When Volatility Is Not Risk: Uncertainty, Disagreement, and Asset Prices*

Anand Bharadwaj†

AUGUST 7, 2026

Abstract

Stocks with high idiosyncratic volatility often earn low subsequent returns, even though investors may require compensation for greater payoff uncertainty when that risk is not fully diversifiable. The apparent contradiction arises because observed volatility can aggregate two distinct features of beliefs. Uncertainty within an investor's posterior weakens risk-bearing demand. Disagreement across posterior means can instead elevate prices when pessimistic investors cannot express their views through sufficiently negative positions. I first derive a common-prior projection bound linking uncertainty and disagreement for a common payoff and horizon. I then solve a rational expectations equilibrium in which the market-clearing price simultaneously reveals information and determines ownership under short-sale constraints. The model produces an ownership-dependent reversal in the price effect of fundamental volatility. I call the boundary between these regimes the *Grossman–Miller Frontier*. Empirically, firm-level IVOL loads positively on both 30-day option-implied uncertainty and analyst earnings forecast disagreement. Among larger, more liquid, and easier-to-borrow firms, IVOL loads more heavily on uncertainty, and the IVOL return coefficient moves toward the risk-compensation sign. As institutional ownership broadens, the IVOL coefficient rises mainly because the uncertainty contribution becomes more positive; under equal weights, the disagreement contribution also becomes less negative. Low returns following jointly high uncertainty and disagreement are strongest among smaller, less readily tradable firms, where pessimistic positions are harder to establish. The return content of volatility therefore depends on the beliefs it reflects and on which investors can act on them.

*I am indebted to Alan Auerbach, Yuriy Gorodnichenko, Martin Lettau, Emi Nakamura, Chris Shannon, and Johan Walden for their guidance on this paper. I am also especially grateful to Michael Woodford, Jón Steinsson, and Eben Lazarus for many suggestions that have greatly improved the paper. For valuable comments and feedback, I also thank Hassan Afrouzi, Robert Barro, Onur Bayar, Nicholas Bloom, Jamie Cross, David Frankel, Jonathon Hazell, Barry Hughes, Michael Jansson, Palani-Rajan Kadapakkam, Roman Kapuscinski, Chen Lian, Demian Pouzo, Andreas Schaab, Benjamin Schoefer, Kannan Sethuraman, Michael Smith, and Jiang Wang, as well as seminar participants at the University of California, Berkeley and the University of Melbourne.

†Department of Economics, University of California, Berkeley. Email: anand_bharadwaj@berkeley.edu.

1 Introduction

Stocks with high idiosyncratic volatility often earn low subsequent returns (Ang et al., 2006). At first glance, this empirical sign runs counter to the standard logic of risk compensation. When investors must bear more uncertainty about a firm's payoff and cannot fully diversify it, they require a lower price and a correspondingly higher expected return. This logic applies when volatility reflects uncertainty within investors' beliefs. Once the same statistic also reflects disagreement about the payoff's expected value, it combines belief components with opposing implications for demand and prices. This distinction provides a direct explanation for the idiosyncratic-volatility puzzle.

The relevant distinction is between dispersion within individual beliefs and dispersion across investors' posterior means. The former weakens risk-bearing demand and raises the compensation required to hold the stock. The latter has no independent effect on the clearing price when long and short positions are unrestricted, because optimistic purchases are offset by pessimistic sales. That symmetry fails once pessimistic investors cannot establish sufficiently negative positions. Ownership then tilts toward investors with relatively high valuations, elevating the current price and potentially depressing subsequent returns. Observed volatility can consequently embody either a risk-compensation effect or an overvaluation effect generated by optimistic selection, depending on the beliefs that produce it and the investors whose views enter market demand.

I study these forces in a rational expectations equilibrium in which the market-clearing price performs two functions. It reveals private information, albeit imperfectly, and it allocates a positive supply of the asset when short positions are constrained. Information revealed through price converts part of investors' private information into common information, compressing the dispersion of posterior means. The short-sale boundary then determines how the residual heterogeneity in beliefs is represented in demand. A mean-preserving increase in fundamental volatility raises both posterior uncertainty and disagreement: the former lowers price by reducing risk-bearing capacity, whereas the latter can raise price by shifting ownership toward relatively optimistic investors.

In a tractable location–scale economy, expected active ownership summarizes the balance between these two forces. Expected active ownership is the fraction of potential investors who hold a positive position. Fundamental volatility lowers price when more than one-half of potential investors hold positive positions and raises price when fewer than one-half do. The one-half ownership boundary is a volatility locus at which risk bearing and optimistic selection exactly offset. An economy with fixed physical float, in which

informed investors and rational hedgers trade for distinct reasons, shows how the same optimistic selection force operates when uncertainty, disagreement, price informativeness, and hedging demand move through separate equilibrium margins.

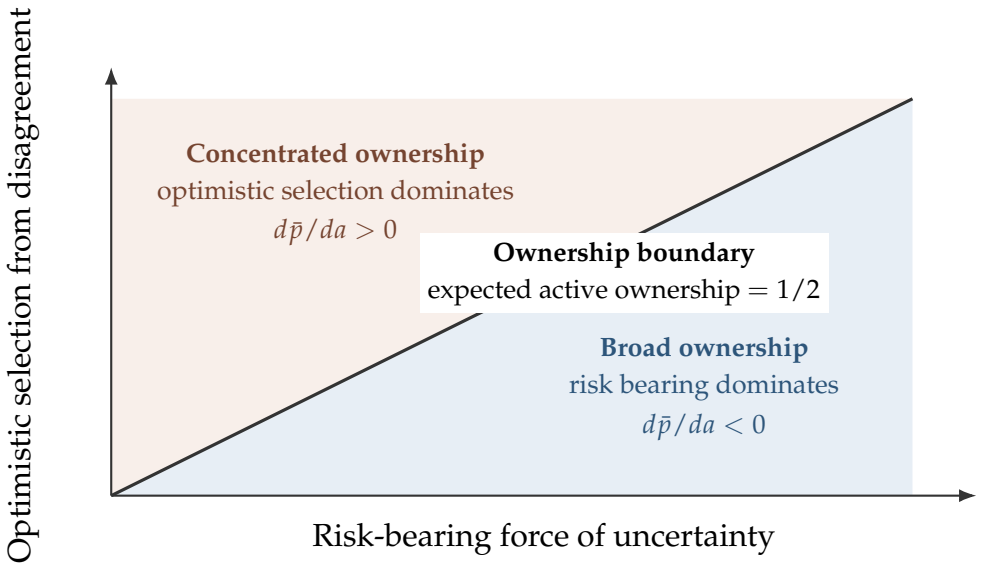


Figure 1: Schematic illustration of the two forces governing the price response to a mean-preserving increase in fundamental volatility. Uncertainty weakens risk-bearing demand, while disagreement strengthens selection toward optimistic owners when pessimists cannot short. Section 3.5 derives the exact ownership boundary in the location–scale economy.

Before either belief moment enters asset demand, their joint behavior is subject to a common-prior projection restriction. For a common payoff and forecast horizon, expected disagreement cannot become arbitrarily large relative to average posterior uncertainty. Define

$$B = \frac{D}{U}, \tag{1}$$

where D denotes dispersion across posterior means and U denotes average uncertainty within individual posteriors. The ratio measures how much of the uncertainty embedded in individual beliefs appears as disagreement across investors. In the Gaussian common/private-information benchmark, additional common precision reduces the ratio because common signals draw posterior means together more rapidly than they narrow each posterior distribution. More generally, the projection theorem bounds the ex ante ratio without imposing a particular signal structure. The restriction also limits how uncertainty and disagreement can move together. In a Gaussian one-factor benchmark, disagreement is proportional to squared uncertainty, so the two belief moments cannot be treated as

unrelated interventions. The economy with fixed physical float provides a broader setting in which learnable fundamental risk, residual payoff risk, price informativeness, and hedging demand can move more independently.

The empirical analysis uses U.S. equities from February 1996 through December 2021. The measure of idiosyncratic volatility is the `idiovol` characteristic constructed by Gu et al. (2020), henceforth IVOL. Uncertainty is annualized 30-day firm-level SVIX^2, constructed from OptionMetrics following Martin (2017) and Martin and Wagner (2019). I evaluate the fitted option surface at a common 30-day maturity and integrate out-of-the-money put prices below the forward price and call prices above it across strikes. The measure therefore uses the fitted 30-day strike slice rather than a single implied-volatility quote. I convert total firm variance into a within-month percentile rank and standardize that rank. Because SVIX^2 is measured under the risk-neutral distribution, it can reflect both expected physical variation and the pricing of variance risk. Disagreement is the dispersion of I/B/E/S forecasts of annual earnings per share. The principal return and continuous market structure regressions use standardized log disagreement. The component-loading table, ownership breadth quintiles, and conditional sign maps use its standardized within-month rank, placing uncertainty and disagreement on the same monotone scale. Subsequent returns are drawn from CRSP. Survey data permit a closer alignment of payoff and horizon for the common-prior ratio.

I use *investability* to summarize how readily investors can establish and maintain a position in a stock. The measure averages within-month ranks of market capitalization, liquidity, and ease of borrowing, then ranks and standardizes the composite; higher values identify larger, more liquid, and easier-to-borrow stocks. Liquidity combines the favorable orientations of the bid–ask spread, Amihud illiquidity, and the frequency of zero trading. A separate measure of weak trading capacity combines small size, a wide bid–ask spread, high Amihud illiquidity, and frequent zero trading. It rises when pessimistic investors face greater difficulty trading, while option-implied borrowing cost is examined separately.

The decomposition of idiosyncratic volatility reveals substantial exposure to both belief components. In the linked option–analyst sample, IVOL loads positively on 30-day firm SVIX variance and analyst earnings forecast disagreement, with respective loadings of 0.692 and 0.141. Together, the two measures explain 62.6 percent of average monthly cross-sectional IVOL variation. Their relative importance changes markedly across the market. Among large, liquid, and readily traded firms, IVOL is closely associated with option-implied uncertainty. Among small and illiquid firms, disagreement accounts for a considerably larger share. As investability rises, IVOL loads more heavily on uncertainty

and less heavily on disagreement. Weak trading capacity and higher implied borrowing costs reverse both movements. Identical values of IVOL can therefore represent materially different configurations of beliefs.

The price implications of IVOL therefore depend both on which belief component it reflects and on the market structure through which those beliefs enter demand. In the linked sample, greater disagreement predicts lower next-month returns after option-implied uncertainty is held fixed, and the interaction between the two belief measures is negative. The relations become more pronounced among small firms, stocks in which pessimistic positions are difficult to establish, and firms with high implied borrowing costs. A continuous interaction specification yields the same ordering. As investability rises, the uncertainty and IVOL coefficients increase and the disagreement coefficient becomes less negative under equal weights. In that cross section, the return shortfall associated with joint high uncertainty and disagreement deepens as trading capacity weakens. Reweighting the linked option–analyst sample toward the size–liquidity distribution of the broader stock sample strengthens the disagreement coefficient and preserves the negative relation in the joint high-uncertainty, high-disagreement state.

The uncertainty premium is most visible among historical S&P 500 constituents, the most liquid stocks, and the most investable firms. These are also the groups in which IVOL is most closely related to firm SVIX variance. A specification patterned on Martin and Wagner (2019) closely reproduces their published S&P 500 estimate and remains positive in the extended sample. Across mutually exclusive investability deciles, the uncertainty and IVOL coefficients rise steadily from the least to the most investable stocks.

I measure institutional ownership breadth as the fraction of reporting S12 institutions that hold a stock, following Chen et al. (2002). From the lowest breadth quintile to the highest, the IVOL coefficient rises by 0.235 percentage points under equal weights and by 0.389 under capped value weights. An exact monthly decomposition assigns most of this movement to a more positive uncertainty contribution. Under equal weights, the disagreement contribution also becomes less negative. This ownership measure reinforces the evidence from size, liquidity, borrowing costs, and composite investability. Across these measures, the IVOL relation moves toward the risk-compensation sign as market participation and risk-bearing capacity expand.

The survey evidence helps distinguish among information structures behind these belief moments. For fitted SPF belief distributions, disagreement relative to uncertainty remains below one in nearly every quarter and moves toward the statewise unit ceiling of the homoskedastic noisy-information benchmark during volatile periods. This dominant

pattern, together with modest dispersion in posterior widths, supports noisy information with stochastic fundamental volatility and finite heterogeneity in signal precision. The isolated post-COVID exceedances can arise within the same heterogeneous-precision noisy-information family, especially its sticky zero/infinite-precision limit after large shocks; heterogeneous priors can generate similar tail behavior. These rare observations do not overturn the noisy-information baseline.

Together, the results provide a conditional account of the IVOL puzzle. The common-prior projection theorem restricts the joint behavior of uncertainty within beliefs and disagreement across beliefs without committing to a specific signal structure. The equilibrium then combines information aggregation with constrained ownership to derive an ownership-dependent sign for the price response to fundamental volatility. Empirically, IVOL contains belief components with distinct return implications, and their relative importance varies systematically with firm size, liquidity, borrowing conditions, investability, and institutional participation. The familiar negative IVOL coefficient is an average across firms for which the same volatility statistic represents different economic forces.

