



HAL
open science

Information aggregation with asymmetric asset payoffs

Elias Albagli, Christian Hellwig, Aleh Tsyvinski

► **To cite this version:**

Elias Albagli, Christian Hellwig, Aleh Tsyvinski. Information aggregation with asymmetric asset payoffs. Journal of Finance, 2024, 79 (4), pp.2715-2758. <10.1111/jofi.13361>. <hal-04867329>

HAL Id: hal-04867329

<https://hal.science/hal-04867329v1>

Submitted on 6 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

WORKING PAPERS

N° 1172

April 2023

“Information Aggregation with Asymmetric Asset Payoffs”

Elias Albagli, Christian Hellwig and Aleh Tsyvinski

Information Aggregation with Asymmetric Asset Payoffs*

Elias Albagli
Central Bank of Chile

Christian Hellwig
Toulouse School of Economics

Aleh Tsyvinski
Yale University

April 17, 2023

Abstract

We study noisy aggregation of dispersed information in financial markets without imposing parametric restrictions on preferences, information, and return distributions. We provide a general characterization of asset returns by means of a risk-neutral probability measure that features excess weight on tail risks. Moreover, we link excess weight on tail risks to observable moments such as forecast dispersion and accuracy, and argue that it provides a unified explanation for several prominent cross-sectional return anomalies. Simple calibrations suggest the model can account for a significant fraction of empirical returns to skewness, returns to disagreement and interaction effects between the two.

*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 asset payoffs and market features 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 an environment, the market-clearing price serves as an endogenous public signal of the asset payoff. In contrast to the existing literature, we don't impose parametric restrictions on the payoff distribution, 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 noisy information aggregation, and show that it generates *excess weight on tail risks*: the pricing kernel or change in probability measure is U-shaped, overweighting probabilities of both very high and low payoff realizations. The resulting price premia can then 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. 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 from the difficulty of tractably dealing with the endogeneity of information contained in the price.

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

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*: the asset price is 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, 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 to highlight the interaction between payoff asymmetry and noisy information aggregation. 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. We show that the difference between the average price and asset payoffs, or the expected 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 relates these comparative statics to several empirical regularities. Excess weight on tail risks is consistent with empirical evidence of non-monotonic or U-shaped pricing kernels that seem at odds with standard risk-based no arbitrage conditions (the *pricing kernel puzzle*) or negative

variance premia (the *variance premium puzzle*), which imply that the risk-neutral variance is a systematically upwards-biased predictor of the underlying return variance.² The other comparative statics translate into empirical predictions for the cross-section of asset returns:

1. *Negative 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 accounts for all three empirical regularities. In particular, it explains why higher 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.

A simple calibration suggests model-implied returns to skewness and disagreement can be quite large. We parametrize asset payoff distributions in the risk-neutral model to match idiosyncratic skewness and volatility metrics for equity returns,⁷ and calibrate 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. While not a perfect measure of investor private information (after all, analyst forecasts are publicly observed!), forecast dispersion is arguably a reasonable proxy of dispersed information between different investors. The analyst forecast sample suggests

²See 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.

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

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

⁵See 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.

⁶Yu (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.

⁷See Boyer et al. (2010).

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 generates 37% of the observed return differential between the highest and lowest skewness quintiles, 70% of the observed return differential between the highest and lowest disagreement quintiles, and an even higher fraction of the interaction effects, with annualized excess returns to skewness (disagreement) varying by as much as 6 percentage points between highest and lowest disagreement (skewness) quintiles.

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, or U-shaped. 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 uninformed traders and multiple securities. Introducing uninformed traders allows us to link excess weight on tail risks to counter-cyclical risk premia, and highlight how returns to skewness are amplified by dispersed information. The second extension generalizes the updating wedge and risk-neutral representation to multi-asset environments.

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. 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) study non-linear models of noisy information aggregation with a single asset market, Malamud (2015) and Chabakauri et al. (2022) non-linear, multi-asset noisy REE models with state-contingent securities and complete or incomplete markets. 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. We instead identify properties of asset prices that apply beyond the specifics of their environment, and link such properties to cross-sectional return anomalies.

⁸He further studies incentives for information acquisition, whereas we take the information structure as given.

Our equilibrium characterization 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.

Section 2 introduces our general model and conveys our main results with three examples. Section 3 discusses how the model’s main comparative statics relate to empirical regularities. In Section 4, we provide generalizations of our main results to arbitrary preferences and distributional assumptions. Section 5 discusses extensions with uninformed traders and multiple securities.

2 The general model and three examples

We begin by describing the information structure and the financial market for our general model. 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 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 $d_L \leq 0 \leq d_H$ and $d_L < d_H$.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

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

¹⁰Monotonicity 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.

¹¹Position limits may be infinite ($(d_L, d_H) = \mathbb{R}$) if investors are strictly risk-averse (so that asset demand is bounded), but must be finite when investors are risk-neutral. Both scenarios may be relevant, depending on context.

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

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 $(\theta, s) \in \mathbb{R} \times [d_L, d_H]$,

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

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

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

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

Investors in our model do not observe signals about asset payoffs, but rather about a fundamental θ , and the asset payoff is a monotone function of θ .¹⁴ This formulation separates 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, leverage and other variables that are often summarized in a single "distance to default" metric. A derivative trader will forecast the underlying. 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 the key ideas described in the introduction that we will generalize throughout the rest of our paper: the *sufficient statistic representation* of the equilibrium price, the *information updating wedge* which makes the price more sensitive to this sufficient statistic than is warranted by the information it conveys about

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

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

the fundamental, and the construction of a risk-neutral probability measure that displays *excess weight on tail risks* (EWTR, henceforth). 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 $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$; (ii) normal distribution of supply: $s \sim \mathcal{N}(\bar{s}, \sigma_s^2)$; (iii) normally distributed private signals, $x_i|\theta \sim \mathcal{N}(\theta, \beta^{-1})$; (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 $d(x, P) = (\chi \text{Var}(\theta|x, P))^{-1} (\mathbb{E}(\theta|x, P) - P)$. Now, define $z(P)$ 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. 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 $d(z(P), P) = \bar{s}$, we can invert the demand relation to obtain the following *sufficient statistic representation* of the equilibrium price as a function of z :

$$P(z) = \mathbb{E}(\theta|x = z, z) - \chi \text{Var}(\theta|x = z, z) \cdot \bar{s}. \quad (3)$$

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 $\text{Var}(\theta|x, z) = (1/\sigma_\theta^2 + \beta + \tau)^{-1}$. Therefore, $\mathbb{E}(\theta|x = z, z) = \hat{\gamma} \cdot z$ and $\text{Var}(\theta|x = z, z) = (1 - \hat{\gamma}) \sigma_\theta^2$, 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 \bar{s} units of the security. The significance of this representation becomes clear if 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 $\text{Var}(\theta|z) = (1 - \gamma) \sigma_\theta^2$, 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 $\chi(1 - \hat{\gamma}) \sigma_\theta^2 \cdot \bar{s}$, which scales with risk aversion, posterior uncertainty, and average supply.

The 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

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

z' conveys positive news about θ through the information contained in the price, which 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 with weight τ , and the additional market-clearing adjustment with weight β . In contrast, the Bayesian posterior of θ given z only includes the first effect.

The market-clearing effect stems from asset demand with dispersed information being finitely elastic, due to risk aversion. Demand (fundamental) and supply shocks then have price impact, which is reflected in 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 $z \sim \mathcal{N}(\theta, \tau^{-1})$. With CARA preferences and taking as given supply s , the asset price with *common information* is $V(z, s) = \mathbb{E}(\theta|z) - \chi \text{Var}(\theta|z) s$.

Risk-neutral measure and excess weight on tail risks: Let $h(\theta|z)$ the pdf of θ conditional on z and $h(\theta|x, z)$ the pdf of θ conditional on x and z . The asset price can equivalently be represented as $P(z) = \mathbb{E}((\theta - R) \cdot m^I(\theta, z) | z)$, where $m^I(\theta, z) \equiv \frac{h(\theta|x=z, z)}{h(\theta|z)}$ summarizes the impact of noisy information aggregation and $R \equiv \chi(1 - \hat{\gamma})\sigma_\theta^2 \cdot \bar{s}$ the expected risk premium. Taking prior expectations, we obtain

$$\mathbb{E}(P(z)) = \mathbb{E}((\theta - R) \cdot m(\theta)) = \int_{-\infty}^{\infty} \theta' d\hat{H}(\theta') \quad (4)$$

where $m(\theta) = \mathbb{E}(m^I(\theta, z) | \theta)$, and $\hat{H}(\theta) = \int_{-\infty}^{\theta} m(\theta' + R) dH(\theta' + R)$. It is straight-forward to check that $\mathbb{E}(m^I(\theta, z) | z) = \mathbb{E}(m(\theta)) = 1$, 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 first-order shift in fundamentals due to the risk adjustment $R = \chi(1 - \hat{\gamma})\sigma_\theta^2 \cdot \bar{s}$ with the informational adjustment $m(\cdot)$.

