are bounded), and π_n the realized payoff. Then the trader's realized utility over gains and losses is given by $U\left(\sum_{n\in\{1,2\}} d_n^i(\pi_n - P_n)\right)$. A *Perfect Bayesian Equilibrium* consists of a set of demand functions $\left(d^I(x,P), d^U(P)\right) \equiv \left\{d_n^I(x,P), d_n^U(P)\right\}_{n=1,2}$ for informed and uninformed traders and for each security, a set of price functions $P(\theta,s) \equiv \{P_n(\theta,s)\}_{n=1,2}$, and posterior beliefs $H(\cdot|P)$ such that (i) $\left(d^I(x,P), d^U(P)\right)$ is optimal given $H(\cdot|x,P)$ and $H(\cdot|P)$; (ii) $P(\theta,s)$ clears the market for all n; and (iii) $H(\cdot|P)$ satisfies Bayes' rule whenever applicable,

[43] The restrictions to two assets is convenient for exposition; extensions to more than two securities are immediate.

[44] Throughout this section, we let $y = (y_1, y_2)$ where y_n, with $n \in \{1,2\}$, denotes a random variable specific to market n. Hence s refers to the vector of supply realizations, $\pi = (\pi_1, \pi_2)$ to the vector of dividends, $P = (P_1, P_2)$ to the vector of prices etc.

i.e., for all P such that $\{(\theta, s) : P(\theta, s) = P\}$ is non-empty. We focus on *price-monotone equilibria* $\{P(\theta, s); d^I(x, P), d^U(P); H(\cdot|P)\}$ in which $d_n^I(x, P)$ and $d_n^U(P)$ decreasing in P_n whenever $d_n^I(x, P), d_n^U(P) \in (d_{L,n}, d_{H,n})$.

Suppose that there exists a unique vector $z(P)$ such that $d^I(z, P) = 0$, and that this vector is invertible, so that we can write $P(\theta, s) = P(z)$.[45] Using the equilibrium price function, we write $z(\theta, s) \equiv z(P(\theta, s))$ as the sufficient statistic vector. Using the informed trader's first-order condition, the equilibrium price for security $k \in \{1, 2\}$ satisfies

$$P_k(z) = \mathbb{E}(\pi_k(\theta_k)|x = z, z)$$

i.e. we obtain a vector representation of the equilibrium price with the updating wedge for informed traders. As before we can thus represent the equilibrium price as the dividend expectation of a marginal trader who chooses not to hold any of the assets, or

$$P_k(z) = \mathbb{E}\left(m^I(\theta, z)\pi_k(\theta_k)|z\right), \text{ where } m^I(\theta, z) = \frac{h(\theta|x = z, z)}{h(\theta|z)}.$$

The expected price premium then takes the form $\mathbb{E}(P(z)) - \mathbb{E}(\pi(\theta)) = Cov\left(\pi(\theta), m^I(\theta)\right)$, where $m^I(\theta) = \mathbb{E}\left(m^I(\theta, z)|\theta\right) = \mathbb{E}\left(\frac{h(\theta|x=z,z)}{h(\theta|z)}|\theta\right)$ and $Cov\left(\pi(\theta), m^I(\theta)\right)$ represents the vector of covariances of $\pi_k(\cdot)$ with $m^I(\cdot)$. These results directly extend the characterization of the updating wedge and the risk-neutral probability measures to multiple securities.

Therefore the key challenge in generalizing the previous analysis to multiple securities is to show that (i) there exists a marginal investor who finds it optimal to not hold any risky assets, and (ii) that the mapping from this marginal investor's private signal to the equilibrium price is invertible. If these conditions are satisfied, we can use this signal vector as a sufficient statistic for the information content of the price vector and the price vector inherits the multi-dimensional analogue of the updating wedge.

But the coexistence of informed and uninformed traders also allows us to derive an alternative equivalent representation of equilibrium asset prices based on the uninformed traders' demand. Specifically, the first-order condition for asset $k \in \{1, 2\}$ by uninformed traders yields:

$$P_k = \frac{\mathbb{E}\left(U'\left(\sum_{n\in\{1,2\}} d_n^U(P)(\pi_n - P_n)\right)\pi_k(\theta_k)|P\right)}{\mathbb{E}\left(U'\left(\sum_{n\in\{1,2\}} d_n^U(P)(\pi_n - P_n)\right)|P\right)} = \mathbb{E}\left(m^U(\theta, P)\pi_k(\theta_k)|P\right)$$

[45] The characterization below generalizes to arbitrary $\bar{D} \in [d_{L,1}, d_{H,1}] \times [d_{L,2}, d_{H,2}]$ along the same lines as in the single-asset case. In that case, the condition $d^I(z, P) = \bar{D}$ must be invertible, and the resulting pricing kernel decomposes into a risk adjustment for given $\bar{D}$ and an adjustment term due to noisy information aggregation. The present case abstracts from the former by setting $\bar{D} = 0$.

where $m^U(\theta, P) = \frac{U'\left(\sum_{n\in\{1,2\}} d_n^U(P)(\pi_n - P_n)\right)}{\mathbb{E}\left(U'\left(\sum_{n\in\{1,2\}} d_n^U(P)(\pi_n - P_n)\right)|P\right)}$. Hence, knowing the exposures of uninformed traders $d^U(P)$ is sufficient to construct a risk-neutral probability measure to price all securities.

Two polar cases: We now discuss two polar cases. In the first, the two securities are conditioned on the same fundamental θ, and informed traders receive a single noisy signal x; formally, θ_1 and θ_2, as well as ε_1^i and ε_2^i are identically distributed and perfectly correlated. This case may represent simultaneous trading of different options on the same underlying security, or stocks and bonds issued by the same company. In the second case, the two securities are independent: fundamentals, supply realizations and private signals are independent across securities. This case may correspond to a situation where traders invest in multiple independent securities, such as shares of different companies.