1.1 Relation to the Literature

The theoretical framework brings together two traditions that have generally emphasized different roles for prices. In noisy rational expectations models, prices reveal information but remain imperfect signals because supply or trading-demand shocks prevent complete revelation, as in Grossman (1976), Grossman and Stiglitz (1980), and Hellwig (1980). Under unconstrained CARA-normal demand, optimistic and pessimistic positions offset: the clearing price reflects the average posterior belief, compensation for posterior uncertainty, and net supply, while disagreement affects the allocation of positions without exerting an independent valuation effect.

Miller-style models identify the complementary force that arises when pessimistic beliefs cannot be expressed symmetrically. Short-sale constraints or resale opportunities place greater weight on optimistic valuations, as in Miller (1977), Harrison and Kreps (1978), and Scheinkman and Xiong (2003). Related dynamic models trace the evolution of heterogeneous beliefs through news, learning, trading, and price information, including Hong and Stein (1999), Banerjee (2011), and Daniel et al. (2023). The equilibrium developed here combines these mechanisms within one market-clearing price. Information aggregation determines how much private disagreement survives, while constrained ownership determines how the surviving dispersion affects valuation.

The common-prior result is also connected to macroeconomic models of expectations under imperfect information. Noisy-information and limited-information-processing models examine how incomplete information shapes expectations and aggregate dynamics, as in Woodford (2003) and Sims (2003). Sticky-information models instead emphasize infrequent updating (Mankiw and Reis, 2002). Empirical work uses forecast errors and cross-sectional disagreement to distinguish these mechanisms (Coibion and Gorodnichenko, 2015). The projection theorem supplies a restriction that is common to these information structures, while the noisy and sticky benchmarks explain why disagreement can remain modest in ordinary periods yet become locally important after large shocks.

The measurement problem has a close antecedent in research separating forecast disagreement from subjective or aggregate uncertainty. Patton and Timmermann (2010) examine the sources of forecast heterogeneity, and Lahiri and Sheng (2010) show that disagreement is one component of aggregate forecast uncertainty rather than a generally valid substitute for it. The analysis below sharpens this distinction by deriving a tight common-prior relation between expected disagreement and posterior uncertainty and then allowing the two belief moments to have opposing effects on asset prices.

The pricing model also builds on rational expectations economies with constrained trading. Diamond and Verrecchia (1987) show that short-sale restrictions suppress informative trades and slow the incorporation of adverse information without implying a mechanical upward price bias in every environment. Related models study private information under trading constraints (Bai et al., 2006). In the economy with fixed float developed here, the same price aggregates private information, remains imperfectly revealing because of hedging demand, and determines which informed investors ultimately hold the asset.

The empirical analysis belongs to the literature on firm-level volatility and expected returns. Campbell et al. (2001) document the importance and changing magnitude of firm-level idiosyncratic volatility, and Ang et al. (2006) establish its negative relation with subsequent returns. Subsequent explanations emphasize expected volatility, reversal, lottery demand, costly arbitrage, and asymmetric limits to arbitrage. These mechanisms appear in Fu (2009), Huang et al. (2010), Bali et al. (2011), Pontiff (2006), and Stambaugh et al. (2015). Option-based research provides forward-looking measures of uncertainty and expected returns, including Martin (2017) and Martin and Wagner (2019), while recent work extracts common disagreement from several proxies and relates it to IVOL and returns (Goulding et al., 2026). The distinctive empirical claim here is that the return content of IVOL varies with the relative importance of uncertainty and disagreement and with the trading environment that determines whose beliefs enter prices.

A complementary literature shows that disagreement becomes more consequential when pessimistic trading is costly or unavailable. Evidence from securities-lending markets, ownership breadth, and short-sale experiments supports this mechanism, as shown by Jones and Lamont (2002), Nagel (2005), Cohen et al. (2007), and Kaplan et al. (2013). The empirical tests examine whether disagreement predicts low returns after uncertainty is held fixed and whether the relation intensifies as the ability to establish pessimistic positions deteriorates.

The distinction between disagreement and uncertainty also has a long history in survey research. Dispersion in point forecasts need not coincide with an individual forecaster's uncertainty, as shown by Zarnowitz and Lambros (1987) and Patton and Timmermann (2010). The survey evidence uses the U.S. Survey of Professional Forecasters (SPF) and the Livingston Survey, both administered by the Federal Reserve Bank of Philadelphia, as well as the CFO Survey conducted jointly by Duke University and the Federal Reserve Banks of Richmond and Atlanta. The Livingston Survey and CFO Survey measures build on the forward-return-expectations framework of Gandhi et al. (2026). Analyst earnings forecast dispersion is widely used as a measure of disagreement (Diether et al., 2002), whereas option-implied variance is a market price under the risk-neutral distribution. The common-prior projection bound applies directly only when the moments refer to the same payoff, horizon, and probability measure. Accordingly, the stock-level tests focus on pricing signs and cross-sectional movements, while the survey exercises examine the ratio between disagreement and uncertainty more directly.

Section 2 develops the common-prior restriction on uncertainty and disagreement. Section 3 introduces prices and derives the ownership frontier. Section 4 translates the model into empirical predictions, and Section 5 presents the stock-level evidence across investability, trading constraints, and institutional participation. Section 6 considers interpretation and robustness. The appendix contains the full derivation for the economy with fixed float, survey evidence, dynamic-information benchmarks, and detailed variable construction.

2 Uncertainty and Disagreement

Before introducing asset demand, I separate the two belief moments that observed volatility can conceal. Let X_t denote a payoff-relevant state and $\mathcal{I}_{it}$ investor i's information set. Average posterior uncertainty is

$$U_t = \mathbb{E}_i[\mathrm{Var}(X_t \mid \mathcal{I}_{it})]\,, \tag{2}$$

whereas disagreement is the cross-sectional variance of posterior means conditional on the realized state,

$$D_t = \mathrm{Var}_i(\mathbb{E}[X_t \mid \mathcal{I}_{it}] \mid X_t)\,. \tag{3}$$

Posterior uncertainty records the dispersion of outcomes within an investor's belief, whereas disagreement records dispersion across investors' posterior means. Conditioning on the realized state excludes variation attributable merely to movements in the underlying fundamental. A structural comparison between D_t and U_t therefore requires both moments to refer to the same payoff, forecast horizon, and probability measure.

2.1 Two Dimensions of Belief Dispersion

A law-of-total-variance decomposition makes the distinction explicit. If Y_t is a representative forecaster's information, then

$$\begin{aligned}\mathrm{Var}(X_t) = &\underbrace{\mathbb{E}[\mathrm{Var}(X_t \mid Y_t)]}_{\text{uncertainty}} + \underbrace{\mathbb{E}[\mathrm{Var}(\mathbb{E}[X_t \mid Y_t] \mid X_t)]}_{\text{expected disagreement}} \\ &+ \underbrace{\mathrm{Var}(\mathbb{E}[\mathbb{E}[X_t \mid Y_t] \mid X_t])}_{\text{forecast variability}}.\end{aligned} \tag{4}$$

The second term isolates heterogeneity across forecasters after the realized state has been held fixed. By contrast, the unconditional variance of posterior means would also absorb ordinary variation in the fundamental and would therefore overstate disagreement.

In the canonical noisy-information benchmark,

$$X_t \sim N(0, \sigma_\epsilon^2), \qquad Y_{it} = X_t + \nu_{it}, \qquad \nu_{it} \sim N(0, \sigma_\nu^2).$$

Gaussian updating gives

$$U^N = \frac{\sigma_\epsilon^2 \sigma_\nu^2}{\sigma_\epsilon^2 + \sigma_\nu^2}, \qquad D^N(x) = \frac{\sigma_\epsilon^4 \sigma_\nu^2}{(\sigma_\epsilon^2 + \sigma_\nu^2)^2}, \tag{5}$$

and hence

$$B^N(x) := \frac{D^N(x)}{U^N} = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_\nu^2} \leq 1 \qquad \text{for every } x. \tag{6}$$

The ratio compares the dispersion of posterior means with the uncertainty that remains within each posterior. It increases with fundamental variance and declines with signal noise. Because posterior variances are homoskedastic in this Gaussian benchmark, disagreement remains below uncertainty in every realized state.

The sticky benchmark is an extreme heterogeneous-precision noisy-information model. A fraction $1 - \lambda$ receives a perfectly precise current signal, while a fraction λ receives no current signal and remains at the prior. Equivalently, current-signal precision is infinite for updaters and zero for nonupdaters. Updated investors have zero posterior variance and stale investors retain variance σ_ϵ^2, so

$$U^S = \lambda \sigma_\epsilon^2, \tag{7}$$

while the dispersion between current and stale forecasts is

$$D^S(x) = \lambda(1 - \lambda)x^2. \tag{8}$$

Thus

$$B^S(x) := \frac{D^S(x)}{U^S} = (1 - \lambda)\frac{x^2}{\sigma_\epsilon^2}, \qquad B^S(x) > 1 \iff x^2 > \frac{\sigma_\epsilon^2}{1 - \lambda}, \tag{9}$$

even though

$$\frac{\mathbb{E}[D^S(X)]}{U^S} = 1 - \lambda \leq 1. \tag{10}$$

Although the ratio remains bounded after averaging over states, sufficiently large shocks can generate statewise exceedances. Updated and stale investors attach different posterior means to the same realized payoff. This does not violate the common-prior projection restriction developed below, which constrains an ex ante ratio averaged over payoff states rather than every realized-state ratio. The zero/infinite-precision case is not unique in this respect: heterogeneous finite signal precisions can also produce statewise exceedances

after sufficiently large realizations while preserving the ex ante projection bound.[3] Heterogeneous priors can generate similar tail behavior after an extreme realization.

The two environments also differ in disagreement about posterior widths. Define

$$U_{it} := \mathrm{Var}(X_t \mid \mathcal{I}_{it}), \qquad G_t := \mathrm{Var}_i(\sqrt{U_{it}}), \qquad \bar{U}_t := \mathbb{E}_i[U_{it}]. \tag{11}$$

The homoskedastic noisy-information model has $G_t = 0$. In the static sticky benchmark,

$$\frac{G_t}{\bar{U}_t} = 1 - \lambda = \frac{\mathbb{E}[D^S(X)]}{U^S}. \tag{12}$$

The sticky-information literature has primarily emphasized staggered updating. For the comparison here, the primitive distinction is cross-sectional heterogeneity in posterior precision. A richer noisy-information model generates $G_t > 0$ through heterogeneous finite signal precisions. Sticky updating is the zero/infinite-precision limit of the same family: information vintages are the histories generated by successive zero-precision draws between perfectly informative updates. The joint behavior of D_t/U_t and G_t/U_t therefore helps locate the data along this precision-heterogeneity spectrum. Appendix E develops the dynamic benchmarks and shows why homoskedastic noisy information with stochastic fundamental volatility approaches the canonical unit ceiling in volatile states, while stronger precision heterogeneity, including the sticky limit, can generate rare exceedances after large shocks.

2.2 A Common-Prior Projection Bound

The Gaussian examples point to a restriction that does not depend on a particular signal structure. Consider the investor–state product space $(\Omega \times \mathcal{I}, \mathcal{F} \otimes \mathcal{B}, \mathbb{P} \otimes \mu)$, with a payoff $X(\omega, i) = X(\omega)$ that is common across investors. The sigma-algebra $\mathcal{H}$ records the identity and information of a randomly sampled investor, so $m = \mathbb{E}[X \mid \mathcal{H}]$ is the corresponding posterior mean. A second sigma-algebra $\mathcal{G}$ contains the realized payoff and any common conditioning variables, but not investor-specific information. The conditional variation of m given $\mathcal{G}$ then coincides with expected cross-investor disagreement.

[3] If $k_i \in [0,1]$ is investor i's posterior loading on a private Gaussian signal and prior variance is V, then, conditional on $X = x$, $D(x) = x^2 \mathrm{Var}_i(k_i) + V\mathbb{E}_i[k_i(1-k_i)]$ and $U = V(1 - \mathbb{E}_i k_i)$. Thus $\mathrm{Var}_i(k_i) > 0$ permits $D(x)/U > 1$ for sufficiently large $|x|$, while $\mathbb{E}[D(X)]/U = \mathbb{E}_i k_i \leq 1$.

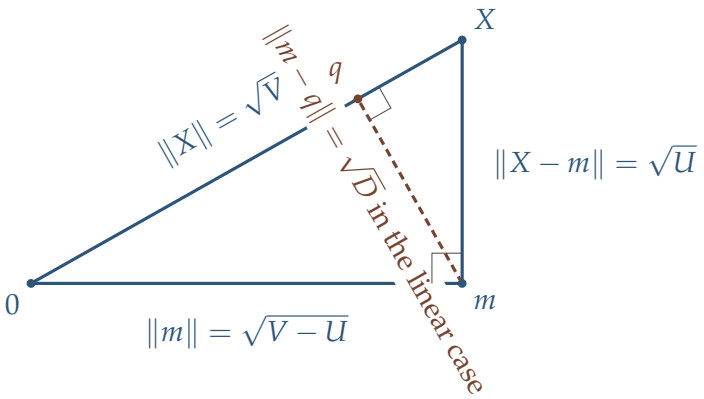


Figure 2: Projection geometry underlying the common-prior projection bound. Lengths represent standard deviations and squared lengths represent variances. In the linear-projection case, the altitude corresponds to expected disagreement. Enlarging the common information set can only reduce the residual projection error.