Equivalently, we can construct $\hat{H}(\cdot)$ by compounding the posterior $\theta|x = z, z \sim \mathcal{N}(\hat{\gamma}z, (1 - \hat{\gamma})\sigma_\theta^2)$, with the prior distribution over z , $z \sim \mathcal{N}(0, \sigma_\theta^2 + \tau^{-1})$, 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 $\hat{\sigma}_\theta^2 \equiv (1 - \hat{\gamma} + \hat{\gamma}^2/\gamma)\sigma_\theta^2 > \sigma_\theta^2$. 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. Mathematically, this result relies on the fact that $\log(m(\theta))$ is convex-quadratic, and thus

U-shaped in the realization of θ , for normal distributions. EWTR distinguishes the risk-neutral measure with noisy information aggregation from an average risk premium, which shifts probability mass from higher to lower fundamentals, 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.

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

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


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

However, $\hat{\sigma}_\theta^2 \approx \sigma_\theta^2 + \sigma_\theta^4 \cdot (\beta + \chi^2 \sigma_s^2)$ is increasing in β for small positive β . Dispersed information thus amplifies EWTR when private signals are sufficiently noisy. 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.¹⁶ Furthermore, when $\sigma_s^2 \rightarrow 0$ but the precision of the price signal $\sqrt{\tau} = \chi \sigma_s / \beta$ is held constant, $\hat{\sigma}_\theta^2 - \sigma_\theta^2$ is of order $\chi \sigma_s$ when $\tau > 0$ but of order $\chi^2 \sigma_s^2$ when $\beta = \tau = 0$. Hence amplification can become arbitrarily large when supply shocks are small.

Asset-pricing predictions: Figure 1 compares the price (thick solid line) with expected dividends (thin solid line), conditional on z (assuming average supply $\bar{s} = 0$). Since $\hat{\gamma} > \gamma$, the price responds more to z than the underlying dividend expectations. We can then derive return

¹⁶At the other extreme, EWTR vanishes ($\hat{\sigma}_\theta^2 - \sigma_\theta^2 \rightarrow 0$) in the limit as $\beta \rightarrow \infty$, since the price then becomes perfectly revealing and must therefore converge to the true fundamental.

Figure 1: Linear payoff model –CARA-normal example



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 difference between the expected dividend (solid horizontal line) and the expected price with $\bar{s} > 0$ (dashed horizontal line) only depends on the expected risk premium R and not on the EWTR property. When $\bar{s} = 0$, the unconditional expectations of price and dividend coincide, as payoff symmetry implies that overpricing when z is positive is exactly offset by underpricing when z is negative.

Below, we argue that equilibrium asset prices with noisy information aggregation display these same properties for a wide range of economic primitives: (i) the realized asset price can be represented as a function of a sufficient statistic z with an information updating wedge that mirrors the one described above, and (ii) this updating wedge then gives rise to a risk-neutral measure that decomposes into a risk premium and the EWTR property, where the latter results from the interaction between supply shocks and investors' private information. At a general level, we show that EWTR maps into log-convexity (or U-shape) of the pricing kernel function. However, once payoffs are asymmetric, over-pricing on the upside and under-pricing on the downside no longer offset each other, and EWTR has first-order implications for average prices and expected 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 are as in section 2.1.¹⁷

With CARA preferences and binary payoffs, an investor with private signal x demands

$$d(x, P) = \frac{1}{\chi} \left(\log \left(\frac{\mu(x, P)}{1 - \mu(x, P)} \right) - \log \left(\frac{P}{1 - P} \right) \right) \quad (5)$$

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 $z(P)$ 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 $z|\theta \sim \mathcal{N}(\theta, \tau^{-1})$ is normally distributed with mean $\theta \in \{0, 1\}$ and precision τ , in which case $\mu(x, P(z))$ takes the form

$$\log \left(\frac{\mu(x, P(z))}{1 - \mu(x, P(z))} \right) = \log \left(\frac{\lambda}{1 - \lambda} \right) + \beta \left(x - \frac{1}{2} \right) + \tau \left(z - \frac{1}{2} \right). \quad (6)$$

Substituting (6) into (5) and inverting the condition $d(z, P) = \bar{s}$ then leads to the following sufficient statistic representation of the equilibrium price:

$$P(z) = \frac{\lambda e^{(\beta+\tau)(z-\frac{1}{2})-\chi\bar{s}}}{\lambda e^{(\beta+\tau)(z-\frac{1}{2})-\chi\bar{s}} + 1 - \lambda}. \quad (7)$$

From the market-clearing condition $s = \int d(x, P) d\Phi(\sqrt{\beta}(x - \theta))$, for $\theta \in \{0, 1\}$, it is then straight-forward to verify that $z = \theta - \chi/\beta \cdot (s - \bar{s})$ and hence $z|\theta \sim \mathcal{N}(\theta, \tau^{-1})$ with $\tau^{-1} = (\chi/\beta)^2 \cdot \sigma_s^2$, which confirms our initial conjecture. As in the CARA-normal model, equation (7) 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 (7). 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”:

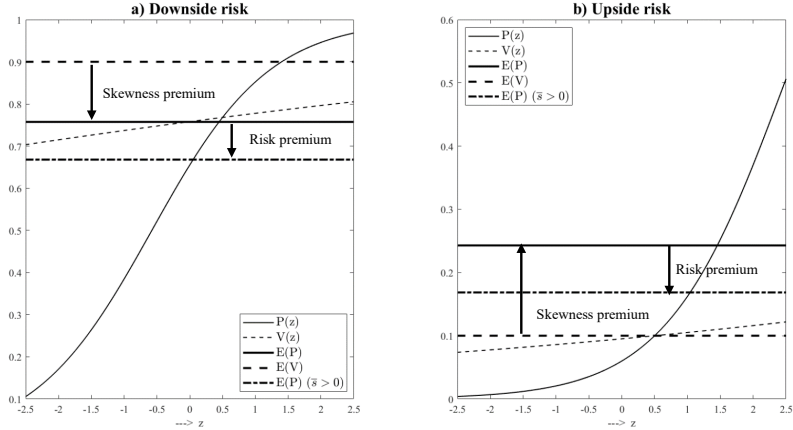
Proposition 1 : *If $\bar{s} = 0$, there exists $\Delta \in (0, 1)$ such that the expected price takes the form*

$$\mathbb{E}(P(z)) = \lambda + \left(\frac{1}{2} - \lambda \right) \Delta. \quad (8)$$

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

¹⁷This example is a special case of example 1 in Breon-Drish (2015).

Figure 2: Non-linear payoffs – CARA-binary model



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.

Figure 2 plots the price and dividend expectation as a function of z , for the cases of downside (Panel a, $\lambda = 0.9$), and upside (Panel b, $\lambda = 0.1$) risks, with $\bar{s} = 0$. 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 ($\lambda = 1/2$, not plotted), and as was the case for the CARA-normal model, the average price coincides with the average payoff. For downside risks, the difference between price and expected dividend is much more pronounced for low realizations of z than for high draws. Positive updates about payoffs tend to have limited effects on investor expectations as the potential for upside gains, relative to prior expectations, was already limited. Negative updates may instead have much larger effects on investor beliefs if they indicate that a high payoff becomes less likely than previously thought. The updating wedge is thus amplified on the downside. Taking expectations across realizations of z , this asymmetry leads to expected prices (thick solid horizontal line) below average payoffs (dashed horizontal line), even when the average supply is set to zero ($\bar{s} = 0$). The opposite asymmetry arises with upside risks.

In addition, Δ 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, just like EWTR in the CARA-

normal model, the skewness premium results from the presence of supply shocks that are orthogonal to fundamentals, but is amplified by private information.

The same forces are at play here, but the additional 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. For upside risks ($\lambda < 1/2$), 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.