1. Common fundamentals: Suppose that fundamentals and signal noise are perfectly correlated, so that $\theta_1 = \theta_2 = \theta$ and $\varepsilon_1^i = \varepsilon_2^i = \varepsilon^i$. In this case, since private information is one-dimensional, there exists a unique $z_k(P)$, and a unique representation

$$P_k(z) = \mathbb{E}\left(\pi_k(\theta_k)\,|x = z_k, z\right)$$

of the equilibrium price from the perspective of the informed traders. Here, we recover that the main spill-over from market 1 into market 2 is informational: the equilibrium price in market 1 depends on the sufficient statistic from market 2 only through the information it conveys about the common fundamental θ. On the other hand, the sufficient statistic from market 1 affects the price in market 1 both for informational reasons, and because of the market-clearing effect we discussed for the single-asset model. Hence, we obtain in this case an asset-specific updating wedge from the informed trader's demand function.[46]

In earlier versions of this paper, we solved a two-asset version of the risk-neutral model with common fundamentals, and showed that differential levels of information frictions in two markets can give rise to a novel rationale for departures from the Modigliani-Miller theorem: by separating upside from downside risks into equity and debt claims, the owner of a cashflow $\pi(\cdot)$ can take advantage of different degrees of information frictions in the two markets. These results are available in the online appendix.

2. Independent assets: Suppose now that θ_1 and θ_2, s_1 and s_2, ε_1^i and ε_2^i are independent across securities. Suppose also that traders have CARA preferences with absolute risk aversion χ. Then there exists an equilibrium in which demand and the price in market 1 are independent from

[46] Chabakauri et al (2021) solve a multi-asset noisy REE model with a finite set of states and complete set of Arrow-Debreu securities using similar insights.

market 2 and vice versa. The equilibrium price for security $k \in \{1,2\}$ satisfies the one-dimensional representations

$$P_k(z_k) = \mathbb{E}\left(m_k^U(\theta_k, z_k)\,\pi_k(\theta_k)\,|z_k\right) = \mathbb{E}\left(m_k^I(\theta_k, z_k)\,\pi_k(\theta_k)\,|z_k\right)$$

where $m_k^U(\theta_k, z_k) = \frac{U'\left(d_k^U(P_k(z_k))\pi_k\right)}{\mathbb{E}\left(U'\left(d_k^U(P_k(z_k))\pi_k\right)|z_k\right)}$ and $m_k^I(\theta_k, z_k) = \frac{h(\theta_k|x_k=z_k,z_k)}{h(\theta_k|z_k)}$, i.e. we recover asset-by-asset the same representation as in the single-asset model. Spill-overs from one asset market to another potentially occur because of information effects and wealth effects. Information effects arise if prices in market 1 convey information about likely returns in market 2. The independence assumption implies that markets are orthogonal to each other, thus eliminating informational spill-overs. Concretely this manifests itself in a multiplicative decomposition of the informed trader's pricing kernel: $m^I(\theta, z) = \prod_{k\in\{1,2\}} m_k^I(\theta_k, z_k)$. Wealth effects arise if prospective gains or losses in market 1 affect the investor's appetite for risk in market 2. The CARA assumption implies that trading in market 1 is independent of gains or losses in market 2. This appears through a multiplicative decomposition of the uninformed traders' pricing kernel: $m^U(\theta, z) = \prod_{k\in\{1,2\}} m_k^U(\theta_k, z_k)$. In Albagli et al. (2017), we solve a variant of the risk-neutral model with a continuum of independent securities (interpreted as equity shares in different firms) to study the interplay between information aggregation and shareholder risk-taking incentives.

Expected Price Premia: Let $C(z) \equiv \mathbb{E}\left(\sum_{n\in\{1,2\}} d_n^U(P)(\pi_n - P_n)\,|z\right)$ denote the uninformed traders' expected portfolio payoff conditional on z. A second-order Taylor expansion of the uninformed traders' first-order condition yields the following approximation of the equilibrium price for security $n \in \{1,2\}$:

$$P_n(z) \approx \mathbb{E}(\pi_n(\theta)\,|z) - \chi(z)\cdot e_n'\Sigma(z)\,d^U(P(z)) + \frac{\alpha(z)}{2}\cdot d^U(P(z))'\,\Psi^n(z)\,d^U(P(z))$$

where $\chi(z) = -\frac{U''(C(z))}{U'(C(z))}$ denotes the absolute risk aversion coefficient, $\alpha(z) = \frac{U'''(C(z))}{U'(C(z))} > 0$ represents a local measure of downside risk aversion,[47] both evaluated at the expected portfolio payoff conditional on z. $e_1' = (1,0)$ and $e_2' = (0,1)$ represent the n-th dimension unit vectors, $\Sigma(z)$

$$\Sigma(z) = \begin{pmatrix} Var(\pi_1(\theta)\,|z) & Cov(\pi_1(\theta), \pi_2(\theta)\,|z) \\ Cov(\pi_1(\theta), \pi_2(\theta)\,|z) & Var(\pi_2(\theta)\,|z) \end{pmatrix},$$

the 2x2 variance-covariance matrix of expected returns, and $\Psi^n(z)$ the n-th third-moment matrix with k, l-th entries

$$\psi_{(k,l)}^n(z) = \mathbb{E}\left((\pi_n(\theta) - \mathbb{E}(\pi_n(\theta)\,|z))(\pi_k(\theta) - \mathbb{E}(\pi_k(\theta)\,|z))(\pi_l(\theta) - \mathbb{E}(\pi_l(\theta)\,|z))\,|z\right).$$

[47] See, e.g. Modica and Scarsini (2005). The measure α can also be represented as the product of prudence $-\frac{U'''(C(z))}{U''(C(z))}$ (Kimball 1990) and risk aversion $\chi(z)$. With CARA preference, $\chi(z) = \chi$ and $\alpha(z) = \chi^2$ are both constants.

The difference between price and expected dividend thus decomposes into a risk adjustment $-\chi(z)\cdot e_n'\Sigma(z)\,d^U(P(z))$ that scales with the uninformed traders' exposure $d^U(P(z))$ and a third-moment adjustment term $\frac{\alpha(z)}{2}\cdot d^U(P(z))'\,\Psi^n(z)\,d^U(P(z))$ that depends on the squares of exposures. The risk adjustment can be rewritten as $e_n'\Sigma(z)\,d^U(P(z)) = Cov\left(\pi_n(\theta),\pi(\theta)\cdot d^U(P(z))\,|z\right)$, where $\pi(\theta)\cdot d^U(P(z))$ represents the uninformed traders' total portfolio return. Abstracting from the second-order (third-moment) terms, the model-implied risk premium thus recovers a standard "CAPM" representation from the perspective of uninformed investors.[48]