Theorem 1 (Common-prior projection bound)**.** *Let $X \in L^2$ with $V := \mathrm{Var}(X) > 0$, let $\mathcal{H}$ be the information sigma-algebra of a randomly sampled investor, and let $\mathcal{G}$ satisfy $\sigma(X) \subseteq \mathcal{G}$. Define*

$$U := \mathbb{E}[\mathrm{Var}(X \mid \mathcal{H})], \qquad D := \mathbb{E}[\mathrm{Var}(\mathbb{E}[X \mid \mathcal{H}] \mid \mathcal{G})].$$

If investors share the common probability measure and form rational conditional expectations, then

$$D \leq U\left(1 - \frac{U}{V}\right) \leq U. \tag{13}$$

Equivalently, whenever $U > 0$,

$$\frac{D}{U} \leq 1 - \frac{U}{V}. \tag{14}$$

After centering X, equality in the first inequality holds if and only if $P_{L^2(\mathcal{G})}m = P_{\mathrm{span}\{X\}}m$. In particular, equality holds when $\mathcal{G} = \sigma(X)$ and $\mathbb{E}[m \mid X] = ((V - U)/V)X$. Thus, when $\mathcal{G} = \sigma(X)$, joint Gaussianity of (m, X) assures equality because the conditional expectation is linear.

Proof sketch. Center X, write $m = P_{\mathcal{H}}X$ and $g = P_{\mathcal{G}}m$, and use orthogonality of conditional expectations. Then $U = \|X - m\|^2 = V - \|m\|^2$ and $\mathbb{E}[mX] = \|m\|^2 = V - U$. Because $\mathrm{span}\{X\} \subseteq L^2(\mathcal{G})$, projecting m onto the richer space cannot leave a larger residual than projecting it onto the line spanned by X. With $q = ((V - U)/V)X$,

$$D = \|m - g\|^2 \leq \|m - q\|^2 = (V - U) - \frac{(V - U)^2}{V} = U\left(1 - \frac{U}{V}\right).$$

Appendix D gives the complete Hilbert-space proof and the conditional formulation. □

The theorem provides a sharper restriction than the familiar inequality $D \leq U$. The static linear Gaussian benchmark in equations (5)–(6) attains the first inequality exactly. More generally, Gaussianity assures equality when the realized payoff is the common conditioning variable; the inequality can be strict when $\mathcal{G}$ contains additional common information that predicts posterior means. As information resolves a larger share of prior variance, the greatest feasible ex ante ratio D/U necessarily declines. Linear Gaussian noisy-information models, including heterogeneous-precision specifications and the sticky zero/infinite-precision limit, attain the projection boundary after integrating over states. Their conditional implications depend on precision dispersion. Homoskedastic noisy information keeps the statewise ratio $B^N(x)$ below one, whereas cross-sectional heterogeneity in posterior loadings can produce temporary exceedances after sufficiently large shocks.

2.3 Empirical Interpretation of the Bound

For beliefs concerning the same payoff and horizon, define

$$B_t := \frac{D_t}{U_t}. \tag{15}$$

Low values indicate that posterior means remain tightly clustered relative to the uncertainty faced by each investor. High values indicate that cross-investor disagreement is quantitatively large. The ratio describes the relation between two belief moments before either enters a pricing equation. Its valuation consequences depend additionally on the information revealed through price and on the extent to which pessimistic investors can take negative positions.

Proposition 1 (Conditional common-prior projection bound)**.** *Let $\mathcal{C}_t$ be common date-t information and apply Theorem 1 under the regular conditional probability given $\mathcal{C}_t$. Define*

$$V_t := \mathrm{Var}(X \mid \mathcal{C}_t), \qquad U_t := \mathbb{E}[\mathrm{Var}(X \mid \mathcal{H} \vee \mathcal{C}_t) \mid \mathcal{C}_t],$$

and

$$D_t := \mathbb{E}[\mathrm{Var}(\mathbb{E}[X \mid \mathcal{H} \vee \mathcal{C}_t] \mid \mathcal{G} \vee \mathcal{C}_t)|\mathcal{C}_t]\,,$$

where $\sigma(X) \vee \mathcal{C}_t \subseteq \mathcal{G} \vee \mathcal{C}_t$. On $\{V_t > 0, U_t > 0\}$,

$$\frac{D_t}{U_t} \leq 1 - \frac{U_t}{V_t} \leq 1 \qquad \textit{almost surely.}$$

The proposition is an ex ante population restriction, formed by averaging over future payoff states conditional on date-t information. It does not require every realized-state ratio to remain below one. Appendix C applies this distinction to point and density forecasts from the U.S. Survey of Professional Forecasters, then compares survey disagreement with market uncertainty using the Livingston Survey and the Duke University/Federal Reserve CFO Survey.

3 Information, Ownership, and Prices

The common-prior projection bound characterizes the belief configurations that can arise, but it does not determine how those beliefs are capitalized into prices. Valuation depends on two additional features of the market: the information investors extract from the market-clearing price and the extent to which pessimistic beliefs can be expressed through short positions. I isolate these mechanisms in turn before allowing the same price to aggregate information and allocate ownership.

3.1 Unconstrained Trading

There are dates 0 and 1, a risk-free asset with gross return one, and a risky payoff

$$X \sim N(\mu_0, a), \qquad s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, b), \tag{16}$$

where signal errors are independent across a unit mass of investors. When investors do not extract information from the current price, Gaussian updating gives

$$U^N = \mathrm{Var}(X \mid s_i) = \frac{ab}{a+b}, \qquad m_i^N = \mathbb{E}[X \mid s_i] = \frac{b\mu_0 + as_i}{a+b}, \tag{17}$$

$$\bar{m}^N(X) := \mathbb{E}_i[m_i^N \mid X] = \frac{b\mu_0 + aX}{a+b}, \qquad D^N := \mathrm{Var}_i(m_i^N \mid X) = \frac{a^2 b}{(a+b)^2}. \tag{18}$$

Thus

$$\frac{D^N}{U^N} = \frac{a}{a+b}, \qquad \frac{\sqrt{D^N}}{U^N} = \frac{1}{\sqrt{b}}. \tag{19}$$

Investors have constant absolute risk aversion (CARA) utility with coefficient $\gamma > 0$. If investor i has posterior $N(m_i, U)$ and the asset price is p, unconstrained demand is

$$q_i^* = \frac{m_i - p}{\gamma U}. \tag{20}$$

Proposition 2 (Unconstrained disagreement)**.** *With fixed positive float S and unrestricted positions, the clearing price is*

$$p^{N,0}(X) = \bar{m}^N(X) - \gamma U^N S. \tag{21}$$

Holding the cross-sectional mean posterior fixed,

$$\frac{\partial p^{N,0}}{\partial U} = -\gamma S < 0, \qquad \frac{\partial p^{N,0}}{\partial D} = 0. \tag{22}$$

Under unrestricted trading, optimistic purchases are matched by pessimistic short sales. An increase in posterior uncertainty reduces the risk-bearing demand of every investor and therefore lowers the clearing price, while a mean-preserving widening of posterior means merely redistributes long and short positions. Disagreement acquires an independent valuation effect only when this symmetry in positions is broken. Appendix A provides the derivation.

3.2 Learning from Prices

A noisy rational expectations equilibrium (REE) allows investors to learn from price. Let net supply pressure be

$$Z = \bar{S} + u, \qquad u \sim N(0, \Delta^2), \tag{23}$$

independent of fundamentals and private signals. In the unique linear equilibrium, observing price is equivalent to observing the public signal

$$Y_P = X - \gamma b u, \qquad \tau_p^H = \frac{1}{\gamma^2 b^2 \Delta^2}, \tag{24}$$

where τ_p^H is the precision of price about X. Writing

$$C_H = \left(a^{-1} + b^{-1} + \tau_p^H\right)^{-1}, \tag{25}$$

posterior uncertainty and state-conditioned disagreement are

$$U_H = C_H, \qquad D_H = \frac{C_H^2}{b}, \qquad \frac{D_H}{U_H} = \frac{C_H}{b}. \tag{26}$$

The full price coefficients and the equilibrium derivation appear in Appendix A.

The price signal augments the information shared by all investors. It narrows each posterior distribution and, at the same time, draws posterior means closer together. More generally, if τ_c denotes common precision and τ_x private precision, then

$$U = \frac{1}{\tau_c + \tau_x}, \qquad D = \frac{\tau_x}{(\tau_c + \tau_x)^2}, \qquad \frac{D}{U} = \frac{\tau_x}{\tau_c + \tau_x}. \tag{27}$$

An informative price therefore reduces both belief moments, with a proportionately larger contraction in disagreement because common information does not itself create cross-sectional dispersion. Yet, so long as short positions remain unrestricted, the average posterior continues to determine valuation. Information aggregation alone cannot confer greater influence on optimistic beliefs than on pessimistic ones.

3.3 Short-Sale Constraints

Return to the no-price-learning beliefs in equations (17)–(18), retain fixed positive float S, and impose $q_i \geq 0$. Demand becomes

$$q_i = \frac{(m_i - p)^+}{\gamma U}. \tag{28}$$

For later use, define

$$\phi(d) = \frac{1}{\sqrt{2\pi}} e^{-d^2/2}, \qquad \Phi(d) = \int_{-\infty}^{d} \phi(z)\, dz, \qquad h(d) = \phi(d) + d\Phi(d). \tag{29}$$

For $Z \sim N(0,1)$, $h(d) = \mathbb{E}[(d+Z)^+] = \phi(d) + d\Phi(d)$, the expected positive part of $d + Z$, and $h'(d) = \Phi(d)$.

Conditional on the common state, posterior means are $N(\bar{m}, D)$. With

$$d = \frac{\bar{m} - p}{\sqrt{D}}, \tag{30}$$

market clearing is

$$\sqrt{D}\, h(d) = \gamma U S. \tag{31}$$

The active ownership mass is $\Phi(d)$.

Lemma 1 (Derivatives of constrained demand)**.** *Let* $A(p, D) = \mathbb{E}[(m_i - p)^+]$. *Then*

$$\frac{\partial A}{\partial p} = -\Phi(d), \qquad \frac{\partial A}{\partial D} = \frac{\phi(d)}{2\sqrt{D}}. \tag{32}$$

Proposition 3 (Uncertainty and disagreement under constrained ownership)**.** *At a no-short clearing price with $U, D, S > 0$,*

$$\left.\frac{\partial p}{\partial U}\right|_{D,\bar{m},S} = -\frac{\gamma S}{\Phi(d)} < 0, \qquad \left.\frac{\partial p}{\partial D}\right|_{U,\bar{m},S} = \frac{\phi(d)}{2\sqrt{D}\,\Phi(d)} > 0. \tag{33}$$

Proof. Differentiate (31), using (32). □

The short-sale boundary breaks the symmetry between the two comparative statics. Greater uncertainty weakens the demand of every investor who remains an owner, whereas greater disagreement spreads valuations around the ownership cutoff. Because the pessimistic tail cannot enter with negative positions while the optimistic tail remains free to buy, the distribution of owners shifts toward higher valuations. The average valuation among owners is

$$\mathbb{E}[m_i \mid m_i > p] = \bar{m} + \sqrt{D}\,\frac{\phi(d)}{\Phi(d)}. \tag{34}$$

The term

$$\lambda(d) := \frac{\phi(d)}{\Phi(d)}$$

is the inverse Mills ratio. It is the optimism wedge created by the ownership cutoff: it measures how far the average owner's valuation lies above the average valuation among all potential investors. As active ownership $\Phi(d)$ contracts, the inverse Mills ratio rises and optimistic selection becomes stronger.

The Gaussian information structure also permits a clean experiment in primitive fundamental volatility. Equation (19) implies $\sqrt{D^N}/U^N = 1/\sqrt{b}$, so clearing fixes a cutoff d_N satisfying

$$h(d_N) = \gamma S\sqrt{b}. \tag{35}$$

The cutoff is independent of fundamental variance a, and

$$\mathbb{E}[p^{N,1}] = \mu_0 - d_N\sqrt{D^N} = \mu_0 - \frac{d_N}{\sqrt{b}}U^N. \tag{36}$$

Proposition 4 (Ownership threshold without price learning)**.** *Holding (b, S, γ, μ_0) fixed, a mean-preserving increase in fundamental variance a raises the expected constrained newswatcher price if and only if*

$$d_N < 0 \iff \Phi(d_N) < \frac{1}{2} \iff S < \frac{\phi(0)}{\gamma\sqrt{b}}. \tag{37}$$

It lowers expected price when the inequalities are reversed.

Proof. Since $dU^N/da = b^2/(a+b)^2 > 0$, differentiating (36) gives a sign opposite to d_N. Because h is strictly increasing and $h(0) = \phi(0)$, equation (35) gives the float equivalence. Symmetry of the Normal distribution gives the ownership equivalence. □

This benchmark isolates the ownership-dependent sign reversal in a setting where posterior beliefs do not respond to the market price. In this respect, investors are analogous to the newswatchers of Hong and Stein (1999): they form beliefs from private news rather than extracting information from the current price. The next model retains the same selection logic in a full rational expectations equilibrium, where price reveals information and clears constrained demand simultaneously.