With noisy private signals, a positive supply shock then increases risk premia 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, while increasing posterior uncertainty for downside risks by raising their (ex ante low) prior that the asset may actually fail to pay off. In other words, for upside risks, 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 $(1/2 - \lambda) \Delta$ 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.¹⁸

In the online appendix, we extend Proposition 1 to the case with $\bar{s} \neq 0$ as follows: Let $\tilde{\lambda} = \lambda e^{-\chi \bar{s}} / (\lambda e^{-\chi \bar{s}} + 1 - \lambda)$ denote a risk-adjusted prior that $\pi = 1$. With $\bar{s} \neq 0$, we can replace λ with $\tilde{\lambda}$ and go through the same steps using the risk-adjusted prior $\tilde{\lambda}$. However, when computing $\mathbb{E}(P(z)) - \lambda$ one must correct for the gap between λ and $\tilde{\lambda}$, which yields the additional risk premium term $(\lambda - \tilde{\lambda}) R$, 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 \bar{s} . The contribution of this term is illustrated in Figure 2 through the dashed horizontal line, corresponding

¹⁸In 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.

to the average price when $\bar{s} = 0$ (thick solid horizontal line) minus the risk premium that arises when $\bar{s} > 0$ (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 even for securities in zero expected supply. However, this example relies heavily on the specific assumptions of 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 and therefore doesn't clearly separate payoff uncertainty and asymmetry. Hence our interpretation of Δ as a price premium linked to skewness is at best suggestive. Our next example relaxes the present payoff assumptions and formally establishes a tight link between the expected price premium and the asymmetry of asset payoffs.

2.3 The risk-neutral, normal model

Our third example generalizes 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 empirical predictions and tie model parameters more closely to observables.

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.¹⁹

We consider the following specialization of the general model: (i) normally distributed fundamentals $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$ and private signals $x_i|\theta \sim \mathcal{N}(\theta, \beta^{-1})$, (ii) risk-neutral preferences $U(w) = w$ with positions limited by $[d_L, d_H] = [0, 1]$, and (iii) stochastic asset supply $s = \Phi(u)$, where $\Phi(\cdot)$ is the cdf of a standard normal distribution, and $u \sim \mathcal{N}(0, \sigma_u^2)$. Importantly, our analysis imposes no restriction on the payoff function $\pi(\cdot)$, which allows us to offer asset pricing predictions for

¹⁹In 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.

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).²⁰ It keeps the updating problem tractable by preserving normality of the investors' posterior beliefs.

Once again, 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 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 $P = \mathbb{E}(\pi(\theta) | \bar{x}, P)$ 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 leads to the following proposition:

Proposition 2 : *The equilibrium price in the unique price-monotone equilibrium is given by*

$$P_\pi(z) = \mathbb{E}(\pi(\theta) | x = z, z) = \int \pi(\theta) d\Phi \left(\frac{\theta - \hat{\gamma}z}{\sqrt{1 - \hat{\gamma}\sigma_\theta}} \right) \quad (9)$$

where $\hat{\gamma} \equiv \frac{\beta + \tau}{1/\sigma_\theta^2 + \beta + \tau}$.

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

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

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

²⁰Goldstein 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.

more strongly to changes in the sufficient statistic z than is warranted purely by the information conveyed through it. This occurs through a shift in the identity of this marginal investor that is required to accommodate supply shocks and clear the market.

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 $\theta|x = z, z \sim \mathcal{N}(\hat{\gamma} \cdot z, \sigma_\theta^2(1 - \hat{\gamma}))$ with the prior $z \sim \mathcal{N}(0, \sigma_\theta^2/\gamma)$ to show that the risk-neutral measure $\hat{H}(\cdot)$ is normal with mean 0 and variance

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

This representation is identical to the one found in the CARA-normal example, and confirms that $\hat{H}(\cdot)$ has the same mean but strictly higher variance than $H(\cdot)$. The change in probability measures $m(\theta) = \hat{h}(\cdot)/h(\cdot)$ is thus, once again, log-quadratic, or U-shaped, in fundamentals. Taking expectations we represent the expected price premium $W(\pi, \hat{\sigma}_\theta) \equiv \mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta))$ as the expected value of the mean-preserving spread between $\hat{H}(\cdot)$ and $H(\cdot)$:

$$W(\pi, \hat{\sigma}_\theta) = \int_{-\infty}^{\infty} \left(\pi\left(\frac{\hat{\sigma}_\theta}{\sigma_\theta}\theta\right) - \pi(\theta) \right) d\Phi\left(\frac{\theta}{\sigma_\theta}\right). \quad (12)$$

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 : (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 when the fundamental and the return are both observable (for example in the case of defaultable bonds or options). In addition, if $H(\cdot)$ is symmetric, upside

and downside risk relate to existing quantile-based measures of skewness that measure whether a distribution of payoffs is more spread out above or below its median. A security that is dominated by upside (downside) risk has positive (negative) third-moment skewness.²¹ 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 1.²²

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 and substantiates our interpretation of $W(\pi; \hat{\sigma}_\theta)$ as a premium for skewness.

Proposition 3 : (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, $\lim_{\hat{\sigma}_\theta \rightarrow \sigma_\theta} |W(\pi; \hat{\sigma}_\theta)| = 0$; and $\lim_{\hat{\sigma}_\theta \rightarrow \infty} |W(\pi; \hat{\sigma}_\theta)| = \infty$, whenever $\lim_{\theta \rightarrow \infty} |\pi(\theta) + \pi(-\theta)| = \infty$.*

(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 relates the expected price premium from noisy information aggregation to asymmetry in the payoff function π and 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 $\hat{\sigma}_\theta \rightarrow \sigma_\theta$. The latter case arises, for example, when private signals become infinitely noisy ($\beta \rightarrow 0$), or the price signal infinitely precise ($\sigma_u^2 \rightarrow 0$). Finally, the price premium has increasing differences in upside risk and EWTR, so the marginal impact of information aggregation frictions is stronger for more asymmetric securities. 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.

²¹With a symmetric prior, a security is an upside (downside risk) according to Definition 1 if and only if $\min_{\theta \geq 0} \chi(\theta) \geq 0$ ($\max_{\theta \geq 0} \chi(\theta) \leq 0$), where $\chi(\theta) = \frac{\pi(\theta) + \pi(-\theta) - 2\pi(0)}{\pi(\theta) - \pi(-\theta)}$. See Groeneveld and Meeden (1984) for mathematical properties of quantile-based skewness coefficients of the form given by $\chi(\theta)$, with Bowley's measure of skewness corresponding to $\theta = H^{-1}(3/4)$ and Kelly's measure of skewness corresponding to $\theta = H^{-1}(0.9)$.

²²Indeed 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.

So far, the risk-neutral model abstracted from average supply effects like the risk premium in the CARA-normal and CARA-binary models. We can reintroduce them by assuming that supply takes the form $s = \Phi(u)$, where $u \sim \mathcal{N}(\bar{u}, \sigma_u^2)$, 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.

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 now argue that EWTR can be represented in terms of two observable 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.

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

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

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:

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

Representation (14) shows that EWTR increases with forecast dispersion and is U-shaped in forecast accuracy $\hat{\gamma}$, reaching a minimum at $\hat{\gamma} = 1/2$, and diverging to infinity when $\hat{\gamma}(1 - \hat{\gamma}) \rightarrow \tilde{D}^2$. However, for most of its range, the effect of $\hat{\gamma}$ on $\hat{\sigma}_\theta/\sigma_\theta$ is very mild as we move away from $\hat{\gamma} = 0.5$.²³ Forecast dispersion \tilde{D} thus emerges as a natural empirical proxy variable for EWTR, and the comparative statics with respect to $\hat{\sigma}_\theta$ in proposition 3 imply 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 disagreement generates a positive option value of resale regardless of a security's return structure.

²³For example, a value of $\tilde{D} = 0.2$ implies a minimum value of $\hat{\sigma}_\theta/\sigma_\theta = 1.0235$, but does not exceed 1.0244 for values of $\hat{\gamma}$ in the range of [0.25, 0.75].

3 Returns to skewness and disagreement

3.1 Empirical evidence

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

Excess Weight on Tail Risks: Starting with Jackwerth (2000) and Ait-Sahalia and Lo (2000), a substantial empirical literature documents that pricing kernels recovered from option prices on stock indices appear to be non-monotone, and in certain cases U-shaped (the *pricing kernel puzzle*). 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) use pricing implications for pure upside risks. Focusing on index options and aggregate market returns, all three find strong evidence in support of pricing kernels that are U-shaped or upwards-sloping on the upside. Although we relate our theory to cross-sectional return premia, these papers provide at least suggestive evidence consistent with our model. Carr and Wu (2008) document that the risk-neutral return variance is an upwards-biased predictor of the true return variance (the *variance premium puzzle*). Christoffersen et al. (2013) discuss the relation between the pricing kernel and variance premium puzzles and argue that the implied variance of returns overestimates the true variance of returns by about 12 to 18%. This 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, which is similar to the estimate we obtain below by matching forecast dispersion and accuracy.