Suppose first that $\alpha(z)=0$ (quadratic preferences, no downside risk aversion), or equivalently, ignore the third-moment terms. Then the expected price premium satisfies

$$\mathbb{E}(P_n(z)) - \mathbb{E}(\pi_n(\theta)) \approx -e_n'\mathbb{E}\left(\chi(z)\,\Sigma(z)\cdot d^U(P(z))\right) = -\mathbb{E}\left(\chi(z)\,Cov\left(\pi_n(\theta),\pi(\theta)\cdot d^U(P(z))\,|z\right)\right)$$

$$= -e_n'\mathbb{E}\left(\chi(z)\,\Sigma(z)\right)D - \mathbb{E}\left(\chi(z)\,e_n'\Sigma(z)\cdot\left(d^U(P(z)) - D\right)\right)$$

where $D=\mathbb{E}\left(d^U(P(z))\right)$. The expected price premium thus decomposes into an average risk premium that scales with risk aversion $\chi(z)$, expected posterior uncertainty $\mathbb{E}(\Sigma(z))$, expected exposure D, and an adjustment due to the co-movement between the exposure, uncertainty and risk aversion. When assets are conditionally independent ($Cov(\pi_1(\theta),\pi_2(\theta)\,|z)=0$), this co-movement term reduces to $\mathbb{E}\left(\chi(z)\,e_n'\Sigma(z)\cdot\left(d^U(P(z))-D\right)\right) = Cov\left(\chi(z)\,Var(\pi_n(\theta)\,|z),d_n^U(P(z))\right)$. If exposure $d_n^U(P(z))$ is everywhere decreasing in z, then the co-movement term is positive (negative) if uncertainty $Var(\pi_n(\theta)\,|z)$, scaled by risk aversion $\chi(z)$, is decreasing (increasing) in z. Therefore, controlling for the average exposure, co-movement generates a positive expected price premium if uncertainty times risk aversion and exposure are both counter-cyclical, and a negative premium if uncertainty times risk aversion is pro-cyclical. This is exactly what return asymmetry generates: for downside risks, a deterioration of z increases the likelihood of adverse tail risks, hence uncertainty is countercyclical. For an upside risk the same deterioration of reduces uncertainty as the positive tail event is less likely to materialize.

This observation leads to an alternative interpretation of our results: generalizing the discussion from section 2, the negative relation between skewness and returns results from the combination of (i) counter-cyclical exposure of uninformed traders, and (ii) pro-cyclical (counter-cyclical) uncertainty of upside (downside) risks, giving rise to an *information-based* skewness premium. The counter-cyclical exposure of uninformed traders emerges naturally from the informed traders' de-

[48]See Andrei et al. (2022) for implications of noisy information aggregation for empirical properties of the CAPM with linear/normal asset returns.

mand and the market-clearing condition: since

$$\kappa_I \int d_I\left(x, P\left(z\right)\right) dF(x-\theta) + \kappa_U d^U\left(P\left(z\right)\right) = s$$

an increase in the fundamental vector θ that raises demand $\int d_I\left(x, P_\pi\left(z\right)\right) dF(x-\theta)$ by the informed traders for all securities must be offset by a reduction in the demand by uninformed traders, resulting in lower exposures for uninformed traders when the fundamental is high, or asset supply is low.[49]

Consider next the case where $\alpha\left(z\right) > 0$. With downside risk aversion, the second-order term $\frac{\alpha(z)}{2} \cdot d^U\left(P\left(z\right)\right)' \Psi^n\left(z\right) d^U\left(P\left(z\right)\right)$ multiplies the investors' attitudes towards downside risk with the squared exposures and asymmetries in returns that are summarized by $d^U\left(P\left(z\right)\right)' \Psi^n\left(z\right) d^U\left(P\left(z\right)\right)$. When assets are conditional independent, the latter term reduces to $Skew\left(\pi_n\left(\theta\right)|z\right) \cdot d_n^U\left(P_n\left(z\right)\right)^2$, where $Skew\left(\pi_n\left(\theta\right)|z\right) = \psi_{(n,n)}^n\left(z\right)$ denotes the conditional skewness of asset payoffs. Taking expectations, this term thus generalizes the observation discussed in section 2 that attitudes towards downside risk introduce a *preference-based* skewness premium in asset prices: because of downside risk aversion, traders require additional compensation for accepting the market-clearing exposure level s on negatively skewed securities, while willing to reduce the risk premium for positively skewed securities.

To summarize, the expected price premium decomposes into an *information-based* skewness premium that depends on the co-movement of the uninformed traders' exposure and posterior uncertainty and a *preference-based* skewness premium that depends on downside risk aversion and the conditional skewness of the security. These two terms generalize the observations discussed in the context of the CARA-normal and CARA-binary models.

Amplification of price premia with dispersed information: We conclude by showing that the amplification results and limit results discussed in section 2 also generalize from the examples to the general model: supply shocks alone are sufficient to generate EWTR and preference-based skewness premia, but these are amplified by noisy information aggregation and

[49] Similar results obtain in the risk-neutral model with a noise trader demand of the form $s = \Phi\left(u + \omega\left(P - \mathbb{E}\left(\pi\left(\theta\right)|P\right)\right)\right)$, which captures the notion that the residual supply available to informed traders increases in the expected price premium $P - \mathbb{E}\left(\pi\left(\theta\right)|P\right)$. In this formulation, the higher is $\omega > 0$ the more actively the uninformed traders arbitrage the perceived price premium, with $P \to \mathbb{E}\left(\pi\left(\theta\right)|P\right)$ as $\omega \to \infty$, akin to free entry by uninformed risk-neutral arbitrageurs. A micro-foundation for this functional form assumption about asset supply can be obtained by assuming that *(i)* the asset supply is normalized to 1, and *(ii)* there is a unit measure of risk-neutral uninformed arbitrageurs, who each have a stochastic cost of $c_i = c + u_i$ of holding the one unit of the asset, where $c \sim \mathcal{N}\left(\bar{c}, \sigma_c^2\right)$ and $u_i \sim \mathcal{N}\left(0, \gamma^{-1}\right)$. In this case, uninformed arbitrageurs buy the security if and only if $c_i + P \leq \mathbb{E}\left(\pi\left(\theta\right)|P\right)$, resulting in a residual supply schedule of $\Phi\left(\sqrt{\gamma}\left(c + P - \mathbb{E}\left(\pi\left(\theta\right)|P\right)\right)\right)$, which confirms the above representation with $\omega = \sqrt{\gamma}$ and $u = \sqrt{\gamma}c \sim \mathcal{N}\left(\sqrt{\gamma}\bar{c}, \gamma\sigma_c^2\right)$.