Finite shorting capacity has an exact representation. If $q_i \geq -\ell$, define

$$\widetilde{q}_i = q_i + \ell, \qquad \widetilde{p} = p - \gamma U \ell. \tag{38}$$

Then

$$\widetilde{q}_i = \frac{(m_i - \widetilde{p})^+}{\gamma U}, \qquad \int_0^1 \widetilde{q}_i \, di = S + \ell. \tag{39}$$

Conditional on beliefs, relaxing the short floor is equivalent to increasing effective float. It admits additional pessimistic positions and therefore lowers the clearing price:

$$\frac{\partial p}{\partial \ell} = \gamma U \left(1 - \frac{1}{\Phi(\widetilde{d})}\right) < 0. \tag{40}$$

3.4 Informative Prices with Constrained Demand

Combining price learning with long-only demand creates a technical tension. Physical supply must remain positive, while the equilibrium price must retain an invertible Gaussian structure if investors are to extract information from it. Conventional Gaussian supply noise can take negative values, whereas truncation destroys linearity. I resolve this problem by mapping a latent Gaussian financing shock into positive float through the same positive-part function that arises when long-only demand is aggregated. The construction preserves an informative price and permits market clearing in every state.

Let

$$X \sim N(\mu_0, \tau_0^{-1}), \qquad s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, \tau_x^{-1}), \tag{41}$$

and let $z \sim N(0, \sigma_z^2)$ be independent. Investors observe (s_i, p), but neither z nor realized supply separately, and choose $q_i \geq 0$. Fix $\kappa > 0$ and $d_0 \in \mathbb{R}$. Physical supply is the strictly positive random variable

$$S(z) = \frac{\sqrt{\tau_x}}{\gamma} h(d_0 - \kappa\sqrt{\tau_x}\, z) > 0. \tag{42}$$

Physical supply is specified directly by equation (42). This functional form guarantees positive supply while preserving a linear price signal.

Define

$$Y = X + \kappa z, \qquad \tau_p = \frac{1}{\kappa^2 \sigma_z^2}, \qquad \tau = \tau_0 + \tau_x + \tau_p. \tag{43}$$

Proposition 5 (Full transformed supply rational expectations equilibrium)**.** *The economy has a rational expectations equilibrium with*

$$U = \mathrm{Var}(X \mid s_i, p) = \frac{1}{\tau}, \qquad D = \mathrm{Var}_i(\mathbb{E}[X \mid s_i, p] \mid X, z) = \frac{\tau_x}{\tau^2}, \tag{44}$$

$$\beta_p = \frac{\tau_x + \tau_p}{\tau}, \qquad \alpha = \frac{\tau_0 \mu_0}{\tau} - d_0\sqrt{D}, \tag{45}$$

and price

$$p(X, z) = \alpha + \beta_p (X + \kappa z). \tag{46}$$

Investor i's posterior mean and demand are

$$m_i = U(\tau_0 \mu_0 + \tau_x s_i + \tau_p Y), \qquad q_i = \frac{(m_i - p)^+}{\gamma U}. \tag{47}$$

The state-contingent ownership cutoff is

$$d(z) = \frac{\mathbb{E}_i[m_i \mid X, z] - p}{\sqrt{D}} = d_0 - \kappa\sqrt{\tau_x}\, z. \tag{48}$$

Consequently, active ownership is $\Phi(d(z))$ and aggregate demand equals $S(z)$ in every state.

Proof. Price is strictly increasing in Y, so observing p is equivalent to observing the Gaussian public signal Y. Bayesian updating gives (44) and (47). Conditional on (X, z), posterior means have cross-sectional mean

$$\bar{m} = U[\tau_0 \mu_0 + (\tau_x + \tau_p) X + \tau_p \kappa z]$$

and variance D. Substituting the price coefficients gives $\bar{m} - p = \sqrt{D}[d_0 - \kappa\sqrt{\tau_x}z]$. Aggregating Gaussian long-only demand then yields

$$\int_0^1 q_i \, di = \frac{\sqrt{D}}{\gamma U} h(d(z)) = \frac{\sqrt{\tau_x}}{\gamma} h(d(z)) = S(z).$$

Hence a single equilibrium price both determines posterior beliefs and clears constrained demand. □

The residual disagreement–uncertainty ratio is

$$\frac{D}{U} = \frac{\tau_x}{\tau_0 + \tau_x + \tau_p}. \tag{49}$$

Without price learning, the corresponding ratio is $\tau_x/(\tau_0 + \tau_x)$. In the noisy rational expectations tradition of Grossman (1976) and Hellwig (1980), observing price adds a common signal and compresses disagreement relative to uncertainty:

Corollary 1 (Price learning attenuates disagreement relative to uncertainty)**.**

$$\frac{D/U}{\tau_x/(\tau_0 + \tau_x)} = \frac{\tau_0 + \tau_x}{\tau_0 + \tau_x + \tau_p} < 1 \quad \textit{when } \tau_p > 0. \tag{50}$$

Thus an informative equilibrium price lowers the disagreement–uncertainty ratio by the factor $(\tau_0 + \tau_x)/(\tau_0 + \tau_x + \tau_p)$.

The equilibrium also shows explicitly how market learning changes measured beliefs. Let

$$U^S = \frac{1}{\tau_0 + \tau_x}, \qquad D^S = \frac{\tau_x}{(\tau_0 + \tau_x)^2}, \qquad \rho_p = \frac{\tau_0 + \tau_x}{\tau_0 + \tau_x + \tau_p}. \tag{51}$$

Then

$$U^P = \rho_p U^S, \qquad D^P = \rho_p^2 D^S, \qquad \frac{D^P}{U^P} = \rho_p \frac{D^S}{U^S}. \tag{52}$$

Price information scales uncertainty by ρ_p and disagreement by ρ_p^2. Consequently, the disagreement–uncertainty ratio contracts by the same factor ρ_p, formalizing the sense in which a common price signal compresses posterior means more strongly than posterior widths.

3.5 Balancing Risk Bearing and Optimistic Selection

The two valuation forces balance along what I term the *Grossman–Miller Frontier*. The name combines the price-learning mechanism of Grossman (1976) with the ownership-selection mechanism of Miller (1977). Grossman's information aggregation determines how much private disagreement survives in posterior beliefs. Miller's optimistic selection determines whether that surviving dispersion changes ownership enough to offset the decline in risk-bearing demand.

Let $a = \tau_0^{-1}$ denote prior fundamental variance and hold $(\tau_x, \kappa, \sigma_z, d_0, \gamma)$ fixed. This mean-preserving volatility experiment leaves the physical supply schedule and absolute price-signal precision unchanged while allowing posterior beliefs and the market-clearing price to adjust. Define

$$\chi^2 = \kappa^2 \tau_x \sigma_z^2 = \frac{\tau_x}{\tau_p}. \tag{53}$$

Expected active ownership and mean supply are

$$\Omega := \mathbb{E}_z[\Phi(d(z))] = \Phi\left(\frac{d_0}{\sqrt{1+\chi^2}}\right), \tag{54}$$

$$\bar{S}(d_0) := \mathbb{E}_z[S(z)] = \frac{\sqrt{\tau_x}}{\gamma}\sqrt{1+\chi^2}\, h\left(\frac{d_0}{\sqrt{1+\chi^2}}\right). \tag{55}$$

Expected price is

$$\bar{p}(a) = \mathbb{E}[p(X,z)] = \mu_0 - d_0\sqrt{D} = \mu_0 - d_0\sqrt{\tau_x}\, U. \tag{56}$$

Proposition 6 (Ownership frontier in the full equilibrium)**.** *In the transformed-supply equilibrium,*

$$\frac{dU}{da} = \frac{U^2}{a^2} > 0, \qquad \frac{dD}{da} = \frac{2\tau_x U^3}{a^2} > 0, \tag{57}$$

$$\frac{d\bar{p}}{da} = -\frac{d_0\sqrt{\tau_x}U^2}{a^2}. \tag{58}$$

Therefore fundamental volatility raises expected price if and only if

$$d_0 < 0 \iff \Omega < \frac{1}{2} \iff \bar{S} < \bar{S}^*, \tag{59}$$

where

$$\bar{S}^* = \frac{\sqrt{\tau_x}}{\gamma}\sqrt{1+\frac{\tau_x}{\tau_p}}\,\phi(0). \tag{60}$$

It lowers expected price when the inequalities are reversed.

Proof. The derivatives follow from $U = (a^{-1} + \tau_x + \tau_p)^{-1}$, $D = \tau_x U^2$, and (56). Equation (54) gives $\Omega < 1/2$ exactly when $d_0 < 0$. Since h is strictly increasing, (55) gives the equivalent mean-float threshold. □

The comparative static holds price precision, the supply schedule, mean supply, and expected active ownership fixed while an increase in a raises both U and D. Under these conditions, active ownership is sufficient to determine which price effect dominates. When $\Omega < 1/2$, the greater influence of optimistic owners more than offsets the contraction in risk-bearing demand, so expected price rises. When $\Omega > 1/2$, the uncertainty channel dominates and expected price falls.

The exact one-half threshold is a consequence of the location–scale geometry rather than of Gaussianity itself.

Proposition 7 (Location–scale ownership switch)**.** *Suppose posterior valuations satisfy*

$$m_i = \mu + cUZ_i, \qquad c > 0, \tag{61}$$

where Z_i has a fixed continuous distribution with a finite first moment and a unique median at zero. Hold (μ, c, S) fixed and let $q_i = (m_i - p)^+/(\gamma U)$. An increase in U raises the clearing price if and only if fewer than one-half of investors hold positive positions.

Proof. Let $H_Z(d) = \mathbb{E}[(d + Z_i)^+]$ and $d = (\mu - p)/(cU)$. Clearing is $cH_Z(d) = \gamma S$, which fixes $d = d^*$ independently of U. Hence $p(U) = \mu - cUd^*$ and $dp/dU = -cd^*$. The zero-median normalization makes $d^* < 0$ equivalent to active ownership below one-half. □

Disclosure precision. The ownership frontier also governs the value of common information. Suppose an issuer commits, before a report is realized, to a public signal of precision $q \in [0, \bar{q}]$, on top of existing common precision τ_c and private precision τ_x. In the constrained Gaussian benchmark,

$$U(q) = \frac{1}{\tau_c + \tau_x + q}, \qquad\qquad D(q) = \frac{\tau_x}{(\tau_c + \tau_x + q)^2},$$

$$\mathbb{E}[p(q)] = \mu - \frac{d\sqrt{\tau_x}}{\tau_c + \tau_x + q}, \qquad \frac{d\mathbb{E}[p(q)]}{dq} = \frac{d\sqrt{\tau_x}}{(\tau_c + \tau_x + q)^2}, \tag{62}$$

where d is fixed independently of q by $\sqrt{\tau_x}h(d) = \gamma S$. Public precision raises expected value when ownership is broad, $d > 0$, and lowers it when ownership is concentrated, $d < 0$. The same frontier that governs the price effect of exogenous volatility therefore governs an issuer's incentive to improve common information. Appendix B derives the optimal precision under convex disclosure costs and develops the ex ante reporting interpretation.

The threshold follows because the experiment fixes the median valuation while scaling the dispersion of valuations in proportion to uncertainty. A more general volatility shock may also shift the average posterior, alter price informativeness or hedging demand, and change the mapping between D and U. Each of these responses enters the total price derivative. Equation (60) illustrates the comparative statics: more precise information in price contracts the region in which ownership selection dominates, whereas greater private precision expands it.

3.6 Fixed Float and Hedging Demand

The closed-form equilibrium relies on a supply schedule that preserves a linear price signal. Appendix F develops a complementary economy with fixed physical float, in which informed investors receive private signals about the payoff and a separate sector trades to hedge nontraded exposure. The hedging motive prevents the price from fully revealing the fundamental.

The extension with fixed physical float allows belief components that move together in the tractable equilibrium to vary separately. Learnable fundamental variance raises both posterior uncertainty and disagreement, while residual payoff variance raises uncertainty without necessarily widening disagreement. It also demonstrates that the short-sale result is not an artifact of the transformed supply schedule. On a stable separating branch, Proposition 9 shows that a mean-preserving spread in informed investors' desired positions weakly raises the clearing price, with a strict increase whenever the spread raises expected demand after the short floor is imposed. Greater physical float or greater shorting capacity has the opposite effect.

The economy with fixed physical float preserves the central ownership mechanism. Once pessimistic positions are truncated, a wider distribution of valuations places greater weight on optimistic demand. The location–scale boundary remains exact for the exper-

iment in Proposition 7. In this economy, a primitive volatility shock can also move the average posterior, price informativeness, hedging demand, and demand elasticity, so the price response reflects the full set of equilibrium margins.

3.7 Return Implications

Across the sequence of models, disagreement reallocates positions without affecting the clearing price when shorting is unrestricted. Information revealed through price reduces both uncertainty and disagreement, with a proportionately larger contraction in the latter. Once pessimistic positions are constrained, however, a wider distribution of valuations can raise price by shifting ownership toward optimists. The price response to a primitive volatility shock consequently reflects the balance between this optimistic selection effect and the accompanying reduction in risk-bearing demand.

The frontier expresses this trade-off as an ownership condition. In the tractable experiment, expected active ownership below one-half identifies the region in which a rise in fundamental volatility elevates price and depresses expected returns. With broader ownership, the risk-bearing effect of uncertainty dominates. Empirically, firm size, liquidity, borrowing conditions, composite investability, and institutional participation provide complementary views of the trading environment in which this trade-off operates.

Purely idiosyncratic risk need not command a premium in a frictionless, fully diversified factor model. The stock-level risk-bearing channel is relevant when incomplete diversification, investor clienteles, or limits on intermediation make firm risk costly to absorb, which is the environment represented by the single-stock rational expectations model.

3.8 Empirical Counterparts

Survey forecasts and market prices reveal different aspects of beliefs. Survey densities and point forecasts record expectations without weighting them by market-clearing positions. Analyst forecasts already incorporate prices and public information, yet they remain reports rather than position-weighted beliefs. Dispersion in earnings forecasts is therefore best interpreted as a price-aware measure of disagreement about cash flows.