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 $W(\pi^{Q_H}; \hat{\sigma}_\theta) - W(\pi^{Q_L}; \hat{\sigma}_\theta)$, 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 $W(\pi; \hat{\sigma}_\theta^{q_H}) - W(\pi; \hat{\sigma}_\theta^{q_L})$, 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, yet pay significantly larger return premia.²⁴

Existing explanations for these empirical findings rely on investor preferences for positive skewness.²⁵ Proposition 3 offers an alternative explanation based on noisy information aggregation and EWTR. In contrast to preference-based theories, this explanation also links returns to skewness and disagreement. We return to the comparison between preference- and information-based returns to skewness in section 5 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. Diether et al. (2002; DMS henceforth) sort stocks by the dispersion of earnings forecasts across analysts

²⁴Huang 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. Other authors have linked credit spreads and 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.

²⁵In 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.

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 evidence as consistent with the hypothesis of Miller (1977) of investor disagreement interacting with short-sales constraints.²⁶ 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 by noting that studies that find negative returns to disagreement focus on securities 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.²⁷ 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

²⁶They rule out a risk-based explanation for the anomaly by controlling for stocks exposure to standard risk factors.

²⁷The 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 measures implied by our theory.

−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 markets.²⁸

For bond markets, GH report that the effect of disagreement on spreads and yields doubles in high leverage or low-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.²⁹

The ability to account for returns to skewness and disagreement simultaneously in equity and bond markets 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 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 connection with Predictions 2 and 3.

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.

²⁸Zhang (2013) documents strong positive correlation between book-to-market ratios and skewness of returns, and shows that book-to-market ratios have significantly lower explanatory power for returns after controlling for skewness.

²⁹Hou 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 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.

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 on asset payoffs and investor disagreement. The objective of the exercise is two-fold: first, to illustrate 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. 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, while keeping in mind that our main contribution remains theoretical.

Forecast dispersion and EWTR in the data:

In line with the empirical literature reviewed above, we interpret the asset fundamental as a firm’s earnings, and use I.B.E.S. data of analyst earnings forecasts to derive measures of normalized forecast dispersion \tilde{D} and forecast accuracy $\hat{\gamma}$ for a cross-section of listed firms (see online appendix for a more detailed description). We base our estimates 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 1987-1998.³⁰ The data report measures of raw forecast dispersion, a consensus or average earnings forecast, and realized earnings per share for each firm-year in the sample. We use the time series of realized earnings and earnings forecast errors to compute estimates of earnings volatility σ_θ , forecast accuracy $\hat{\gamma}$, and (normalized) forecast dispersion \tilde{D} for each firm, and substitute these estimates into equation (14) to construct a firm-level estimate of EWTR.³¹

This approach 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 If public disagreement in analyst forecasts is broadly representative of dispersion in private investor forecasts and positively correlated with the latter, the qualitative

³⁰We thank these authors, as well as Ludwig Straub and Robert Ulbricht, for sharing the data used in this section.

³¹For robustness, we replicate our estimates on a second sample from Straub and Ulbricht (2023), who use a much longer 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 to reduce noise in our estimates of forecast dispersion and accuracy. Estimates of forecast dispersion, accuracy and EWTR are qualitatively similar in the two samples, which gives us some confidence in the robustness of our numerical examples. See the online appendix for further details.

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

	Mean	10%	30%	Median	70%	90%
Dispersion (\tilde{D})	0.192	0.058	0.116	0.173	0.246	0.3659
Accuracy ($\hat{\gamma}$)	0.747	0.402	0.667	0.814	0.915	0.975
EWTR ($\frac{\hat{\sigma}_e}{\sigma_e}$)	1.102	1.002	1.009	1.023	1.055	1.195

Table 1: Forecast dispersion, accuracy, and implied EWTR

predictions discussed above will remain valid.

But using analyst’s earnings forecasts to quantify EWTR further requires 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, median, and 10th, 30th, 70th and 90th percentiles 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 EWTR is very low. However, forecast dispersion and EWTR are both fairly significant in the top quintile of the distribution: the 90th percentile of forecast dispersion and EWTR is about twice as high as the mean. 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%, while median EWTR is estimated around 2.3%.

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 consistent with the cross-sectional estimates from GH (Table 1).

³³Consistent with our view that analyst forecast disagreement proxies for forecast dispersion among a wider pool of investors, Kovbasyuk and Pagano (2022) provide anecdotal and empirical evidence to suggest that analysts have an incentive to disclose their information after taking certain positions to realize gains from informed trading.

			Mean	Q1	Q3	Q5	Q3-Q1	Q5-Q1
	<i>Targets</i>	<i>Skew (BMV)</i>	<i>0.851</i>	<i>0.167</i>	<i>0.565</i>	<i>1.629</i>		
	<i>Disp (GH)</i>	Data	BMV				-9	-67
Mean	<i>0.315</i>	GH	<i>-9.5</i>	<i>-0.8</i>	<i>-4.6</i>	<i>-25.8</i>	<i>-3.8</i>	<i>-25.0</i>
Q3-Q1	<i>0.173</i>	-18	<i>-1.5</i>	<i>-0.1</i>	<i>-0.7</i>	<i>-4.0</i>	<i>-0.6</i>	<i>-3.9</i>
Q5-Q1	<i>0.359</i>	-26	<i>-18.2</i>	<i>-1.5</i>	<i>-8.8</i>	<i>-49.2</i>	<i>-7.4</i>	<i>-47.8</i>

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

Table 2 compares the empirical and model-implied returns.³⁴ The columns refer to different return distributions which are respectively calibrated to match the mean, bottom, median and top quintile bins of skewness reported in BMV (Table 1 and Table 3, column 4). For each of these securities bins we also match the idiosyncratic volatility reported in BMV (Table 1 and Table 3, column 5). The last two columns refer to excess return of median and top quintile securities over the bottom quintile. For these securities, we report empirical excess returns from BMV (Table 3, column 1) in the row labelled “Data”. The rows in Table 2 then refer to different levels of EWTR. Formally, we set $\hat{\gamma}$ to the sample mean of 0.75, and vary \tilde{D} to match the sample mean of EWTR ($\frac{\hat{\sigma}_\theta}{\sigma_\theta} = 1.1$) and the midpoints of the different quintile bins of the cross-sectional distribution of forecast dispersion in Table 1.³⁵ The corresponding dispersion levels are reported in the column labeled “Targets”. Since return premia in the lowest dispersion quintile bin are negligible, we only report excess returns for the median and top quintiles over the bottom. We report empirical counterparts for these excess returns from GH (Table 12, Panel A) in the column labeled “Data”. The numbers in italics then report model-implied expected returns for securities matching the skewness target by column and the dispersion target by row.

At the mean level of dispersion, our model generates excess 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. These returns to skewness vary substantially across forecast dispersion levels, from no more than a couple of basis points when dispersion is below its median all the way to 49 bp in the top dispersion quintile.

Similarly, at the mean level of skewness, the model implied returns to disagreement are -18.2

³⁴For conciseness we only report the top, median and bottom quintiles, along with estimates for mean skewness and mean forecast dispersion. The complete version of the table is in the online appendix.

³⁵Due to the strong convexity of EWTR in \tilde{D} the mean EWTR appears to capture better an average level of frictions than the average forecast dispersion.

bp, or about 70% of the returns to disagreement of reported for equity in GH (26 bp). But these average returns to disagreement mask a wide variation, from less than 2 bp in the lowest skewness quintile, all the way to 49 bp in the highest.

These results suggest that noisy information aggregation may account for a sizeable proportion of observed returns to skewness and disagreement. They further suggest very strong interaction effects, since most of the return premia are concentrated in the top skewness and top dispersion quintiles, generating cross-sectional variation in returns to disagreement and returns to skewness of up to 48 bp, which correspond to 5.9% 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 to the general model set-up introduced in section 2. Suppose that there exists a price-monotone equilibrium $\{P(\theta, s); d(x, P); H(\cdot|P)\}$.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 $Pr(P \leq P'|\theta) = Pr(D(\theta, P) \geq D(\theta, P')) = Pr(s \geq D(\theta, P')) = 1 - G(D(\theta, P'))$. Therefore conditional on θ , z is distributed according to

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

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

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

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

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

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

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

³⁷The 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.

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 $\frac{U'(\bar{D}(\pi(\theta) - P_\pi(z)))}{\mathbb{E}(U'(\bar{D}(\pi(\theta) - P_\pi(z)))m^I(\theta, z)|z)}$ that weighs states according to the investor's marginal utility of 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).³⁸ 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} : $\mathbb{E}(P_\pi(z)) = \hat{\mathbb{E}}(\pi(\theta)) = \mathbb{E}(\pi(\theta) \cdot m(\theta))$, where $m(\theta) \equiv \frac{\hat{h}(\theta)}{h(\theta)} = \mathbb{E}(m(\theta, z)|\theta)$, with $\mathbb{E}(m(\theta)) = 1$.