the information-based skewness premium. Suppose that $\mathbb{E}(s)=0$ and consider first the limit without informed traders, as κ_I goes to 0. In this case, $d^U(P(z))$ must converge to s/κ_U, P must become completely uninformative, and price fluctuations are exclusively due to supply shocks. Therefore, $\Sigma(z)$ converges to the prior variance-covariance matrix Σ which is independent of $d^U(P(z)) = s/\kappa_U$, and hence the information-based premium vanishes: $\mathbb{E}\left(e_n'\Sigma(z)\cdot\left(d^U(P(z))-D\right)\right) = 0$. However $\mathbb{E}\left(d^U(P(z))'\,\Psi^n(z)\,d^U(P(z))\right)$ converges to $\frac{1}{\kappa_U^2}\mathbb{E}(s'\Psi^n s)$, where Ψ^n is the unconditional third-moment matrix. This last expression is positive (negative) whenever Ψ^n is positive (negative)-definite; with independent assets, this limit is $Skew(\pi_n(\theta))\,\mathbb{E}\left(s_n^2\right)$. This limit thus highlights that the preference-based skewness premium scales with the variance of supply shocks. *preference-based* skewness premium that scales with the variance of supply shocks:

Compare this limit with the alternative in which $\kappa_I \to 0$ and the distribution of supply shocks is also scaled by κ_I, i.e. $s=\kappa_I\tilde{s}$, where $\tilde{s}$ is distributed according to some fixed distribution $\tilde{G}$, with $\mathbb{E}(\tilde{s})=0$. This limit is equivalent to the limit where $\kappa_U\to\infty$, holding $\kappa_I=1$ and the distribution of supply shocks constant (i.e. we have entry of risk-averse, uninformed investors). In this limit, z remains informative and conveys information about $\int d_I(x,P(z))\,dF(x-\theta)-s$, i.e. $\Sigma(z)$ converges to a finite limit, and $d^U(P(z))$ must then scale with κ_I to satisfy market-clearing. The expected price premium satisfies

$$\mathbb{E}(P_n(z))-\mathbb{E}(\pi_n(\theta)) = -\chi\cdot cov\left(Var\left(\pi_n(\theta)\,|z\right),s_n\right)+o\left(Var(s_n)\right)$$

This last expression, and hence the information-based skewness premium, scales with κ_I, or the standard deviation of supply shocks, and conditional second moments of returns (variance), rather than third moments (skewness). By the same argument, letting κ_U tend to infinity implies that for large κ_U, the expected price premium vanishes at a rate $1/\kappa_U$, equal to inverse of the measure of uninformed traders. Hence sufficient entry by uninformed (but risk-averse) traders will result in convergence of prices to expected dividends - even if the information aggregated through the price remains noisy.

To conclude, for small values of κ_I and $\sigma^2_{s_n}$, the skewness premium becomes an order of magnitude larger with dispersed information than without. Relative to the benchmark with no dispersed information, noisy information aggregation thus amplifies the price premium for skewness.[50]

[50] In the online appendix, we show that the same result holds very generally when supply shocks are small: the expected price premium is of order σ_s^2 if there are no informed traders, but of order σ_s, when private information vanishes along-side supply shocks while keeping the informativeness of the price constant, or equivalently the mass of uninformed traders converges to infinity while keeping noise trading and informed trading constant. Hence an economy with small supply shocks and some privately informed traders may generate substantially larger average

6 Concluding Remarks

We have developed a theory of asset price formation based on dispersed information and its aggregation in asset markets. This theory ties expected asset returns to properties of their return distribution and the market's information structure. The theory imposes no restrictions on asset payoffs, investor information and asset supply and therefore speaks to much wider asset classes than most of the prior literature on noisy information aggregation. Finally, our theory is tractable and easily lends itself to applications as well as quantitative evaluation of asset pricing puzzles by calibrating model parameters to moments of forecast dispersion. In particular we show that our theory can account for a rich set of empirical facts regarding returns to skewness and forecast dispersion in equity and bond markets.

Future work will have to explore the quantitative implications of dispersed information for excess price volatility and return predictability, as well as other asset pricing puzzles. In Albagli et al. (2014) we use our framework to develop a dynamic model of corporate credit spreads. A second direction is to explore the effects of public news and information disclosures. A third direction consists in exploring how market frictions influence real decision-making by firms, households or policy makers. Using variants of our model, Bassetto and Galli (2019) compare information sensitivity of domestic and foreign debt and provide a theory of "original sin", and Gaballo and Galli (2022) develop a theory of quantitative easing based on information frictions and limits to arbitrage between bond and money holdings. In Albagli et al. (2017), we study the interplay between noisy information aggregation and risk-taking incentives. In an earlier version of this paper, we applied our model to security design and capital structure questions. These applications already suggest that our model may be useful to shed light on other economic phenomena well beyond empirical asset pricing puzzles.

References

[1] Ait-Sahalia, Yacine and Andrew Lo (2000). "Nonparametric risk management and implied risk aversion," Journal of Econometrics, 94 (1-2), 9-51.

price premia from skewness than the same economy in which all traders are uninformed. Another way to note the distinction between information- and preference-based skewness premia is to observe that with quadratic preferences, $\alpha = 0$, and therefore the preference-based skewness premium vanishes when agent's risk preferences do not display downside risk aversion. The information-based skewness premium instead vanishes if the price conveys no information about fundamentals and hence $Var\left(\pi\left(\theta\right)|z\right) = Var\left(\pi\left(\theta\right)\right)$.

[2] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2011). "A Theory of Asset Prices based on Heterogeneous Information," NBER Working paper No. 17548, revised 2015.

[3] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2014). "Dynamic Dispersed Information and the Credit Spread Puzzle," NBER Working paper No. 19788.

[4] Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2017). "Imperfect Financial Markets and Shareholder Incentives in Partial and General Equilibrium," NBER Working paper No. 23419.

[5] Andrei, Daniel, Julien Cujean and Mungo Wilson (2022). "The Lost Capital Asset Pricing Model," Review of Economic Studies, forthcoming.

[6] Audrino, Francesco, Robert Huitema, and Markus Ludwig (2021). "An Empirical Implementation of the Ross Recovery Theorem as a Prediction Device," Journal of Financial Econometrics, 19 (2), 291-312.

[7] Bai, Jennie, Robert S. Goldstein, and Fan Yang (2020). "Is the Credit Spread Puzzle a Myth?," Journal of Financial Economics, 137 (2), 297-319.