Option-implied variance is forward-looking and market-based, but it is defined under the risk-neutral distribution. Total firm SVIX variance serves as the stock-level proxy for uncertainty. The main analysis studies its total return content. Conditioning on physical-

volatility measures would instead identify a residual option-variance component and answer a different question.

When posterior moments are aligned by payoff and horizon, $B = D/U$ measures disagreement relative to uncertainty. Firm-level earnings dispersion and option-implied equity variance do not form such a structural ratio because they refer to different payoffs and probability measures. I therefore use them to study pricing signs and cross-sectional movements, while the survey analysis examines more closely aligned measures of uncertainty and disagreement.

4 Empirical Predictions

The empirical predictions concern two related margins: what idiosyncratic volatility contains and how its components are priced. Uncertainty and disagreement should carry different return coefficients, and their relative importance should vary with the trading environment. Firm SVIX variance and earnings forecast dispersion enter separately because one is a risk-neutral measure of equity variance and the other records dispersion in reported cash-flow forecasts. The common-prior ratio is examined most directly using within-survey point and density forecasts, where payoff, horizon, and probability measure are aligned.

Two measures summarize the trading environment. Investability averages the within-month ranks of size, liquidity, and ease of borrowing, with higher values denoting stocks that are larger, more liquid, and easier to borrow. The separate index of weak trading capacity combines small size, a wide bid–ask spread, high Amihud illiquidity, and frequent zero trading; higher values indicate greater difficulty expressing pessimistic views. Implied borrowing cost provides a direct complementary measure of shorting conditions.

4.1 Belief Composition and Investability

Holding disagreement fixed, greater uncertainty weakens risk-bearing demand and should be followed by a higher return. Holding uncertainty fixed, greater disagreement can instead raise the current price when pessimistic investors cannot establish sufficiently negative positions, implying a lower subsequent return. The same distinction should appear inside IVOL. Among firms that are easy to hold and trade, option-implied uncertainty should account for more of the characteristic and disagreement for less.

Hypothesis 1 (Investability and the IVOL relation)**.** *As investability rises, IVOL loads more heavily on uncertainty and less heavily on disagreement. The uncertainty and IVOL return coefficients rise, while the disagreement coefficient becomes less negative.*

4.2 Constraints on Pessimistic Trading

The valuation effect of disagreement depends on the limited influence of pessimistic beliefs. Its return coefficient should become more negative among small, illiquid, expensive-to-borrow, or otherwise difficult-to-short stocks. The ownership mechanism motivates examining whether the joint high-uncertainty, high-disagreement state is most adverse in these same parts of the market. In such states, a broad distribution of valuations meets a binding short-sale boundary, giving relatively optimistic owners greater influence on the clearing price.

Hypothesis 2 (Constraints on pessimistic trading)**.** *The disagreement coefficient becomes more negative as the ability to establish pessimistic positions deteriorates. The return shortfall associated with the joint high-uncertainty, high-disagreement state is concentrated in the same markets. The uncertainty and disagreement gradients in Hypothesis 1 reverse as trading capacity weakens or borrowing becomes more costly.*

Firm size, liquidity, trading activity, and implied borrowing costs provide complementary views of this margin. The tests use both a composite measure of trading capacity and the option-implied borrowing rate.

4.3 Ownership and Market Participation

The ownership frontier concerns the total price effect of fundamental volatility. In the location–scale economy, this effect changes sign as active ownership crosses one-half. Institutional ownership breadth provides an empirical measure of how widely a stock is held. As breadth increases, the valuations of a small optimistic subset should exert less influence on the clearing price. The IVOL coefficient should therefore move toward the positive sign associated with risk compensation, even if breadth and investability do not change the composition of IVOL in identical ways.

Hypothesis 3 (Ownership and participation)**.** *The idiosyncratic-volatility coefficient rises with institutional ownership breadth. Across breadth groups, the uncertainty contribution becomes more positive, while the disagreement contribution becomes less negative.*

4.4 Survey Implications

For beliefs concerning the same payoff and horizon, $B = D/U$ measures disagreement relative to uncertainty, and the common-prior theorem bounds this ratio ex ante. Homoskedastic Gaussian noisy information keeps the ratio below one in every state, while stochastic fundamental volatility moves it toward that unit ceiling in volatile periods. Heterogeneous finite signal precision generates modest dispersion in posterior widths. Sticky updating is the zero/infinite-precision limit of this family: some forecasters receive a perfectly precise current signal while others receive no new signal. The resulting precision heterogeneity can produce statewise exceedances after large shocks. As common information about a fixed event accumulates, disagreement should contract more rapidly than average posterior uncertainty.

Hypothesis 4 (Information structure)**.** *Within-survey measures for a common payoff and horizon satisfy the common-prior restriction on average. In the canonical homoskedastic noisy-information benchmark, the ratio remains below one state by state and moves toward one in volatile periods. Occasional exceedances after unusually large shocks are consistent with heterogeneous signal precision, especially the sticky zero/infinite-precision limit, or with heterogeneous priors.*

5 Empirical Results

The empirical analysis begins by asking what economic content is embedded in idiosyncratic volatility and then examines how that content is priced across different trading environments. I first decompose the characteristic into exposures to option-implied uncertainty and analyst disagreement. I then estimate their return coefficients at a common 30-day horizon and trace how they vary with trading conditions, composite investability, and institutional participation. Complementary within-stock evidence describes the extreme time-series tail, while survey data permit a more direct comparison of uncertainty and disagreement for a common payoff and horizon.

5.1 Data and Timing

The stock-level sample extends from February 1996 through December 2021. Monthly returns, prices, shares outstanding, and identifiers are obtained from CRSP. Firm characteristics come from the public data associated with Gu et al. (2020). The sole measure of idiosyncratic volatility is their `idiovol` characteristic. The baseline, continuous market

structure, and conditional sign regressions use standardized, winsorized log IVOL. The ownership breadth quintile regressions instead use its standardized within-month rank.

Uncertainty is annualized 30-day firm-level $SVIX^2$, constructed from OptionMetrics option surfaces following Martin (2017) and Martin and Wagner (2019). On the final valid surface date in month t, I evaluate the OptionMetrics fitted surface at a common 30-day maturity, infer the forward price from put–call parity, and numerically integrate out-of-the-money puts below the forward and calls above it across strikes. The construction therefore incorporates the fitted 30-day strike slice rather than a single implied-volatility quote. I annualize the resulting total firm variance, convert it into a within-month percentile rank, and standardize the rank. As an option-implied risk-neutral variance measure, $SVIX^2$ can embed both expected physical variation and pricing-kernel or variance-risk-premium effects. The baseline analysis retains the total measure.

Disagreement is the dispersion of current I/B/E/S analyst forecasts of annual earnings per share. Forecasts are linked to CRSP, screened for staleness, and scaled by price. The principal pricing and continuous market structure regressions log, winsorize, and standardize the measure within month. The component-loading table, ownership breadth quintiles, and conditional sign maps instead use its standardized within-month percentile rank. The measure captures differences across analysts about corporate earnings. The resulting option–analyst sample contains 595,181 stock-months over 311 months. All signals are observed by the end of month t and are used to predict the total stock return in month $t + 1$. Unless otherwise indicated, coefficients are averages of monthly Fama–MacBeth regressions and inference is based on six-lag Newey–West standard errors.

The baseline return regression relates subsequent returns to firm SVIX variance and forecast disagreement jointly. A companion specification adds standard return predictors. Capped value weights truncate market capitalization at the monthly 99th percentile before normalization. Portfolio breakpoints and weights are reconstructed within each monthly estimation sample. Tables state explicitly when signal ranks are also recomputed within a restricted sample.

5.2 IVOL, Uncertainty, and Disagreement

Table 1 reports the benchmark for the broader stock sample. Across these stocks, the return difference under equal weights between the highest and lowest IVOL deciles is -1.215 percentage points per month ($t = -1.96$). The spread is most negative among the smallest stocks and within the constrained tercile. Among the largest and least constrained stocks, it

is close to zero or positive under capped value weights. The unconditional IVOL relation is materially weaker in the linked option–analyst data, where uncertainty and disagreement can be observed jointly.

Table 1: The IVOL return relation in the broader and linked samples

Sample	FM IVOL slope	EW D10–D1	Capped VW D10–D1	Months
Broader stock sample	-0.249 (-1.29)	-1.215 (-1.96)	-0.720 (-1.15)	311
Linked option–analyst sample	-0.121 (-0.57)	-0.443 (-0.75)	-0.216 (-0.36)	311
Smallest size quintile	-0.316 (-1.57)	**-1.525 (-2.29)**	**-2.065 (-3.12)**	311
Largest size quintile	-0.083 (-0.38)	-0.213 (-0.39)	0.117 (0.20)	311
Weakest shorting ability	-0.145 (-0.70)	-0.885 (-1.39)	**-1.681 (-2.60)**	311
Strongest shorting ability	-0.046 (-0.25)	-0.181 (-0.36)	0.083 (0.14)	311

The first result column is the no-control monthly Fama–MacBeth slope on standardized log GKX `idiovol`. The next two columns are high-minus-low IVOL decile returns in percentage points per month. Parentheses contain six-lag Newey–West t-statistics. The broader sample contains 1,443,616 stock-months. The linked option–analyst sample contains 595,181. Capped value weights truncate market capitalization at its monthly 99th percentile before normalization.

To establish whether both belief components are embedded in the characteristic, I project IVOL each month on 30-day firm SVIX variance and earnings forecast disagreement. Table 2 reports an uncertainty loading of 0.692 ($t = 35.42$) and a disagreement loading of 0.141 ($t = 6.60$) in the linked sample. Together, the two measures account for 62.6 percent of average monthly cross-sectional IVOL variation, and both loadings remain positive under capped value weights. These loading regressions establish that the observed characteristic contains substantial exposure to each belief component before their return implications are examined.

Table 2: Uncertainty and disagreement as components of idiosyncratic volatility

	Equal weight			Capped value weight		
Sample	U loading	D loading	R^2	U loading	D loading	R^2
Entire linked sample	**0.692 (35.42)**	**0.141 (6.60)**	0.626	**0.739 (65.38)**	**0.090 (4.98)**	0.591
Historical S&P 500	**0.643 (38.71)**	**0.138 (7.05)**	0.531	**0.646 (62.01)**	**0.088 (5.15)**	0.532
Smallest size quintile	**0.415 (18.87)**	**0.214 (7.23)**	0.346	**0.426 (18.09)**	**0.208 (6.88)**	0.338
Largest size quintile	**0.668 (48.50)**	**0.094 (5.22)**	0.521	**0.655 (63.41)**	**0.074 (4.32)**	0.541
Least liquid quintile	**0.460 (18.05)**	**0.250 (8.93)**	0.431	**0.476 (16.41)**	**0.238 (9.84)**	0.401
Most liquid quintile	**0.645 (32.98)**	**0.090 (4.62)**	0.493	**0.684 (52.58)**	**0.076 (4.03)**	0.576
Bottom investability quintile	**0.439 (18.30)**	**0.230 (8.01)**	0.388	**0.456 (16.41)**	**0.240 (8.02)**	0.384
Top investability quintile	**0.702 (50.57)**	**0.064 (3.71)**	0.549	**0.678 (61.71)**	**0.059 (3.50)**	0.564

Each row reports monthly cross-sectional regressions of standardized log GKX `idiovol` on the standardized within-sample ranks of total 30-day firm SVIX variance and EPS disagreement. Ranks are recomputed among the stocks in each indicated sample. Coefficients and average monthly R^2 values are averaged across 311 months. Parentheses contain six-lag Newey–West t-statistics. The regressions contain no controls. The table describes IVOL's belief composition; return pricing is examined separately.

The subgroup estimates reveal a systematic shift in the composition of IVOL across trading environments. Table 3 estimates this shift directly using a composite of size,

liquidity, and ease of borrowing. A one-standard-deviation increase in this investability measure raises the IVOL loading on uncertainty by 0.039 ($t = 6.05$) and lowers the loading on disagreement by 0.043 ($t = -6.20$). Under capped value weights, the estimates are 0.058 and -0.047. When trading capacity weakens or implied borrowing costs rise, both movements reverse. IVOL becomes less closely associated with uncertainty and more closely associated with disagreement.

Table 3: The composition of IVOL across market structure

	Equal weight		Capped value weight	
Market characteristic M	$U \times M$	$D \times M$	$U \times M$	$D \times M$
Composite investability	**0.039 (6.05)**	**-0.043 (-6.20)**	**0.058 (4.42)**	**-0.047 (-6.40)**
Trading-capacity constraint	**-0.034 (-5.25)**	**0.034 (5.14)**	-0.014 (-1.08)	**0.037 (4.44)**
Implied borrowing cost	**-0.031 (-7.51)**	**0.045 (7.94)**	**-0.061 (-6.50)**	**0.040 (8.60)**

Each row reports the change in the uncertainty and disagreement loadings associated with a one-standard-deviation increase in M. IVOL is standardized winsorized log `idiovol`, U is the standardized within-month rank of total 30-day firm SVIX variance, and D is standardized log EPS disagreement. Higher investability denotes easier participation; higher values of the other measures denote greater difficulty in establishing pessimistic positions. Parentheses contain six-lag Newey–West t-statistics. Appendix Table 14 reports the full specifications.