Excess Weight on Tail Risks: 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 $\lim_{z \rightarrow -\infty} H(z+k|z) = 0$ and $\lim_{z \rightarrow \infty} H(z+k|z) = 1$, for any finite k . Then $\lim_{\theta \rightarrow \infty} m(\theta) = \lim_{\theta \rightarrow -\infty} m(\theta) = \infty$.*

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 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

³⁸Since $\bar{D} \in (d_L, d_H)$, $\bar{D} = 0$ requires that $d_L < 0 < d_H$.

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 $\lim_{\theta \rightarrow \infty} m(\theta) = \lim_{\theta \rightarrow -\infty} m(\theta) = \infty$. 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 $\bar{\tau}_\psi \geq -\left(\frac{\psi''}{\psi} - \left(\frac{\psi'}{\psi}\right)^2\right) \geq \underline{\tau}_\psi > 0$. Suppose further that*

$$\frac{\underline{\tau}_f + \underline{\tau}_\psi}{\underline{\tau}_f + \underline{\tau}_\psi + \underline{\tau}_h} \underline{\tau}_h > \bar{\tau}_h \frac{\bar{\tau}_\psi}{\bar{\tau}_\psi + \bar{\tau}_h}. \quad (17)$$

Then, $m(\cdot)$ is strictly log-convex. 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).

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 satisfied whenever condition (17) holds. In addition, Proposition 5 provides a continuity result for $\frac{d}{d\theta} \frac{m'(\theta)}{m(\theta)}$ which implies that equation (11) provides a good approximation for excess weight on tail risks, if fundamental and signal densities are approximately Gaussian.

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

These sufficient conditions can be decomposed into three parts. 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, condition (17) imposes a lower bound on the log-curvature, or equivalently, the informativeness of the traders' private signal densities.⁴⁰ Variation in log-curvature implies that posterior uncertainty may vary across states and signal realizations. Condition (17) insures that 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. These same bounds on log-curvature of the distribution can then be used to establish bounds on $\frac{d}{d\theta} \frac{m'(\theta)}{m(\theta)}$ and show that $\frac{d}{d\theta} \frac{m'(\theta)}{m(\theta)}$ converges to its linear-normal counterpart given by equation (11) as the fundamental and signal distributions converge to normals.

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 $P'_\pi(z) \frac{\partial z(\theta, s)}{\partial \theta} = -D_\theta(\theta, P_\pi(z)) / D_P(\theta, P_\pi(z))$. In general, $P'_\pi(z) = -d_x(z, P) / d_P(z, P)$ 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 overturn the implication of condition (17) that the risk-neutral posterior responds more strongly to changes in z than the objective posterior.

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

⁴⁰If the prior is normal ($\bar{\tau}_h = \underline{\tau}_h$), condition (17) simplifies to $\underline{\tau}_f > \bar{\tau}_\psi - \underline{\tau}_\psi$, i.e. the log-curvature of private signal is uniformly higher than the variation in log-curvature of the sufficient statistic z .

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 $m(\cdot)$ 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 m , we write the expected price premium as $W(\pi, m) = \mathbb{E}(P_\pi(z)) - \mathbb{E}(\pi(\theta))$. Applying the partial order on returns given by Definition 1, we can then directly generalize the comparative statics predictions of Proposition 3 if $H(\cdot)$ and $m(\cdot)$ are symmetric and centered around the same mean (say, around 0):

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 $W(\pi_1, m_2) \geq W(\pi_1, m_1) \geq 0$ ($W(\pi_1, m_2) \leq W(\pi_1, m_1) \leq 0$). Moreover, $\lim_{m \rightarrow 1} W(\pi, m) = 0$, and $|W(\pi, m)|$ grows arbitrarily large if m has arbitrarily large excess weight on tail risk and $\lim_{\theta \rightarrow \infty} |\pi(\theta) + \pi(-\theta)| = \infty$.*

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

We can still obtain variants of these comparative statics results upon relaxing the symmetry and equal means assumptions underlying Theorem 2. For example, the theorem continues to hold with equal means but asymmetric distributions, 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, 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:

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

where $\delta \equiv \int \theta d\hat{H}(\theta) - \int \theta dH(\theta)$. 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 $\int_{-\infty}^{\infty} (H(\theta) - \hat{H}(\theta + \delta)) d\pi(\theta)$ instead captures the excess weight on tail risks implied by the risk-neutral distribution, controlling for the

difference in means. 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 $\delta \geq 0$ for convex return functions (upside risks) or $\delta \leq 0$ for convex return functions (downside risks). If $\hat{H}(\theta + \delta)$ has EWTR over $H(\theta)$, we can then apply the above comparative statics results to this second term.

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 results 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 reasons. First, short of having explicit sufficient conditions for EWTR based on exogenous model primitives, this procedure provides a method for verifying, at least numerically, whether the model-implied risk-neutral measure satisfies the sufficient conditions for EWTR described in Proposition 4. 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: $\Psi^{(1)}(P|\theta) \equiv 1 - G(D^{(0)}(\theta, P'))$. Conditional price distribution $\Psi^{(1)}$ is then used 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 solve a model with binary payoffs and CRRA preferences numerically, and confirm the analytical results for the CARA-binary and risk-neutral

⁴¹Bernardo 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. We instead solve explicitly 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.

examples in section 2: on average downside risk is under-priced, upside risk is over-priced.⁴²

5 Extensions: uninformed traders and multiple securities

We extend our model to include uninformed traders, as well as multiple securities. Our general set up is the same as in section 4, with two modifications to generalize prior assumptions.

(i) **Assets:** There are $N < \infty$ securities with payoff vector $\pi(\theta) \in \mathbb{R}^N$ where $\theta \in \mathbb{R}^N$ is distributed according to a smooth prior cdf. $H(\cdot) : \mathbb{R}^N \rightarrow [0, 1]$. The payoff functions $\pi(\theta)$ is non-decreasing in θ . The asset supply $s \in \mathbb{R}^N$ is drawn from an N -dimensional rectangle $\mathcal{D} = \{s \in \mathbb{R}^N : s_n \in [d_{L,n}, d_{H,n}]\} \subseteq \mathbb{R}^N$ according to a smooth cdf. G . We assume that $\text{int}(\mathcal{D})$ is non-empty and $d_{L,n} \leq 0 \leq d_{H,n}$ for all n .

(ii) **Investors:** A measure $\kappa_I \in [0, 1]$ of traders are informed, and a measure $\kappa_U \geq 0$ are uninformed, with $\kappa_U = 1 - \kappa_I$ unless otherwise noted. Informed traders receive a signal vector $x = \theta + \varepsilon^i \in \mathbb{R}^N$, where ε^i is iid across agents, component-wise independent,⁴³ and distributed according to cdf F with unbounded support and strictly log-concave and unbounded pdf f . Informed and uninformed traders' preferences are given by increasing concave utility function $U((\pi(\theta) - P)'d)$ over realized gains and losses $(\pi(\theta) - P)'d$, where the investors' portfolio $d \in \mathcal{D}$ is restricted by position limits $[d_{L,n}, d_{H,n}]$, for each security n , and $P \in \mathbb{R}^N$ represents the price vector.

A *Perfect Bayesian Equilibrium* consists of a set of portfolios $d^I(x, P)$, $d^U(P) \in \mathcal{D}$ for informed and uninformed traders, a vector-valued price function $P(\theta, s)$, and posterior beliefs $H(\cdot|x, P)$ and $H(\cdot|P)$ such that (i) $d^I(x, P)$ and $d^U(P)$ are optimal given $H(\cdot|x, P)$ and $H(\cdot|P)$; (ii) $P(\theta, s)$ clears all N markets; and (iii) $H(\cdot|P)$ and $H(\cdot|x, P)$ satisfy Bayes' rule whenever applicable.

This generalization nests our model from section 4 when $N = 1$ (single asset) and $\kappa_U = 0$ (no uninformed traders). We now relax each of these assumptions in turn.

Uninformed Traders

Introducing uninformed traders into the model enables us to generalize our earlier observation that EWTR is generated by stochastic risk premia due to random supply shocks, which are then amplified by the presence of private information.

Suppose that $N = 1$ (single asset) but allow for the presence of uninformed traders ($\kappa_U > 0$).

⁴²We make our matlab code available for the CRRA, binary payoff case. Aggregate demand monotonicity w.r.t. prices is not guaranteed for generic preferences and payoff structures (see, e.g. Barlevy and Veronesi, 2003). But without strict monotonicity, $Pr(P \leq P'|\theta) = Pr(u \geq D(\theta, P'))$ no longer holds for all prices, and our solution method would not work. We verify numerically that demand monotonicity is satisfied for the CRRA, binary example.

⁴³We say a random vector y is *component-wise independent* if y_n is independent of $y_{n'}$ for all $n \neq n'$.