[8] Bakshi, Gurdip, Dilip Madan, and George Panayotov (2010). "Returns of Claims on the Upside and the Viability of U-shaped Pricing Kernels," Journal of Financial Economics 97 (1), 130-154.

[9] Barberis, Nicholas and Ming Huang (2008). "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," American Economic Review, 98 (5), 2066 - 2100.

[10] Barlevy, Gadi and Pietro Veronesi (2003). "Rational Panics and Stock Market Crashes," Journal of Economic Theory, 110 (2), 234-263.

[11] Bassetto, Marco and Carlo Galli (2019). "Is Inflation Default? The Role of Information in Debt Crises," American Economic Review, 109 (10), 3556 - 3584.

[12] Bernardo, Antonio E., and Kenneth L. Judd (2000). "Asset market equilibrium with general tastes, returns, and informational asymmetries." Journal of Financial Markets 3 (1), 17-43.

[13] Bhamra, Harjoat S., Lars-A. Kuehn, and Ilya A. Stebulaev (2010). "The Levered Equity Risk Premium and Credit Spreads: A Unified Framework," Review of Financial Studies, 23 (2), 645-703.

[14] Boyer, Brian, Todd Mitton, and Keith Vorkink (2010). "Expected Idiosyncratic Skewness," Review of Financial Studies, 23 (1), 169-202.

[15] Breon-Drish, Bradyn (2015). "On Existence and Uniqueness of Equilibrium in a Class of Noisy Rational Expectations Models." Review of Economic Studies, 82(3), 868-921.

[16] Brunnermeier, Markus. (2001). *Asset Pricing under Asymmetric Information - Bubbles, Crashes, Technical Analysis and Herding.* Oxford University Press.

[17] Brunnermeier, Markus and Jonathan Parker (2005). "Optimal Expectations," American Economic Review, 95 (4), 1092-1118.

[18] Brunnermeier, Markus, Christian Gollier, and Jonathan Parker (2007). "Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," American Economic Review, 97 (2), 159-165.

[19] Carlin, Bruce I., Francis A. Longstaff, and Kyle Matoba (2014). "Disagreement and Asset Prices," Journal of Financial Economics, 114 (2), 226-238.

[20] Carr, Peter, and Liuren Wu (2009). "Variance Risk Premiums," Review of Financial Studies 22 (3), 1311-1341.

[21] Chabakauri, Georgy, Kathy Yuan, and Konstantinos Zachariadis (2021). "Multi-Asset Noisy Rational Expectations Equilibrium with Contingent Claims," Review of Economic Studies, forthcoming.

[22] Chen, Hui, 2010, "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure," Journal of Finance, 65 (6), 2171–2212.

[23] Chen, Hui, Rui Cui, Zhiguo He and Konstantin Milbradt (2018). "Quantifying Liquidity and Default Risks of Corporate Bonds over the Business Cycle." Review of Financial Studies, 31 (3), 852-897.

[24] Christoffersen, Peter, Steven Heston and Kris Jacobs (2013). "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel," Review of Financial Studies 26 (8), 1963-2006.

[25] Conrad, Jennifer, Robert Dittmar, and Eric Ghysels (2013). "Ex Ante Skewness and Expected Stock Returns," Journal of Finance, 68 (1), 85-124.

[26] Diamond, Douglas and Robert Verrecchia (1981). "Information Aggregation in a Noisy Rational Expectations Economy," Journal of Financial Economics, 9 (3), 221-235.

[27] Diether, Karl, Christopher Malloy, and Anna Scherbina (2002). “Differences of Opinion and the Cross Section of Stock Returns,” Journal of Finance, 57 (5), 2113-2141.

[28] Feldhütter, Peter, and Stephen Schaefer (2018). “The Myth of the Credit Spread Puzzle.” Review of Financial Studies, 31 (8), 2897-2942.

[29] Gaballo, Gaetano, and Carlo Galli (2022). “Asset Purchases and Default-Inflation Risks in Noisy Financial Markets.” Working paper, HEC Paris and U. Carlos III.

[30] Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan (2001). “Toward an Implied Cost of Capital,” Journal of Accounting Research, 39(1), 135-76.

[31] Goldstein, Itay, Emre Ozdenoren, and Kathy Yuan (2013). “Trading frenzies and their impact on real investment,” Journal of Financial Economics, 109(2), 566-582.

[32] Green, Clifton and Byoung-Hyoun Hwang (2012). “Initial Public Offerings as Lotteries: Skewness Preference and First-Day Returns,” Management Science, 58 (2), 432-444.

[33] Grossman, Sanford J. and Joseph E. Stiglitz (1980). “On the Impossibility of Informationally Efficient Markets,” American Economic Review, 70 (3), 393-408.

[34] Güntay, Levent, and Dirk Hackbarth (2010). “Corporate Bond Credit Spreads and Forecast Dispersion.” Journal of Banking and Finance, 34 (10), 2328-2345.

[35] He, Zhiguo and Konstantin Milbradt (2014). “Endogenous Liquidity and Defaultable Bonds.” Econometrica, 82 (4), 1443-1508.

[36] Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski (2006). “Self-Fulfilling Currency Crises: The Role of Interest Rates,” American Economic Review, 96 (5), 1769-1787.

[37] Hellwig, Martin (1980). “On the aggregation of Information in Competitive Markets,” Journal of Economic Theory, 22 (3), 477-498.

[38] Hou, Kewei, Chen Xue and Lu Zhang (2020). “Replicating Anomalies,” Review of Financial Studies, 33 (5), 2019-2133.

[39] Huang, Jing-zhi, and Ming Huang (2012). “How much of the Corporate-Treasury Yield Spread is due to Credit Risk?” Review of Asset Pricing Studies, 2 (2), 153-202.

[40] Jackwerth, Jens (2000). “Recovering Risk Aversion from Option Prices and Realized Returns” Review of Financial Studies, 13 (2), 433–451.

[41] Kovbasyuk, Sergey, and Marco Pagano (2022). "Advertizing Arbitrage," Review of Finance, 26 (4), 799–827.

[42] Kremer, Ilan (2002). "Information aggregation in common value auctions." Econometrica, 70 (4), 1675-1682.

[43] Malamud, Semyon (2015). "Noisy Arrow-Debreu Equilibria". Working paper, EPFL.

[44] McQuade, Timothy (2018). "Stochastic Volatility and Asset Pricing Puzzles". Working paper, Stanford GSB.