Figure 3 traces this rotation across investability deciles. Under equal weights, the uncertainty loading rises from 0.607 in the first decile to 0.730 in the tenth, while the disagreement loading falls from 0.202 to 0.067. Capped value weights produce the same ordering. IVOL's composition changes throughout the cross section, with no dependence on a single size or liquidity cutoff.

5.3 Short-Sale Constraints and Disagreement

Table 4 reports the principal common-horizon return regressions. In the linked sample, the disagreement coefficient conditional on uncertainty is -0.181 percentage points per month ($t = -1.77$), while the uncertainty coefficient conditional on disagreement is -0.055. The interaction between the two measures is more precisely estimated at -0.173 ($t = -3.52$). In the broad cross section, returns are especially low when uncertainty and disagreement are simultaneously high. The positive uncertainty premium emerges in the large and investable segments in which IVOL is most closely connected to firm SVIX variance.

The full-sample average conceals substantial heterogeneity. Conditional disagreement is -0.428 ($t = -3.45$) in the smallest size quintile, -0.363 ($t = -2.71$) among stocks with the weakest capacity for pessimistic trading, and -0.276 ($t = -2.43$) in the highest-

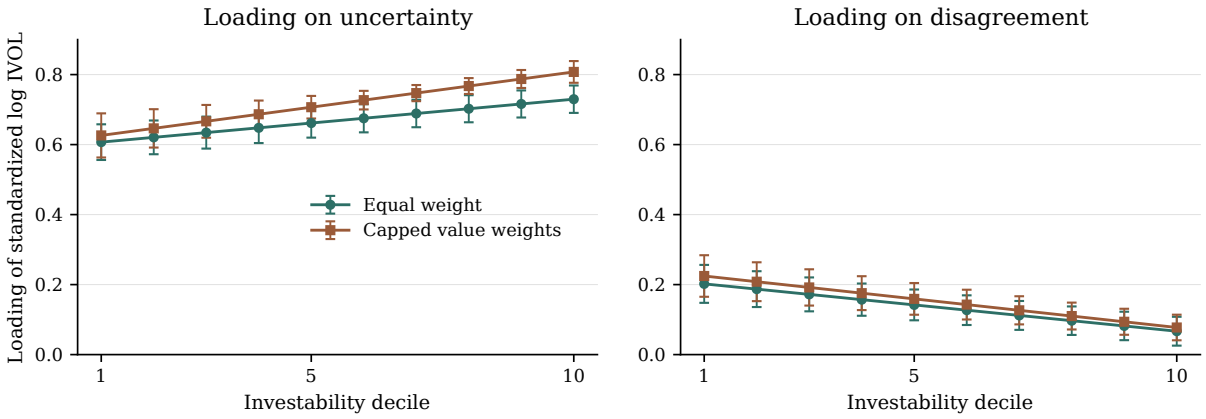


Figure 3: The composition of idiosyncratic volatility across investability. The curves report implied uncertainty and disagreement loadings from monthly interaction regressions evaluated across investability deciles. The two panels use a common vertical scale. Error bars denote 95 percent Newey–West confidence intervals.

borrow-cost quartile. In the corresponding large, readily traded, and low-borrow-cost groups, the coefficient is close to zero. The uncertainty–disagreement interaction follows the same ordering and remains negative when total firm SVIX variance is replaced by a market-adjusted 30-day measure.

Standard return predictors leave the central pattern intact. The conditional disagreement coefficient becomes -0.185 ($t = -2.24$), and the uncertainty–disagreement interaction is -0.144 ($t = -3.45$). Panel B of Table 4 shows how the estimates change as conventional controls and the market adjustment are introduced. The market-adjusted measure subtracts index variance at the same horizon, scaled by the stock's squared beta, from firm SVIX variance. The slightly more negative interaction shows that common market variance does not drive the joint return pattern. Appendix Table 12 reports the full sequence under both weighting schemes.

Figure 4 presents the joint state without imposing a linear functional form. Stocks are sorted first into disagreement quintiles and then into uncertainty quintiles within each disagreement group. Under equal weights, the average return in the high-uncertainty, high-disagreement portfolio is 0.67 percent, compared with 1.18 percent in the high-uncertainty, low-disagreement portfolio. Capped value weights preserve the same ordering. The surface gives the negative interaction a clear economic interpretation. The return associated with uncertainty varies materially with the prevailing level of disagreement.

A continuous interaction specification produces the same gradients. Table 5 shows that a one-standard-deviation increase in investability raises the uncertainty coefficient by

Table 4: Uncertainty and disagreement in next-month returns

Panel A: Baseline estimates and cross-sectional heterogeneity

Sample	U only	D only	$U \mid D$	$D \mid U$	$U \times D$
Entire linked sample	-0.138 (-0.76)	-0.165 (-1.26)	-0.055 (-0.32)	-0.181 (-1.77)	**-0.173 (-3.52)**
Entire linked sample, capped VW	0.037 (0.20)	-0.042 (-0.33)	0.093 (0.51)	-0.112 (-1.03)	**-0.147 (-3.04)**
Smallest size quintile	**-0.562 (-2.17)**	**-0.478 (-3.20)**	-0.264 (-1.14)	**-0.428 (-3.45)**	**-0.267 (-3.31)**
Largest size quintile	0.027 (0.13)	-0.061 (-0.46)	0.069 (0.33)	-0.104 (-0.98)	-0.156 (-1.94)
Weakest shorting-ability quintile	-0.357 (-1.37)	**-0.368 (-2.38)**	-0.085 (-0.37)	**-0.363 (-2.71)**	**-0.259 (-2.78)**
Strongest shorting-ability quintile	0.073 (0.34)	-0.007 (-0.06)	0.107 (0.52)	-0.051 (-0.47)	-0.105 (-1.16)
Highest-borrow-cost quartile	-0.354 (-1.76)	**-0.320 (-2.31)**	-0.183 (-1.00)	**-0.276 (-2.43)**	**-0.239 (-3.79)**
Lowest-borrow-cost quartile	0.028 (0.17)	0.010 (0.08)	0.049 (0.30)	-0.064 (-0.51)	-0.043 (-0.68)

Panel B: Selected control sequence

Uncertainty measure and terms	Controls	U	D	$U \times D$
Total 30-day U and D	None	-0.055 (-0.32)	-0.181 (-1.77)	–
Total 30-day U and D	Non-volatility	-0.144 (-0.86)	**-0.185 (-2.24)**	–
Total 30-day U, D, and interaction	Non-volatility	-0.169 (-1.00)	-0.145 (-1.82)	**-0.144 (-3.45)**
Market-adjusted 30-day U, D, and interaction	Non-volatility	-0.143 (-1.24)	-0.142 (-1.49)	**-0.164 (-4.37)**

Monthly Fama–MacBeth coefficients in percentage points per month, with six-lag Newey–West t-statistics in parentheses. The outcome is the next-month total return. U is the standardized within-month rank of total 30-day firm SVIX variance and D is standardized log EPS disagreement. Panel A reports univariate and joint regressions without controls. Signal transformations are formed across the linked sample each month before subgroup selection; capped market-cap weights are normalized within each indicated sample. Panel B reports equal-weight estimates. The non-volatility controls are size, book-to-market, momentum, reversal, profitability, investment, and log analyst coverage. IVOL, realized volatility, MAX, and skewness are excluded. Appendix Table 12 reports the full sequence under both weighting schemes. The linked sample contains 595,181 stock-months over 311 months.

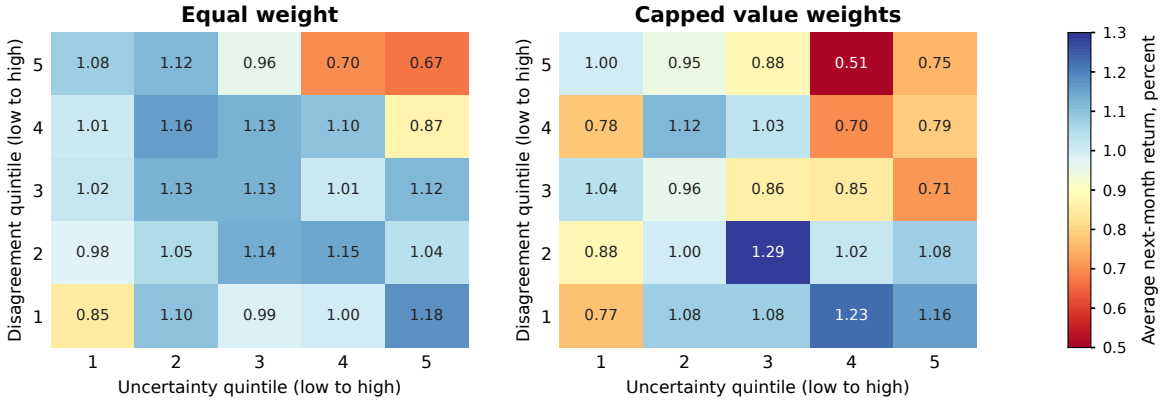


Figure 4: Average next-month returns across uncertainty and disagreement quintiles. The sort uses total 30-day firm SVIX variance and earnings forecast disagreement in the linked sample from February 1996 through December 2021. Within each month, stocks are sorted first on disagreement and then on uncertainty. The two panels use a common color scale.

0.108 ($t = 2.73$) and makes the disagreement coefficient 0.103 less negative ($t = 2.17$). The IVOL coefficient rises by 0.153 ($t = 4.17$). The uncertainty and disagreement movements reverse as pessimistic trading becomes harder or borrowing costs rise. Under capped value weights, the positive uncertainty and IVOL interactions survive, whereas the disagreement movement is concentrated in the cross section under equal weights.

Table 5: Market structure and the pricing of uncertainty, disagreement, and IVOL

Panel A: Continuous pricing gradients

	Equal weight			Capped value weight		
Market characteristic M	$U \times M$	$D \times M$	$IVOL \times M$	$U \times M$	$D \times M$	$IVOL \times M$
Composite investability	**0.108 (2.73)**	**0.103 (2.17)**	**0.153 (4.17)**	**0.131 (2.52)**	-0.048 (-0.82)	**0.131 (3.01)**
Raw S12 breadth	**0.240 (2.94)**	0.064 (0.98)	0.080 (1.15)	0.064 (0.70)	0.044 (0.71)	-0.033 (-0.46)
Size-neutral S12 breadth	**0.114 (2.33)**	0.057 (1.78)	**0.075 (2.09)**	0.068 (1.35)	-0.005 (-0.15)	0.031 (0.63)
Trading-capacity constraint	**-0.087 (-2.08)**	**-0.100 (-2.15)**	–	**-0.138 (-2.38)**	0.050 (0.71)	–
Implied borrowing cost	**-0.090 (-3.40)**	**-0.094 (-2.20)**	–	**-0.091 (-2.14)**	0.036 (0.85)	–

Panel B: Trading capacity and the joint high-uncertainty, high-disagreement state

Controls	Weighting	$U \times D$	$D \times C$	$U \times D \times C$
None	Equal weight	**-0.139 (-3.01)**	0.011 (0.24)	**-0.060 (-1.98)**
Non-volatility	Equal weight	**-0.117 (-2.65)**	0.002 (0.06)	**-0.074 (-2.58)**
Non-volatility	Capped value weight	**-0.125 (-2.29)**	0.140 (1.81)	-0.022 (-0.56)

In both panels, U is the standardized within-month rank of total 30-day firm SVIX variance and D is standardized log EPS disagreement; IVOL is standardized winsorized log `idiovol`. In Panel A, each entry is the change in a return coefficient associated with a one-standard-deviation increase in M. Higher investability and breadth denote broader participation. Higher values of the capacity constraint and borrowing cost denote greater difficulty in expressing pessimistic views. Raw-breadth regressions also include size and the corresponding signal-by-size interactions. Panel B defines C so that it rises as size falls and bid–ask spreads, Amihud illiquidity, and zero trading increase. The triple coefficient summarizes whether the negative joint high-U, high-D return relation becomes more pronounced as pessimistic trading capacity deteriorates. Parentheses contain six-lag Newey–West t-statistics. Appendix Tables 13 and 16 report the full specifications.

The interaction specification asks whether the low returns associated with jointly high uncertainty and disagreement become more pronounced as stocks become smaller, less liquid, and less frequently traded. The triple interaction is -0.060 ($t = -1.98$) under equal weights and -0.074 ($t = -2.58$) after standard return controls are added. The corresponding estimates under capped value weights are close to zero. Under equal weights, the return shortfall is therefore deepest where constraints on pessimistic trading are most severe. Panel B of Table 5 places these estimates beside the continuous pricing gradients. Appendix Table 16 reports the full specification.

Within this specification, the constraint gradient is concentrated in high-uncertainty states rather than in the disagreement slope at average uncertainty. Under capped value weights the effect is muted, consistent with its concentration among smaller firms in which pessimistic constraints are more likely to bind.

5.4 Readily Tradable Stocks

The risk-bearing channel is more readily identified in the panel design of Martin and Wagner (2019), who show that firm-level SVIX variance predicts expected stock returns. Table 6 applies their raw 30-day firm SVIX variance specification without variance controls. Over their historical S&P 500 sample from 1996 to 2014, the estimated gamma is 0.321, close to the published estimate of 0.301. It rises to 0.361 over 1996–2021 and to 0.648 ($t = 2.08$) during 2012–2021 after disagreement is introduced. The smaller coefficient in the entire linked sample is consistent with the decomposition results: the risk-bearing content of IVOL is most visible among established and investable firms.