To simplify the exposition, we also assume that $\mathbb{E}(s) = 0$ throughout this part. In the online appendix, we extend the CARA-normal and CARA-binary examples to include uninformed traders and show that the same equilibrium characterization applies with the private signal precision β replaced by $\beta\kappa_I$. In other words, a single composite parameter summarizes the extent of informed trading both on the extensive (fraction of informed) and the intensive (signal precision) margin. The magnitude of the updating wedge, the excess weight on tail risks and, in the binary case, the skewness premium, now depend on the measure of informed traders, as well as their signal precision. We can then extend the earlier limiting results for the CARA-normal and CARA-binary models to κ_I (rather than β) going to 0. As $\kappa_I \rightarrow 0$, the price becomes completely uninformative, and EWTR in the CARA-normal and skewness premia in the CARA-binary models remain but scale with $\chi^2\sigma_s^2$, the variance of supply shocks. In the alternative limit in which $\kappa_I \rightarrow 0$, but price informativeness τ is kept constant, EWTR and skewness premia scale with κ_I or $\chi\sigma_s$, or the standard deviation of supply shocks.⁴⁴ *Ceteris paribus*, if σ_s^2 is sufficiently small, adding a small measure of informed traders to a market without informed traders therefore amplifies EWTR and skewness premia.

The uninformed traders' first-order condition in the CARA-normal case with $\bar{s} = 0$ further allows us to decompose the price into an expected dividend and a risk premium:

$$P(z) = \mathbb{E}(\theta|z) - \chi d^U(P(z)) \text{Var}(\theta|z) \quad (18)$$

where $\text{Var}(\theta|z) = (1/\sigma_\theta^2 + \tau)^{-1}$. Moreover, since $\mathbb{E}(\theta|x=z, z) > P(z) > \mathbb{E}(\theta|z)$ when $z > 0$ and $\mathbb{E}(\theta|x=z, z) < P(z) < \mathbb{E}(\theta|z)$ when $z < 0$, the uninformed traders take the opposite side of the market compared to the marginal informed trader: $d^U(P(z)) < 0 < d^I(z, P(z))$ when $z > 0$ and $d^U(P(z)) > 0 > d^I(z, P(z))$ when $z < 0$. The updating wedge embedded in the price therefore constitutes a *counter-cyclical risk premium* for the uninformed traders since their exposure $d^U(P(z))$ is negatively correlated with the sufficient statistic z , and hence with fundamentals. This result arises even if $\kappa_I \rightarrow 1$ and there are no uninformed traders in the market.

We now generalize these observations and limit results to our general model. Using a second-order Taylor expansion of the uninformed traders' first-order condition and then taking prior expectations yields the following approximation of the expected price premium:

$$W(\pi; m) \approx -\chi \text{cov}(\text{Var}(\pi(\theta)|z), d^U(P(z))) + \frac{\alpha}{2} \cdot \mathbb{E}\left(\text{Skew}(\pi(\theta)|z) \cdot d^U(P(z))^2\right) \quad (19)$$

⁴⁴This second limit is equivalent to keeping the measure of informed traders $\kappa_I = 1$ and the distribution of supply shocks fixed but letting the measure of uninformed traders κ_U tend to infinity. Keeping price informativeness constant is thus perhaps best interpreted as a case with unlimited entry by uninformed but risk-averse arbitrageurs. In this case, the uninformed traders positions, and the skewness premia, must scale with $1/\kappa_U$.

where $Skew(\pi(\theta)|z) = \mathbb{E}\left((\pi(\theta) - \mathbb{E}(\pi(\theta)|z))^3|z\right)$ denotes expected skewness of asset payoffs, $\chi = -\frac{U''(0)}{U'(0)}$ the absolute risk aversion coefficient, and $\alpha = \frac{U'''(0)}{U''(0)} > 0$ represents downside risk aversion.⁴⁵ Equation (19) decomposes the expected price premium into two components, an *information-based* (first term) and a *preference-based* skewness premium (second term).

The *information-based* skewness premium $-\chi cov(Var(\pi(\theta)|z), d^U(P(z)))$ depends on the uninformed traders' risk aversion χ and the comovement between their exposure $d^U(P(z))$ and posterior uncertainty $Var(\pi(\theta)|z)$. $d^U(P(z))$ is counter-cyclical if posterior beliefs are sufficiently well-behaved.⁴⁶ For downside risks, posterior uncertainty is counter-cyclical resulting in a positive covariance with exposure and a negative price premium, while for upside risks, uncertainty is pro-cyclical thus resulting in a negative covariance and a positive price premium.

The *preference-based* skewness premium $\frac{\alpha}{2} \cdot \mathbb{E}\left(Skew(\pi(\theta)|z) \cdot d^U(P(z))^2\right)$ depends on downside risk aversion α , squared exposure $d^U(P(z))^2$, and expected skewness $Skew(\pi(\theta)|z)$. The parameter α captures investors' attitudes towards asymmetric risks, while the expected skewness identifies relative magnitudes of upside and downside risks.

We now revisit the two limit cases discussed above. If $\kappa_I \rightarrow 0$, while holding other primitives fixed, $\lim_{\kappa_I \rightarrow 0} d^U(P(z)) = s$, z becomes completely uninformative, and $\mathbb{E}(\pi(\theta)|z)$, $Var(\pi(\theta)|z)$ and $Skew(\pi(\theta)|z)$ converge to their unconditional counterparts. The information-based premium $cov(Var(\pi(\theta)|z), d^U(P(z)))$ vanishes, and $\lim_{\kappa_I \rightarrow 0} W(\pi; m) \approx \frac{\alpha}{2} \cdot Skew(\pi(\theta)) \cdot \sigma_s^2$: with downside risk aversion ($\alpha > 0$), downside risks generate a negative price premium and upside risks a positive one, which scales with σ_s^2 . The preference-based premium vanishes if risks are symmetric ($Skew(\pi(\theta)) = 0$), or investors are neutral towards downside risk ($\alpha = 0$, quadratic preferences).

In contrast, if $\kappa_I \rightarrow 0$ and the distribution of supply shocks is also scaled by κ_I (or $\kappa_I = 1$ and G are fixed but $\kappa_U \rightarrow \infty$), z remains informative about θ , $Var(\pi(\theta)|z)$ converges to a non-deterministic limit, and $d^U(P(z))$ must then scale with κ_I or $1/\kappa_U$ to satisfy market-clearing. $W(\pi; m) \approx -\chi \cdot cov(Var(\pi(\theta)|z), d^U(P(z)))$ scales with κ_I , or equivalently σ_s : In this limit, $W(\pi; m)$ is primarily driven by the information-based premium.

In the online appendix, we show very generally that the expected price premium is of order σ_s^2 if $\kappa_I = 0$, but of order σ_s , when $\sigma_s \rightarrow 0$ but price remain informative. Informed trading thus amplifies the skewness premium relative to a benchmark with no dispersed information. This

⁴⁵See, e.g. Modica and Scarsini (2005). The measure α can also be represented as the product of prudence $-\frac{U'''(0)}{U''(0)}$ (Kimball 1990) and risk aversion χ . With CARA preference, $\alpha = \chi^2 > 0$.

⁴⁶Well-behaved in the sense that $\mathbb{E}(\theta|x=z, z) - \mathbb{E}(\theta|z)$ is monotonic, where z is constructed as the private signal of a marginal trade with 0 asset holdings, $d^I(z, P(z)) = 0$.

helps explain why we found quantitatively significant returns to skewness and disagreement in the numerical example in section 3.

Multiple Securities

Here we allow for $N > 1$, extending our analysis of noisy information aggregation to multi-asset markets. This becomes relevant when multiple securities are linked to the same fundamentals, but also helps to clarify conditions under which trading is independent across different securities.

Consider a price monotone equilibrium with informed trader demand $d^I(z, P)$ and fix $d_0 \in \text{int}(\mathcal{D})$. Suppose that there exists a unique vector $z(P) \in \mathbb{R}^N$ such that $d^I(z, P) = d_0$, and that this vector is invertible. Using the equilibrium price function, we write $z(\theta, s) \equiv z(P(\theta, s))$ as the sufficient statistic vector. Invertibility of $z(P)$ coupled with the informed traders' first-order condition evaluated at $d^I(z, P) = d_0$ then implies that $P(z) = \mathbb{E}(\pi(\theta) m(\theta, z) | z)$, where

$$m(\theta, z) = \frac{U'((\pi(\theta) - P(z))' d_0)}{\mathbb{E}(U'((\pi(\theta) - P(z))' d_0) | x = z, z)} m^I(\theta, z)$$

with $m^I(\theta, z) = \frac{h(\theta|x=z,z)}{h(\theta|z)} = \frac{f(z-\theta)}{\mathbb{E}(f(z-\theta)|z)}$, and $\mathbb{E}(m(\theta, z) | z) = \mathbb{E}(m^I(\theta, z) | z) = 1$. In other words, if $z(P)$ is invertible, we recover a sufficient statistic representation of the equilibrium price vector with a vector-valued sufficient statistic which inherits the same updating wedge $m^I(\theta, z)$ as its scalar counterpart in the one asset model, along with a risk adjustment factor. Moreover, the same pricing kernel $m(\theta, z)$ is used to price all securities simultaneously. The expected price vector satisfies $\mathbb{E}(P(z)) = \mathbb{E}(\pi(\theta) m(\theta))$, where $m(\theta) = \mathbb{E}(m(\theta, z) | \theta)$ and $\mathbb{E}(m(\theta)) = 1$. For $d_0 = 0$, we obtain $m(\theta, z) = m^I(\theta, z)$ and $P(z) = \mathbb{E}(\pi(\theta) | x = z, z)$.