[45] Milgrom, Paul (1979). "A convergence theorem for competitive bidding with differential information." Econometrica, 47 (3), 679-688.

[46] Milgrom, Paul (1981). "Rational Expectations, Information Acquisition, and Competitive Bidding," Econometrica, 49 (4), 921-943.

[47] Miller, Edward (1977). "Risk, uncertainty, and divergence of opinion." Journal of Finance 32 (4), 1151-1168.

[48] Mitton, Todd and Keith Vorkink (2007). "Equilibrium Underdiversification and the Preference for Skewness", Review of Financial Studies, 20 (4), 1255-1288.

[49] Peress, Joel (2004). "Wealth, information acquisition, and portfolio choice." Review of Financial Studies, 17 (3), 879-914.

[50] Perry, Motty, and Philip J. Reny (2006). "Toward a strategic foundation for rational expectations equilibrium." Econometrica, 74 (5), 1231-1269.

[51] Pesendorfer, Wolfgang, and Jeroen M. Swinkels (1997). "The loser's curse and information aggregation in common value auctions." Econometrica, 65 (6), 1247-1281.

[52] Saumard, Adrien, and Jon Wellner (2014). "Log-Concavity and Strong Log-Concavity: a Review," working paper (arXiv:1404.5886).

[53] Straub, Ludwig, and Robert Ulbricht (2015). "Endogenous Uncertainty and Credit Crunches," working paper, MIT and TSE.

[54] Veldkamp, Laura (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.

[55] Vives, Xavier (2008). *Information and Learning in Markets.* Princeton University Press.

[56] Wilson, Robert (1977). "A bidding model of perfect competition," Review of Economic Studies, 44 (3), 511-518.

[57] Yuan, Kathy (2005). "Asymmetric Price Movements and Borrowing Constraints: a Rational Expectations Equilibrium Model of Crises, Contagion and Confusion," Journal of Finance, 60(1), 379- 411.

[58] Yu, Jialin (2011). "Disagreement and Return Predictability of Stock Portfolios," Journal of Financial Economics, 99 (1), 162-183.

[59] Zhang, Xiao-Jun (2013). "Book-to-Market Ratio and Skewness of Stock Returns." The Accounting Review, 88 (6), 2213-2240.

7 Appendix: Proofs

Proof of Proposition 1:

See Online Appendix.

Proof of Proposition 2:

The price function $P_{\pi}(z) = \mathbb{E}(\pi(\theta)|x = z, z)$ given by (9) is continuous and strictly increasing in z. It then follows from arguments given in the text that when coupled with the threshold $\hat{x}(P) = z$ and the associated posterior beliefs, $P_{\pi}(z)$ constitutes an equilibrium in which $d(x, P)$ is non-increasing in P. Moreover, by market-clearing, $z = \hat{x}(P_{\pi}(z))$ and $z' = \hat{x}(P_{\pi}(z'))$, and therefore $z = z'$ if and only if $P_{\pi}(z) = P_{\pi}(z')$. Therefore, the equilibrium conjectured above is the only equilibrium, in which P is informationally equivalent to z.

It remains to be shown that there exists no other equilibrium in which demand is non-increasing in P. In any equilibrium, in which $d(x, P)$ is non-increasing in P, $\hat{x}(P)$ must be non-decreasing in P. Moreover, $\hat{x}(P)$ must be continuous – otherwise, if there were jumps, then there would be certain realizations for z, for which there is no P, such that $\hat{x}(P) = z$, implying that the market cannot clear at these realizations of z. Now, if $\hat{x}(P)$ is strictly increasing in P, it is invertible, and we are therefore back to the equilibrium that we have already characterized. Suppose therefore that $\hat{x}(P)$ $\hat{x}(P) = \hat{x}(P') = \hat{x}(P'')$ for $P \in (P', P'')$ and $P'' > P'$. Suppose further that for sufficiently low

$\varepsilon > 0$, $\hat{x}(P)$ is strictly increasing over $(P' - \varepsilon, P')$ and $(P'', P'' + \varepsilon)$, and hence uniquely invertible.[51] But then for $z \in (\hat{x}(P' - \varepsilon), \hat{x}(P')$ and $z \in (\hat{x}(P''), \hat{x}(P'' + \varepsilon))$, $P(z)$ is uniquely defined, so we have $P' \geq lim_{z\uparrow\hat{x}} P(z) = lim_{z\uparrow\hat{x}} \mathbb{E}(\pi(\theta)|x = z, z)$ and $P'' \leq lim_{z\downarrow\hat{x}} P(z) = lim_{z\downarrow\hat{x}} \mathbb{E}(\pi(\theta)|x = z, z)$. But since $\mathbb{E}(\pi(\theta)|x = z, z)$ is continuous, it must be that

$$P'' \leq lim_{z\downarrow\hat{x}} \mathbb{E}(\pi(\theta)|x = z, z) = lim_{z\uparrow\hat{x}} \mathbb{E}(\pi(\theta)|x = z, z) \leq P',$$

which yields a contradiction.

Proof of Proposition 3:

Part (i) follows directly from applying Definition 1 in equation (12) and from taking the derivative w.r.t. $\hat{\sigma}_\theta$. Part (ii) follows from additivity (for given $\hat{\sigma}_\theta$, $W(\pi_1, \hat{\sigma}_\theta) - W(\pi_2, \hat{\sigma}_\theta) = W(\pi_1 - \pi_2, \hat{\sigma}_\theta)$) and applying part (i) to $\pi_1 - \pi_2$. For the limit as $\hat{\sigma}_\theta \to \infty$, note that $lim_{\hat{\sigma}_\theta \to \infty} \int_{-\infty}^{\infty} \left(\pi\left(\frac{\hat{\sigma}_\theta}{\sigma_\theta}\theta\right)\right) d\Phi\left(\frac{\theta}{\sigma_\theta}\right) = lim_{\theta\to\infty} \frac{1}{2}(\pi(\theta) + \pi(-\theta))$, and therefore $lim_{\hat{\sigma}_\theta\to\infty} |W(\pi, \hat{\sigma}_\theta)| = lim_{\theta\to\infty} \frac{1}{2} |\pi(\theta) + \pi(-\theta)|$.