Table 6: The 30-day uncertainty premium across broad and readily tradable stock samples

Sample	Period	U only	$U \mid D$	$D \mid U$
Historical S&P 500	1996-2014	0.321 (1.42)	0.314 (1.36)	0.032 (0.25)
Historical S&P 500	1996-2021	0.361 (1.81)	0.368 (1.80)	-0.027 (-0.24)
Historical S&P 500	2012-2021	0.571 (1.76)	**0.648 (2.08)**	-0.236 (-1.34)
Entire linked sample	1996-2021	0.028 (0.43)	0.042 (0.64)	-0.132 (-1.21)
Entire linked sample	2012-2021	0.035 (0.61)	0.056 (1.07)	-0.207 (-1.33)
Top size decile	2012-2021	0.140 (0.57)	0.166 (0.72)	-0.075 (-0.50)

Pooled panel regressions without controls, using equal observation weights and stock- and month-clustered standard errors. Returns are next-month stock returns in excess of the contemporaneous value-weighted sample return. U is firm 30-day SVIX variance minus the sample's value-weighted average firm SVIX variance, so its coefficient is the Martin–Wagner gamma rather than a percentage-point slope per signal standard deviation. D is standardized log EPS disagreement and is reported in monthly percentage points per standard deviation. The 1996–2014 S&P 500 estimate can be compared with the published gamma of 0.301.

Figure 5 traces the three return coefficients across mutually exclusive investability deciles. Under equal weights, the uncertainty coefficient rises from -0.171 in the first decile to 0.166 in the tenth. The disagreement coefficient moves from -0.289 to 0.031, while the IVOL coefficient rises from -0.346 to 0.130. The uncertainty and IVOL coefficients under capped value weights follow the same upward path, while the disagreement coefficient does not. The coefficients change gradually across deciles rather than at a single breakpoint.

The same pattern appears in samples defined by several pre-specified tradability screens. The uncertainty coefficient is positive among historical S&P 500 constituents, the most liquid stocks, and the highest investability group, while disagreement is small in those samples. Capped value-weighted IVOL spreads are positive across most of these screens, although individual confidence intervals are wide. Appendix Tables 23, 25, and 26 report the full estimates across alternative definitions of readily tradable stocks.

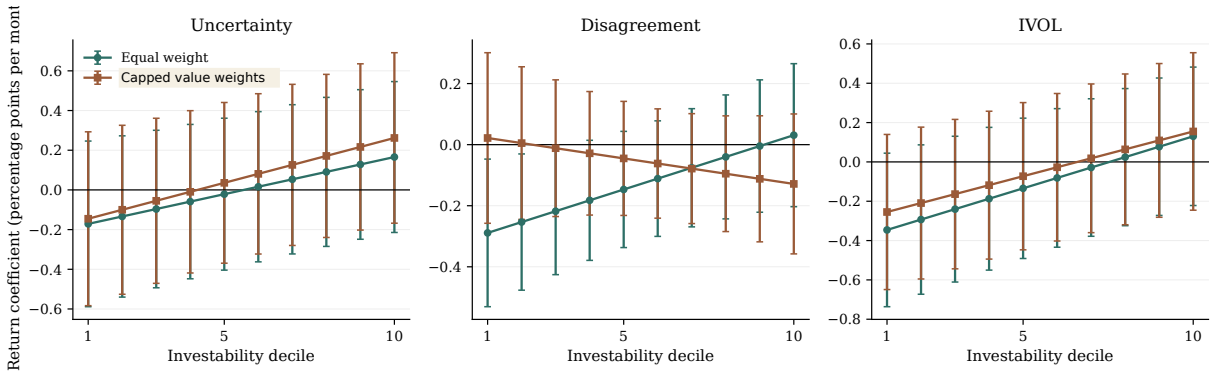


Figure 5: Return coefficients across investability. The curves report implied coefficients from monthly interaction regressions evaluated across investability deciles. A positive disagreement gradient means that its negative return coefficient becomes less negative as investability rises. Error bars denote 95 percent Newey–West confidence intervals.

5.5 Ownership and Market Participation

Size, liquidity, and borrowing conditions already trace the trading environment in which pessimistic views enter prices. Institutional ownership supplies a complementary measure of how broadly positions are distributed. Following Chen et al. (2002), I measure breadth in the S12 institutional-holdings data as the number of reporting institutions holding a stock divided by the total number of reporting institutions. Holdings are lagged before they are assigned to monthly returns. I use both raw breadth and a size-neutral measure formed by ranking breadth separately within each month and size decile.

The IVOL coefficient rises as ownership broadens. From the lowest breadth quintile to the highest, it increases by 0.235 percentage points under equal weights and by 0.389 under capped value weights. The interaction estimates in Table 5 point in the same direction. With size-neutral breadth, the IVOL interaction is 0.075 ($t = 2.09$) without controls and 0.109 ($t = 3.72$) after standard return predictors are added. The uncertainty coefficient also becomes more positive with breadth under equal weights. Interactions under capped value weights are smaller, consistent with breadth carrying its greatest variation outside the largest firms.

Table 7 separates the movement in the IVOL coefficient into uncertainty, disagreement, and a residual component. The allocation averages over the two possible orders in which uncertainty and disagreement can enter, so their covariance is not assigned by an arbitrary sequence. The decomposition is exact within month and breadth quintile. Appendix G gives the construction.

From the lowest breadth quintile to the highest, the uncertainty contribution increases by 0.200 percentage points ($t = 2.66$), while the disagreement contribution becomes 0.108 percentage points less negative ($t = 2.75$). The residual contribution declines by 0.073 percentage points, a change that is statistically indistinguishable from zero ($t = -1.06$). Under capped value weights, the IVOL coefficient rises by 0.389 ($t = 3.22$), driven mainly by a 0.311-percentage-point increase in the uncertainty contribution.

Table 7: The IVOL return slope across ownership breadth

Breadth quintile	Weighting	IVOL slope	U contribution	D contribution	Residual	Sum
1	Equal weight	-0.253 (-1.39)	-0.187 (-1.65)	**-0.139 (-3.11)**	0.073 (0.98)	-0.253 (-1.39)
1	Capped value weight	-0.206 (-1.19)	-0.199 (-1.61)	**-0.076 (-2.24)**	0.070 (1.03)	-0.206 (-1.19)
2	Equal weight	-0.077 (-0.42)	-0.036 (-0.30)	-0.031 (-0.81)	-0.010 (-0.14)	-0.077 (-0.42)
2	Capped value weight	-0.059 (-0.34)	-0.081 (-0.65)	-0.028 (-0.92)	0.050 (0.67)	-0.059 (-0.34)
3	Equal weight	-0.132 (-0.67)	-0.066 (-0.50)	-0.040 (-1.13)	-0.026 (-0.34)	-0.132 (-0.67)
3	Capped value weight	-0.121 (-0.63)	-0.070 (-0.50)	-0.035 (-1.21)	-0.016 (-0.23)	-0.121 (-0.63)
4	Equal weight	-0.047 (-0.25)	0.029 (0.21)	-0.015 (-0.49)	-0.060 (-0.72)	-0.047 (-0.25)
4	Capped value weight	-0.017 (-0.09)	0.019 (0.13)	-0.013 (-0.44)	-0.024 (-0.27)	-0.017 (-0.09)
5	Equal weight	-0.018 (-0.09)	0.013 (0.09)	-0.031 (-0.80)	0.000 (0.00)	-0.018 (-0.09)
5	Capped value weight	0.184 (0.90)	0.111 (0.73)	-0.035 (-1.19)	0.108 (1.33)	0.184 (0.90)
5 minus 1	Equal weight	**0.235 (2.32)**	**0.200 (2.66)**	**0.108 (2.75)**	-0.073 (-1.06)	**0.235 (2.32)**
5 minus 1	Capped value weight	**0.389 (3.22)**	**0.311 (3.95)**	0.041 (1.31)	0.038 (0.45)	**0.389 (3.22)**

IVOL is standardized winsorized log `idiovol`, U is ranked total 30-day firm SVIX variance, and D is standardized log EPS disagreement. The uncertainty and disagreement columns are order-invariant Shapley contributions. The residual completes the exact monthly decomposition, so the three components sum to the IVOL slope before decomposition. Entries are monthly percentage points with six-lag Newey–West t-statistics. Appendix G gives the construction.

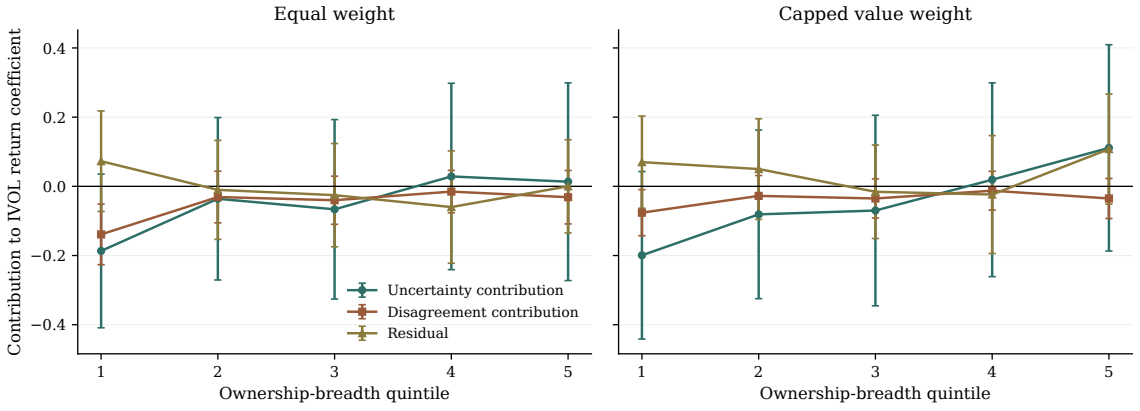


Figure 6: An exact Shapley decomposition of the IVOL return coefficient across ownership breadth. The uncertainty, disagreement, and residual contributions sum to the IVOL coefficient within every month and breadth quintile. The panels report estimates under equal and capped value weights. Error bars denote 95 percent Newey–West confidence intervals.

Composite investability and institutional breadth capture distinct aspects of market participation. Investability shifts both the composition of IVOL and the prices attached to its components. Conditional on size, breadth is more closely related to the pricing of uncertainty than to its loading within IVOL, while the disagreement loading continues to decline. The return decomposition shows that broader ownership raises the uncertainty contribution and weakens optimistic selection. The breadth gradient is an empirical analogue of the model's ownership margin; the numerical one-half threshold is a feature of the location–scale benchmark. Size neutralization and control residualization preserve this ordering. Appendix Tables 20–22 report the full results.

5.6 Extreme Within-Stock IVOL States

The cross-sectional tests compare firms at a given point in time. As complementary evidence, I examine what follows when a stock's own IVOL becomes unusually high relative to its prior 60-month history, using only lagged observations. In the full sample, average next-month returns decline from 0.81 percent in the $(-1, 0]$ state to -0.12 percent when the past-only IVOL score exceeds three. The temporal tail spread is -0.94 percentage points ($t = -4.94$).

Because the far-tail decline appears among both small and large stocks and in groups with stronger and weaker shorting ability, it is not a clean test of the ownership frontier, but it provides complementary evidence on nonlinear state dependence. Noisy information with stochastic fundamental volatility can raise uncertainty and disagreement together in extreme episodes; heterogeneous signal precision, including the sticky zero/infinite-precision limit, or heterogeneous priors can generate additional rare belief spikes, paralleling the SPF exceedances. Reversal and lottery demand remain competing explanations. Appendix H.5 reports the full construction, table, and figure.

5.7 Survey Evidence

Survey evidence allows a closer comparison of uncertainty and disagreement for a common payoff and horizon. The Livingston Survey and the CFO Survey conducted by Duke University and the Federal Reserve Banks of Richmond and Atlanta measure dispersion in forecasts of a common annual market outcome. Each survey uses its source-specific one-year option denominator: Livingston is paired with constructed 365-day index SVIX variance, while CFO is paired with squared Cboe VIX1Y. Their mean ratios are 0.096 and

0.057, respectively, and every observation lies below one. Risk-neutral variance is used directly, without separating expected physical variance from the variance-risk component.

Table 8: Aggregate survey disagreement relative to market-implied uncertainty

Source	Period	T	Med. N	Mean	Median	Max.	> 1
Livingston Survey	1996-06–2021-12	52	19	0.096	0.085	0.225	0
CFO Survey	2020Q3–2025Q4	22	283	0.057	0.054	0.117	0

Constructed from respondent-level survey records with a survey-specific option denominator. Livingston Survey disagreement is the sample variance across respondents of $\log(SPIF_{12M}/SPIF_{ZM})$ and is paired within four calendar days with the nearest prior value of `svix1y_rp` from `spf_h.txt`, the constructed 365-day index SVIX variance series. CFO Survey disagreement is the sample variance of cleaned most-likely 12-month S&P 500 return forecasts, with values outside $[-100, 100]$ percentage points removed, and is paired on the official survey close date with `vix1y` from `richmond_mean_q.txt`; Cboe VIX1Y is squared to express it in variance units. The uncertainty measures are risk neutral; this comparison does not separate physical variance from the variance-risk component.