Therefore, if $z(P)$ exists and is invertible, the characterization of the updating wedge and the risk-neutral probability measures to multiple securities. Further assumptions on primitives are only required to (i) guarantee the existence of a marginal investor who finds it optimal to hold portfolio d_0 , and (ii) show that this marginal investor's private signal vector is informationally equivalent to the equilibrium price. By setting $d_0 \in \text{int}(\mathcal{D})$, we implicitly assume that informed investors freely arbitrage across the different securities: in other words, they are not specialists in specific markets. This cross-market arbitrage is key for insuring that there exists a common pricing kernel for all securities.⁴⁷ Under these conditions, the marginal trader's private signal vector z is a

⁴⁷Similarly, the characterization of a common pricing kernel $m^U(\theta, z) = \frac{U'((\pi(\theta) - P(z))' d^U(P(z)))}{\mathbb{E}(U'((\pi(\theta) - P(z))' d^U(P(z))) | z)}$ from the uninformed traders' first-order condition $P(z) = \mathbb{E}(\pi(\theta) m^U(\theta, z) | z)$ implicitly assumes that uninformed traders can freely arbitrage across all markets. In the online appendix we use this characterization of $m^U(\theta, z)$ to generalize our previous discussion of counter-cyclical risk premia and information vs. preference based skewness premia from one to many securities. We also obtain a CAPM-style representation of asset prices similar to Andrei et al. (2022) who study empirical properties of the CAPM with linear/normal asset returns and noisy information aggregation.

sufficient statistic for the information content of the price vector and the price vector inherits the multi-dimensional analogue of the updating wedge and excess weight on tail risks.

To conclude we discuss two interesting special cases.

1. *Independent assets:* We first identify conditions under which information aggregation and trading is independent across assets. Sufficient statistics z and prices $P(z)$ are component-wise independent, if fundamentals θ , supply s and signal noise ε^i are all component-wise independent, $\frac{\partial \pi_n}{\partial \theta_{n'}} = 0$ for $n \neq n'$, and traders have CARA preferences. In other words, prices are determined independently in each market and aggregate only asset-specific information, if each asset has its own fundamental (π_n only depends on θ_n), fundamentals and signals are independent across markets, and all traders have constant absolute risk aversion. Independence of prices across markets arises if $m(\theta, z)$ can be decomposed into a product of market-specific pricing kernels $m_n(\theta_n, z_n)$. For $m^I(\theta, z)$ this requires fundamentals and private and market signals to be independent across markets, i.e. there are no informational spill-overs from one market to another. CARA preferences further imply that trading gains and losses in one market do not affect investors' willingness to take risks in other markets, hence the risk adjustment also admits a multiplicative decomposition.

This case is especially useful for studying cross-sectional return implications from noisy information aggregation. In Albagli et al. (2023), we solve a 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.

2. *Common fundamentals:* Consider now a limiting case where all fundamentals are perfectly correlated. For example, suppose that $\theta_n = \nu_n(1 - \rho) + \frac{\rho}{N} \sum_{m=1}^N \nu_m$, where ν_m are component-wise independent, and let $\rho \rightarrow 1$. Suppose also that supply shocks are component-wise independent to abstract from correlation in noise trading. Then, in the limit, each market price provides a separate independent signal about the common fundamental $\frac{1}{N} \sum_{m=1}^N \nu_m$. Markets then have informational spill-overs on each other, since a lower price in one market is perceived as conveying bad news about the likely payoffs in the other markets. In other words, z conveys a vector of independent signals about θ that is factored into prices. At the same time, the marginal investors private signal in each market responds more strongly to the market-specific shocks for any $\rho > 0$ to absorb market-specific supply variation that we assumed to be independent across markets. Hence we obtain a market-specific updating wedge in each market but informational spill-overs across markets.

Application (Modigliani-Miller Theorem): In the online appendix, we solve a two-asset version of the risk-neutral model with common fundamentals and apply it to a security design problem. Because the price premium is additive in cash flows for given m (i.e. $W(\pi_1 + \pi_2, m) = W(\pi_1, m) +$

$W(\pi_2, m)$), a version of the Modigliani-Miller Theorem obtains if information frictions are the same in both markets, or if there exists a common pricing kernel m . However, if the two markets are segmented and EWTR is larger in one market than the other, then the increasing difference property implies that the total market value of the cash flow is not invariant to how upside and downside risks are allocated to the two markets: the originator of a cash flow can take advantage of different degrees of information frictions in the two markets by separating upside from downside risks into equity and debt claims, and selling the equity claim in the market with the highest frictions, while either keeping the debt claim or selling it in the market with the lowest level of information frictions. These results follow directly from Proposition 3 and Theorem 2.

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 EWTR 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 how market frictions influence real decision-making by firms, households or policy makers. Albagli et al. (2023) study the interplay between noisy information aggregation and risk-taking incentives. Bassetto and Galli (2019) use our model to 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. These applications already suggest that our model may be useful to shed light on economic phenomena beyond asset prices.

References

- [1] Ait-Sahalia, Yacine and Andrew Lo (2000). “Nonparametric risk management and implied risk aversion,” *Journal of Econometrics*, 94 (1-2), 9-51.
- [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 (2023). “Imperfect Financial Markets and Investment Inefficiencies,” *American Economic Review*, forthcoming.
- [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 (2022). “Multi-Asset Noisy Rational Expectations Equilibrium with Contingent Claims,” *Review of Economic Studies*, 89 (5), 2445–2490.
- [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] Groeneveld, Richard A. and Glen Meeden (1984). “Measuring Skewness and Kurtosis.” *Journal of the Royal Statistical Society. Series D (The Statistician)*, 33 (4), 391-399.
- [34] Grossman, Sanford J. and Joseph E. Stiglitz (1980). “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70 (3), 393-408.
- [35] Güntay, Levent, and Dirk Hackbarth (2010). “Corporate Bond Credit Spreads and Forecast Dispersion.” *Journal of Banking and Finance*, 34 (10), 2328-2345.
- [36] He, Zhiguo and Konstantin Milbradt (2014). “Endogenous Liquidity and Defaultable Bonds.” *Econometrica*, 82 (4), 1443-1508.
- [37] Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski (2006). “Self-Fulfilling Currency Crises: The Role of Interest Rates,” *American Economic Review*, 96 (5), 1769-1787.
- [38] Hellwig, Martin (1980). “On the aggregation of Information in Competitive Markets,” *Journal of Economic Theory*, 22 (3), 477-498.

- [39] Hou, Kewei, Chen Xue and Lu Zhang (2020). “Replicating Anomalies,” *Review of Financial Studies*, 33 (5), 2019-2133.
- [40] 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.
- [41] Jackwerth, Jens (2000). “Recovering Risk Aversion from Option Prices and Realized Returns” *Review of Financial Studies*, 13 (2), 433–451.
- [42] Kovbasyuk, Sergey, and Marco Pagano (2022). “Advertizing Arbitrage,” *Review of Finance*, 26 (4), 799–827.
- [43] Kremer, Ilan (2002). “Information aggregation in common value auctions.” *Econometrica*, 70 (4), 1675-1682.
- [44] Malamud, Semyon (2015). “Noisy Arrow-Debreu Equilibria”. Working paper, EPFL.
- [45] McQuade, Timothy (2018). “Stochastic Volatility and Asset Pricing Puzzles”. Working paper, Stanford GSB.
- [46] Milgrom, Paul (1979). “A convergence theorem for competitive bidding with differential information.” *Econometrica*, 47 (3), 679-688.
- [47] Milgrom, Paul (1981). “Rational Expectations, Information Acquisition, and Competitive Bidding,” *Econometrica*, 49 (4), 921-943.
- [48] Miller, Edward (1977). “Risk, uncertainty, and divergence of opinion.” *Journal of Finance* 32 (4), 1151-1168.
- [49] Mitton, Todd and Keith Vorkink (2007). “Equilibrium Underdiversification and the Preference for Skewness”, *Review of Financial Studies*, 20 (4), 1255-1288.
- [50] Peress, Joel (2004). “Wealth, information acquisition, and portfolio choice.” *Review of Financial Studies*, 17 (3), 879-914.
- [51] Perry, Motty, and Philip J. Reny (2006). “Toward a strategic foundation for rational expectations equilibrium.” *Econometrica*, 74 (5), 1231-1269.
- [52] Pesendorfer, Wolfgang, and Jeroen M. Swinkels (1997). “The loser’s curse and information aggregation in common value auctions.” *Econometrica*, 65 (6), 1247-1281.