Derivation of equation 14:

Simple algebra shows that

$$\frac{\hat{\sigma}_\theta^2}{\sigma_\theta^2} = 1 + \hat{\gamma}\frac{\hat{\gamma} - \gamma}{\gamma} = 1 + \hat{\gamma}\frac{\beta/\sigma_\theta^2}{\left(1/\sigma_\theta^2 + \beta + \tau\right)\tau} = 1 + \frac{\beta/\sigma_\theta^2}{\left(1/\sigma_\theta^2 + \beta + \tau\right)^2}\frac{\beta + \tau}{\tau}$$

Since $\tilde{D}^2 = \frac{\beta/\sigma_\theta^2}{\left(1/\sigma_\theta^2 + \beta + \tau\right)^2}$ and $\hat{\gamma}(1 - \hat{\gamma}) = \frac{(\beta+\tau)/\sigma_\theta^2}{\left(1/\sigma_\theta^2 + \beta + \tau\right)^2}$, it follows that

$$\frac{\hat{\sigma}_\theta^2}{\sigma_\theta^2} = 1 + \tilde{D}^2 \frac{\hat{\gamma}(1 - \hat{\gamma})}{\hat{\gamma}(1 - \hat{\gamma}) - \tilde{D}^2}.$$

Proof of Theorem 1:

We begin with two useful lemmas:

Lemma 1 *Suppose that θ is distributed according to cdf. $H(\cdot)$ and that $f(\cdot)$ is log-concave and $f'(\cdot)/f(\cdot)$ unbounded. Then $H(\theta|x) \equiv \int_{-\infty}^{\theta} f(x - \theta')\, dH(\theta') / \int_{-\infty}^{\infty} f(x - \theta')\, dH(\theta')$ is decreasing in x, with $\lim_{x\to-\infty} H(\theta|x) = 1$ and $\lim_{x\to\infty} H(\theta|x) = 0$.*

[51] It cannot be flat everywhere, because then informed demand would be completely inelastic, and there would be no way to absorb supply shocks.

Proof. Notice that

$$\frac{H(\theta|x)}{1-H(\theta|x)} = \frac{\int_{-\infty}^{\theta} f(x-\theta')\,dH(\theta')}{\int_{\theta}^{\infty} f(x-\theta')\,dH(\theta')} = \frac{\int_{-\infty}^{\theta} \frac{f(x-\theta')}{f(x-\theta)}dH(\theta')}{\int_{\theta}^{\infty} \frac{f(x-\theta')}{f(x-\theta)}dH(\theta')} = \frac{H(\theta)}{1-H(\theta)} \frac{\mathbb{E}\left(\frac{f(x-\theta')}{f(x-\theta)}|x,\theta'\le\theta\right)}{\mathbb{E}\left(\frac{f(x-\theta')}{f(x-\theta)}|x,\theta'>\theta\right)}$$

Log-concavity and MLRP of f imply that whenever $\theta' < \theta$, $f(x-\theta')/f(x-\theta)$ is decreasing in x with $lim_{x\to-\infty} f(x-\theta')/f(x-\theta) = \infty$ and $lim_{x\to\infty} f(x-\theta')/f(x-\theta) = 0$. It follows that the second ratio is strictly decreasing in x and converges to 0 as $x\to\infty$ and ∞ as $x\to-\infty$. ■

Lemma 2 *In any equilibrium, and for any P on the interior of the support of $\pi(\theta)$, there exist $x_L(P)$ and $x_H(P)$, such that $d(x,P) = d_L$ for all $x \le x_L(P)$, $d(x,P) = d_H$ for all $x \ge x_H(P)$, and $d(x,P)$ is strictly increasing in x for $x \in (x_L(P), x_H(P))$*

Proof. For any D, consider the risk-adjusted cdf

$$H(\cdot|P;D) = \frac{\int_{-\infty}^{\theta} U'(D(\pi(\theta)-P))\,dH(\theta|P)}{\int_{-\infty}^{\infty} U'(D(\pi(\theta)-P))\,dH(\theta|P)},$$

and let $H(\cdot|x,P;D)$ and $\mathbb{E}(\pi(\theta)|x,P;D) \equiv \int \pi(\theta)\,dH(\theta|x,P;D)$ denote the cdf and conditional expectations after updating conditional on a private signal x. By Lemma 1, $H(\cdot|x,P;D)$ is strictly decreasing in x, $\mathbb{E}(\pi(\theta)|x,P;D)$ is strictly increasing in x and $\lim_{x\to-\infty}\mathbb{E}(\pi(\theta)|x,P;D) < P < \lim_{x\to\infty}\mathbb{E}(\pi(\theta)|x,P;D)$ for any P on the interior of the support of $\pi(\cdot)$. But then there exist $x_L(P)$ s.t. $\mathbb{E}(\pi(\theta)|x_L(P),P;d_L) = P$, which implies that $d(x,P) = d_L$ for all $x \le x_L(P)$, and $x_H(P)$s.t. $\mathbb{E}(\pi(\theta)|x_H(P),P;d_H(P)) = P$, which implies that $d(x,P) = d_H$ for all $x \ge x_H(P)$.

For $x \in (x_L(P), x_H(P))$ and $x' > x$, Lemma 1 implies that $P = \mathbb{E}(\pi(\theta)|x,P;d(x;P)) < \mathbb{E}(\pi(\theta)|x',P;d(x;P))$, or equivalently $\mathbb{E}((\pi(\theta)-P)|x',P;d(x,P)) > 0$. Since the LHS of this condition is strictly decreasing in d, it follows that $d(x',P) > d(x,P)$. ■

Lemmas 1 and 2, and $d(x,P)$ strictly decreasing in P imply that there exists a unique $z(P) \in (x_L(P), x_H(P))$; s.t. $d(z(P),P) = \bar{D}$, or equivalently, $P = P_\pi(z) = \mathbb{E}\left(\pi(\theta)|x = z(P), P; \bar{D}\right)$. Combining with the equilibrium price function, we then define a candidate sufficient statistic function $z(\theta,u) = z(P(\theta,u))$, and since $z(P)$ is invertible, z must be a sufficient statistic for the information contained in P. Therefore we obtain the representation (16), along with representation (15) of equilibrium beliefs.

Proof of Proposition 4:

Set $\bar{D} = 0$ (wolg). Fix θ_n and $z_n = \theta_n + k$. We first show that $\lim_{\theta_n \to \infty} m(\theta_n, z_n) = \lim_{\theta_n \to -\infty} m(\theta_n, z_n) = \infty$. To see this, write

$$\begin{aligned}
m(\theta, z)^{-1} &= \frac{\mathbb{E}(f(z-\theta')|z)}{f(z-\theta)} = \frac{f(z-\theta) - \int_\theta^\infty (1 - H(\theta'|z)) f'(z-\theta') d\theta' + \int_{-\infty}^\theta H(\theta'|z) f'(z-\theta') d\theta'}{f(z-\theta)} \\
&= 1 - \frac{\int_\theta^\infty (1 - H(\theta'|z)) f'(z-\theta') d\theta'}{\int_\theta^\infty f'(z-\theta') d\theta'} - \frac{\int_{-\infty}^\theta H(\theta'|z) f'(z-\theta') d\theta'}{\int_{-\infty}^\theta f'(z-\theta') d\theta'} \\
&= 1 - \frac{\int_{-\infty}^{z-\theta} (1 - H(z-u'|z)) f'(u') du'}{\int_{-\infty}^{z-\theta} f'(u') du'} - \frac{\int_{z-\theta}^\infty H(z-u'|z) f'(u') du'}{\int_{z-\theta}^\infty f'(u') du'}
\end{aligned}$$

Now, with $z_n = \theta_n + k$, we have

$$m(\theta_n, z_n) = 1 - \frac{\int_{-\infty}^{k} (1 - H(z_n-u'|z_n)) f'(u') du'}{\int_{-\infty}^{k} f'(u') du'} - \frac{\int_{k}^\infty H(z_n-u'|z_n) f'(u') du'}{\int_{k}^\infty f'(u') du'}$$

and since $\lim_{z\to\infty} H(z-u'|z) = 1$, it follows that $\lim_{\theta_n\to\infty} m(\theta_n, z_n) = \infty$. Likewise, since $\lim_{z\to-\infty} H(z-u'|z) = 0$, it also follows that $\lim_{\theta_n\to-\infty} m(\theta_n, z_n) = \infty$. In addition, it is straightforward to check that $\lim_{z\to-\infty} m(\theta, z) = \lim_{z\to\infty} m(\theta, z) = \infty$, for fixed θ.

Next notice that for fixed k, $Pr(x > z_n|\theta_n) = 1 - F(z_n - \theta_n) = 1 - F(k)$. It follows that $F(k) d_L \le D(\theta_n, P(z_n)) \le (1 - F(k)) d_H$, and therefore $D(\theta_n, P(z_n))$ is strictly bounded away from d_H and d_L. But then, $\Psi(z_n|\theta_n) = 1 - G(D(\theta_n, P(z_n)))$ is strictly and uniformly bounded away from 0 and 1 for all θ_n.

Finally, consider $m(\theta) = \int m(\theta, z) \psi(z|\theta) dz$. Since $\lim_{z\to\infty} m(\theta, z) = \infty$ it must be the case that

$$m(\theta_n) \ge m(\theta_n, z_n)(1 - \Psi(z_n|\theta_n)) + \int_{-\infty}^{z_n} m(\theta_n, z) \psi(z|\theta_n) dz.$$

for θ_n sufficiently high. Since $\lim_{\theta_n\to\infty} \Psi(z_n|\theta_n) < 1$ and $\lim_{\theta_n\to\infty} m(\theta_n, z_n) = \infty$, it follows that $\lim_{\theta_n\to\infty} m(\theta_n) = \infty$. Likewise, $m(\theta_n) \ge m(\theta_n, z_n) \Psi(z_n|\theta_n)$ for θ_n sufficiently low, and it follows that $\lim_{\theta_n\to-\infty} m(\theta_n) = \infty$.

Proof of Proposition 5:

See Online Appendix.

Proof of Theorem 2:

With symmetry, we have $\hat{H}_k(\theta) = \int_{-\infty}^\theta m_k(\theta') h(\theta') d\theta' = \int_{-\infty}^\theta m_k(-\theta') h(-\theta') d\theta' = \int_\theta^\infty m_k(\theta') h(\theta') d\theta' =$

$1-\hat{H}_k(\theta)$. Moreover,

$$W(\pi,\hat{m}_k)=\int_{-\infty}^{\infty}\Big(H(\theta)-\hat{H}_k(\theta)\Big)\,d\pi(\theta)=\int_0^{\infty}\Big(H(\theta)-\hat{H}_k(\theta)\Big)\,d\left(\pi(\theta)+\pi(-\theta)\right).$$

Finally, by log-convexity of $m_2(\cdot)$ and $m_1(\cdot)/m_1(\cdot)$, $\hat{H}_2(\theta)\geq\hat{H}_1(\theta)\geq H(\theta)$ for $\theta\leq 0$ and $\hat{H}_2(\theta)\leq\hat{H}_1(\theta)\leq H(\theta)$ for $\theta\geq 0$.

Part (i) then follows from Definition 1 and the ordering of distributions. Part (ii) follows from additivity, $W(\pi_1,m)-W(\pi_2,m)=W(\pi_1-\pi_2,m)$, and applying part (i) to $\pi_1-\pi_2$.

To complete the proof, we show that $|W(\pi,m)|$ may become arbitrarily large if $\hat{H}(\cdot)$ converges to an improper distribution characterized by $\hat{H}(\theta)\to\frac{1}{2}$. In this limit case, we rewrite

$$|W(\pi,m)|=\int_0^{\infty}\Big(H(\theta)-\hat{H}(\theta)\Big)\,d\hat{\pi}(\theta),$$

where $\hat{\pi}(\theta)=|\pi(\theta)+\pi(-\theta)|$. For some (small) $\varepsilon>0$, there exist $\bar{\theta}$ and $\underline{\theta}$, such that $H(\theta)-\hat{H}(\theta)\geq\varepsilon$ for $\theta\in\left(\underline{\theta},\bar{\theta}\right)$, and therefore

$$|W(\pi,\hat{m})|\geq\int_{\underline{\theta}}^{\bar{\theta}}\Big(H(\theta)-\hat{H}(\theta)\Big)\,d\hat{\pi}(\theta)\geq\varepsilon\left(\hat{\pi}\left(\bar{\theta}\right)-\hat{\pi}\left(\underline{\theta}\right)\right)$$

As $\hat{H}(\theta)\to\frac{1}{2}$, $\bar{\theta}\to\infty$ and $\underline{\theta}\to H^{-1}\left(\frac{1}{2}+\varepsilon\right)$. It follows that $\lim_{\hat{H}(\theta)\to\frac{1}{2}}|W(\pi,\hat{m})|=\lim_{\bar{\theta}\to\infty}\varepsilon\hat{\pi}\left(\bar{\theta}\right)=\infty$.