- [53] Saumard, Adrien, and Jon Wellner (2014). “Log-Concavity and Strong Log-Concavity: a Review,” working paper (arXiv:1404.5886).
- [54] Straub, Ludwig, and Robert Ulbricht (2023). “Endogenous Uncertainty and Credit Crunches,” working paper, MIT and Boston College.
- [55] Veldkamp, Laura (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.
- [56] Vives, Xavier (2008). *Information and Learning in Markets*. Princeton University Press.
- [57] Wilson, Robert (1977). “A bidding model of perfect competition,” *Review of Economic Studies*, 44 (3), 511-518.
- [58] 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.
- [59] Yu, Jialin (2011). “Disagreement and Return Predictability of Stock Portfolios,” *Journal of Financial Economics*, 99 (1), 162-183.
- [60] 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 $\hat{x}(P) = \hat{x}(P') = \hat{x}(P'')$ 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.⁴⁸ 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 $\hat{\sigma}_\theta \rightarrow \infty$, note that $\lim_{\hat{\sigma}_\theta \rightarrow \infty} \int_{-\infty}^{\infty} \left(\pi \left(\frac{\hat{\sigma}_\theta \theta}{\sigma_\theta} \right) \right) d\Phi \left(\frac{\theta}{\sigma_\theta} \right) = \lim_{\theta \rightarrow \infty} \frac{1}{2} (\pi(\theta) + \pi(-\theta))$, and therefore $\lim_{\hat{\sigma}_\theta \rightarrow \infty} |W(\pi, \hat{\sigma}_\theta)| = \lim_{\theta \rightarrow \infty} \frac{1}{2} |\pi(\theta) + \pi(-\theta)|$.

Derivation of equation 14:

Simple algebra shows that

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

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

⁴⁸It cannot be flat everywhere, because then informed demand would be completely inelastic, and there would be no way to absorb supply shocks.

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 $\lim_{x \rightarrow -\infty} H(\theta|x) = 1$ and $\lim_{x \rightarrow \infty} H(\theta|x) = 0$.

Proof. Notice that

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

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

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 $x \leq x_L(P)$, $d(x, P) = d_H$ for all $x \geq x_H(P)$, 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 $\lim_{x \rightarrow -\infty} \mathbb{E}(\pi(\theta)|x, P; D) < P < \lim_{x \rightarrow \infty} \mathbb{E}(\pi(\theta)|x, P; D)$ 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 $x \leq x_L(P)$, 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 $x \geq x_H(P)$.

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, $P = P_{\pi}(z) = \mathbb{E}(\pi(\theta)|x = z(P), P; \bar{D})$. 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:

Fix k and set $z_n = \theta_n + k$. Since $d(x, P(z_n)) \geq \bar{D}$ for $x \geq z_n$, we have $d_H - D(\theta_n, P(z_n)) \geq (d_H - \bar{D}) Pr(x < z_n | \theta_n) = F(k) > 0$ and $D(\theta_n, P(z_n)) - d_L \geq (\bar{D} - d_L)(1 - F(k)) > 0$. There therefore exists $\varepsilon > 0$ such that $\Psi(z_n | \theta_n) = 1 - G(D(\theta_n, P(z_n))) \in (\varepsilon, 1 - \varepsilon)$ for all θ_n . It now follows from

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

that $\lim_{z \rightarrow -\infty} m(\theta, z) = \lim_{z \rightarrow \infty} m(\theta, z) = \infty$, for fixed θ . Setting $z_n = \theta_n + k$, we obtain

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

Therefore, $\lim_{z \rightarrow \infty} H(z - u' | z) = 1$ for any finite u' implies that $\lim_{\theta_n \rightarrow \infty} m(\theta_n, z_n) = \infty$. Since $m(\theta) = \int m(\theta, z) \psi(z | \theta) dz$ and $\lim_{z \rightarrow \infty} m(\theta, z) = \infty$, we have $m(\theta_n) \geq m(\theta_n, z_n)(1 - \Psi(z_n | \theta_n)) \geq \varepsilon \cdot m(\theta_n, z_n)$ for θ_n sufficiently high, and therefore $\lim_{\theta_n \rightarrow \infty} m(\theta_n) = \lim_{\theta_n \rightarrow \infty} m(\theta_n, z_n) = \infty$.

Similarly, $\lim_{z \rightarrow -\infty} H(z - u' | z) = 0$ for any finite u' implies that $\lim_{\theta_n \rightarrow -\infty} m(\theta_n, z_n) = \infty$. Since $\lim_{z \rightarrow -\infty} m(\theta, z) = \infty$ it must be the case that $m(\theta_n) \geq m(\theta_n, z_n) \Psi(z_n | \theta_n) \geq \varepsilon \cdot m(\theta_n, z_n)$ for θ_n sufficiently low, and therefore $\lim_{\theta_n \rightarrow -\infty} m(\theta_n) = \infty$.

Proof of Proposition 5: See Online Appendix.

Proof of Theorem 2:

With symmetry, $\hat{H}_k(\theta) = \int_{-\infty}^{\theta} m_k(\theta') h(\theta') d\theta' = \int_{-\infty}^{\theta} m_k(-\theta') h(-\theta') d\theta' = \int_{\theta}^{\infty} m_k(\theta') h(\theta') d\theta' = 1 - \hat{H}_k(-\theta)$ for $k \in \{1, 2\}$. Moreover, by log-convexity of $m_1(\cdot)$ and $m_2(\cdot)/m_1(\cdot)$, $\hat{H}_2(\theta) \geq \hat{H}_1(\theta) \geq H(\theta)$ for $\theta \leq 0$ and $\hat{H}_2(\theta) \leq \hat{H}_1(\theta) \leq H(\theta)$ for $\theta \geq 0$. Write $W(\pi, m_k)$ as

$$W(\pi, m_k) = \int_{-\infty}^{\infty} (H(\theta) - \hat{H}_k(\theta)) d\pi(\theta) = \int_0^{\infty} (H(\theta) - \hat{H}_k(\theta)) d\hat{\pi}(\theta),$$

where $\hat{\pi}(\theta) = \pi(\theta) + \pi(-\theta)$. 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 $\hat{H}(\theta) \rightarrow \frac{1}{2}$. For given m , we write

$$W(\pi, m) = \int_0^\infty (H(\theta) - \hat{H}(\theta)) d\hat{\pi}(\theta) = \int_0^\infty (\hat{\pi}(\theta) - \hat{\pi}(0)) (m(\theta) - 1) h(\theta) d\theta.$$

Since m is symmetric and log-convex, there exists a unique $\bar{\theta} > 0$, such that $m(\theta) > 1$ if and only if $\theta > \bar{\theta}$ or $\theta < -\bar{\theta}$. Therefore

$$W(\pi, m) \geq \int_{\bar{\theta}}^\infty (\hat{\pi}(\theta) - \hat{\pi}(0)) (m(\theta) - 1) h(\theta) d\theta \geq (\hat{\pi}(\bar{\theta}) - \hat{\pi}(0)) (H(\bar{\theta}) - \hat{H}(\bar{\theta}))$$

when $\hat{\pi}(\theta)$ is increasing, and $W(\pi, m) \leq (\hat{\pi}(\bar{\theta}) - \hat{\pi}(0)) (H(\bar{\theta}) - \hat{H}(\bar{\theta}))$ when $\hat{\pi}(\theta)$ is decreasing. By construction, $\bar{\theta}$ is set to maximize $H(\theta) - \hat{H}(\theta)$, which implies that $\bar{\theta} \rightarrow \infty$ and $H(\bar{\theta}) - \hat{H}(\bar{\theta}) \rightarrow \frac{1}{2}$ as $\hat{H}(\cdot) \rightarrow \frac{1}{2}$. It follows that $\lim_{\hat{H}(\theta) \rightarrow \frac{1}{2}} |W(\pi, \hat{m})| = \lim_{\bar{\theta} \rightarrow \infty} \frac{1}{2} |\hat{\pi}(\bar{\theta}) - \hat{\pi}(0)| = \infty$ whenever $\lim_{\theta \rightarrow \infty} |\hat{\pi}(\theta)| = \infty$.