The U.S. Survey of Professional Forecasters (SPF), administered by the Federal Reserve Bank of Philadelphia, provides a direct within-survey comparison because respondent point forecasts and probability histograms refer to the same outcome and date. After excluding point forecasts that are incompatible with the reported histogram, median D/U is 0.232 for core CPI, 0.218 for core PCE, and 0.061 for unemployment. A small number of inflation observations exceed one, primarily following large shocks. More generally, the ratio declines sharply from the first to the fourth quarter as information about the same calendar-year outcome becomes common. Appendix C presents the complete construction and the corresponding figures.

6 Interpretation and Robustness

The principal specifications use one measure of idiosyncratic volatility, a common 30-day option horizon, and next-month returns. The exercises below vary the conditioning set, sample composition, signal definition, and weighting scheme, and trace the IVOL coefficient across standard screens for readily tradable stocks.

6.1 Cardinal and Ordinal Signal Definitions

The principal specification represents uncertainty by its within-month rank and disagreement by its standardized, winsorized logarithm. The first transformation discards distances while preserving the ordering of firm option variance; the second preserves proportional

differences while compressing the much longer right tail of price-scaled earnings dispersion. Appendix Table 11 asks whether the results depend on either choice. It combines seven economically interpretable transformations of each signal: winsorized levels, square roots, inverse hyperbolic sines, raw and winsorized logarithms, percentile ranks, and inverse-normal ranks. The resulting grid contains 49 otherwise identical regressions without additional controls. Holding the sample, horizon, weighting scheme, and regression design fixed isolates the consequences of measuring signal magnitudes rather than merely their cross-sectional ordering.

The disagreement coefficient is negative in all 49 specifications under both equal and capped value weights. Its median is -0.168 under equal weights and -0.120 under capped value weights. The broad uncertainty coefficient remains sensitive to weighting rather than to signal scale: it is negative in every equal-weight specification and positive in 44 of 49 capped value-weight specifications. The uncertainty–disagreement interaction is negative and statistically significant under matched asinh, log, percentile rank, and inverse-normal-rank definitions and under both weighting schemes. The conditional return pattern therefore does not depend on preserving cardinal distances or on replacing them with ordinal information.

6.2 Return Controls and the Volatility Estimand

Firm SVIX variance combines expected physical variance with the compensation investors require for bearing variance risk. The baseline coefficient therefore refers to total option-implied firm variance. Conditioning on realized volatility, IVOL, MAX, or skewness would instead identify a residual option-variance component. Standard return predictors leave the conditional disagreement coefficient at -0.185 ($t = -2.24$) and the uncertainty–disagreement interaction at -0.144 ($t = -3.45$). Appendix Table 12 reports the complete sequence.

A separate market adjustment subtracts index variance at the same horizon, scaled by the stock's squared beta, from firm SVIX variance. The uncertainty–disagreement interaction becomes slightly more negative, showing that common market variance does not drive the joint return pattern. The result remains strongest among small, difficult-to-trade, and expensive-to-borrow stocks.

6.3 Sample Composition

The linked option–analyst sample contains fewer of the smallest and least liquid stocks, where the unconditional IVOL anomaly is strongest. I therefore reweight these observations to reproduce the monthly size–liquidity distribution of the broader stock sample. The procedure gives greater influence to firms that are sparsely represented in the linked data. Appendix G describes the construction and its coverage.

The conditional disagreement coefficient changes from -0.181 to -0.333 ($t = -2.55$). With capped weights, it is -0.364 ($t = -2.79$). The uncertainty–disagreement interaction remains negative, and the positive IVOL loadings on both belief components are preserved. The disagreement result is stronger under the broader market composition. Appendix Table 18 reports the complete estimates and documents which size–liquidity bins are represented in the linked sample.

6.4 The IVOL Sign across Investability

The interaction regressions establish that the IVOL coefficient rises with investability. Portfolio returns show the economic magnitude. Figure 7 forms IVOL quintiles separately within ten mutually exclusive investability groups. In the least investable decile, the high-minus-low spread under equal weights is -0.919 percentage points per month ($t = -2.15$). By the highest decile, it is close to zero under equal weights and 0.424 under capped value weights. The decile path is nonmonotonic, while the endpoint shift is clear: the pronounced negative spread among difficult-to-trade firms gives way to a positive sign among the most investable.

Appendix Tables 24–26 show that the endpoint shift is not tied to a particular tradability screen. Neighboring large-cap, liquidity, and borrowing-cost definitions yield the same broad ordering, although individual confidence intervals remain wide. The broad IVOL anomaly and the positive estimates among readily tradable stocks therefore arise in different parts of the cross section: equal weighting emphasizes small firms, where disagreement is most strongly associated with low subsequent returns, whereas samples of large, readily tradable stocks place more weight on firms whose IVOL is closely connected to option-implied uncertainty. Changing the sample changes the economic mixture represented by IVOL and, with it, the estimated coefficient.

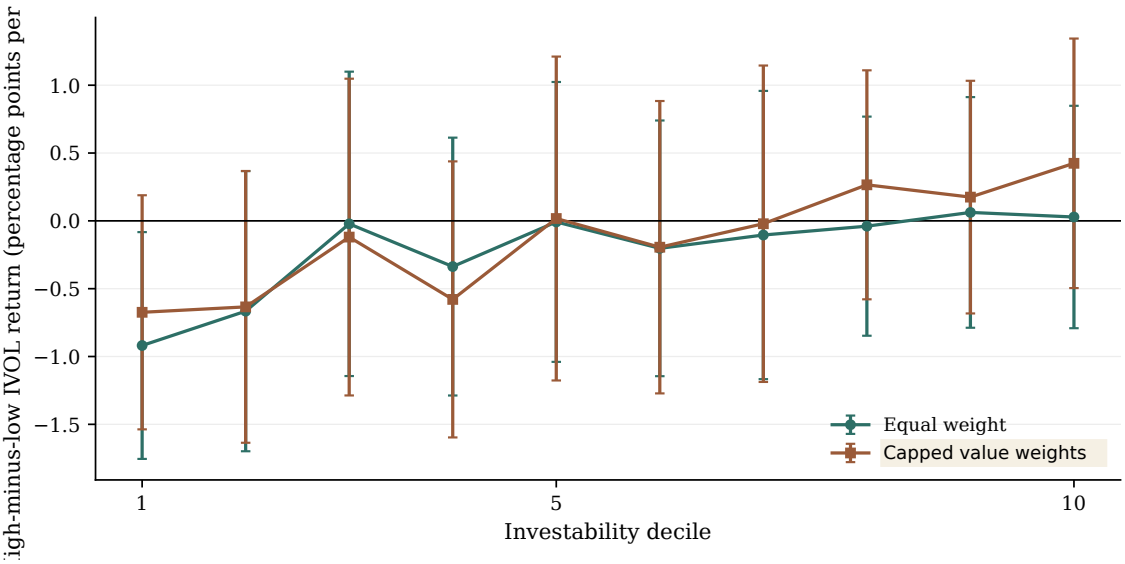


Figure 7: High-minus-low IVOL returns across investability deciles. IVOL quintiles and weights are formed separately within each monthly investability decile. Error bars denote 95 percent Newey–West confidence intervals.

6.5 Market Participation, Size, and Ownership

Institutional breadth is closely related to size, liquidity, analyst coverage, option availability, and the presence of diversified institutions. The raw breadth interaction therefore combines several dimensions of market participation. Ranking breadth within size deciles and adding standard return predictors narrows the comparison. Under equal weights, the IVOL coefficient continues to rise with size-neutral breadth, and the uncertainty coefficient becomes more positive. The corresponding estimates under capped value weights are smaller.

Breadth and composite investability reveal different aspects of the mechanism. Investability changes both IVOL's composition and the prices attached to its components. Conditional on size, breadth is more strongly related to component pricing than to the uncertainty loading within IVOL, while the disagreement loading continues to fall. The Shapley decomposition shows that the uncertainty contribution becomes more positive and, under equal weights, the disagreement contribution becomes less negative. Broader ownership therefore reduces optimistic selection through the pricing of the two belief components.

Control residualization and size neutralization each retain part of this ordering; the estimates weaken when both adjustments are imposed together. Appendix Tables 20–22

report the estimates. Together with the size, liquidity, and borrowing-cost results, the breadth analysis shows how the pricing balance changes as market participation expands.

6.6 Residual IVOL

Uncertainty and disagreement account for a large share of monthly IVOL variation. The remaining component has a mildly negative coefficient in the linked sample and a small, statistically weak coefficient among readily tradable stocks. It can include extreme recent returns, lottery demand, reversal, arbitrage costs, and incomplete diversification (Fu, 2009; Huang et al., 2010; Bali et al., 2011; Pontiff, 2006; Stambaugh et al., 2015). The belief decomposition adds to these explanations by showing that IVOL contains two components with opposing price implications and that their importance changes across the market.

6.7 Interpreting the Survey Evidence

The surveys measure belief moments more directly. Forecasts from the Livingston Survey and the CFO Survey conducted by Duke University and the Federal Reserve Banks of Richmond and Atlanta show that reported disagreement is small relative to one-year market-implied uncertainty. Within the U.S. Survey of Professional Forecasters, disagreement contracts sharply as forecasters acquire more common information about a fixed calendar-year outcome. Both patterns are consistent with the model's information-aggregation mechanism: common signals draw posterior means together more rapidly than they narrow a typical posterior distribution.

The fitted SPF distributions distinguish more sharply between information models. Disagreement relative to uncertainty lies below one in nearly every quarter and approaches the statewise ceiling of the homoskedastic noisy-information benchmark when fundamental volatility is high. Posterior widths differ across forecasters but only modestly, supporting noisy information with stochastic fundamental volatility and finite heterogeneity in signal precision. The isolated post-COVID exceedances can arise from stronger precision heterogeneity within the same noisy-information family, especially its sticky zero/infinite-precision limit, or from heterogeneous priors; they do not describe the central tendency of the data.

The market comparisons use risk-neutral variance without separating expected physical variance from the variance-risk component. Every Livingston Survey and CFO Survey ratio lies below one, while the SPF ratio has a small upper tail after major news and

declines as the fixed-event horizon shortens. This contrast is consistent with price pooling attenuating disagreement relative to uncertainty.

7 Conclusion

The negative relation between idiosyncratic volatility and subsequent stock returns is puzzling only if volatility is equated with risk. A volatile stock may be hard for each investor to value because its payoff distribution is wide, or valued differently because posterior means diverge. Both features raise IVOL but have opposing price effects. When firm-specific risk is costly to absorb, uncertainty weakens risk-bearing demand and can raise required compensation; disagreement can sustain a higher price when pessimistic views cannot be expressed through sufficiently negative positions.

The common-prior projection theorem bounds feasible combinations of these belief moments for a common payoff and horizon. The survey evidence points mainly to noisy information with stochastic fundamental volatility and modest precision heterogeneity: the statewise ratio approaches the homoskedastic benchmark's ceiling in volatile periods, while rare exceedances can arise within the same family, especially its sticky zero/infinite-precision limit after large shocks, or from heterogeneous priors. Once prices and positions enter, common market information compresses disagreement and the ownership cutoff determines whose valuations enter price.

In the location–scale economy, a mean-preserving increase in fundamental volatility lowers expected price when more than one-half of potential investors hold positive positions, but raises it when fewer than one-half do. The one-half ownership boundary is the Grossman–Miller Frontier. The fixed-float economy shows that optimistic selection survives when posterior centers, price informativeness, hedging demand, and demand elasticities move.

Empirically, IVOL loads positively on both 30-day option-implied firm uncertainty and analyst earnings forecast disagreement. As investability rises, the uncertainty loading strengthens, the disagreement loading contracts, and the IVOL coefficient rises. Disagreement predicts especially low returns where pessimistic positions are difficult to establish, whereas the uncertainty premium is clearest among historical S&P 500 constituents, liquid firms, and other investable groups. Broader institutional participation raises the IVOL coefficient mainly through a more positive uncertainty contribution. More broadly, dispersion measures combine uncertainty and disagreement, while market institutions determine whose beliefs prices reflect. The IVOL relation is therefore conditional rather than anomalous.

A Unconstrained Benchmarks

A.1 Private-Signal Investors

Under equation (20), unrestricted market clearing requires

$$\int_0^1 q_i^* \, di = \frac{\bar{m}^N - p}{\gamma U^N} = S.$$

Solving gives equation (21). Differentiating while holding the mean posterior fixed gives equation (22). The expected price is

$$\mathbb{E}[p^{N,0}] = \mu_0 - \gamma U^N S.$$

A mean-preserving increase in fundamental variance therefore lowers expected price in the unconstrained benchmark.

A.2 The Hellwig Equilibrium

Conjecture a linear price

$$P = A + BX - Cu.$$

Observing price is equivalent to observing

$$Y_P = \frac{P - A}{B} = X - \kappa u, \qquad \kappa = \frac{C}{B}.$$

Conditional on (s_i, Y_P), posterior variance is

$$C_H = \left(a^{-1} + b^{-1} + \tau_p^H\right)^{-1}, \qquad \tau_p^H = \frac{1}{\kappa^2 \Delta^2},$$

and the cross-sectional average posterior mean is

$$\bar{m}_H = C_H \left[a^{-1}\mu_0 + b^{-1}X + \tau_p^H (X - \kappa u)\right].$$

Unconstrained clearing requires

$$P = \bar{m}_H - \gamma C_H (\bar{S} + u).$$

Matching the ratio of the coefficients on u and X in the conjectured and clearing prices yields

$$\kappa = \frac{\tau_p^H \kappa + \gamma}{b^{-1} + \tau_p^H},$$

so $\kappa = \gamma b$ and $\tau_p^H = (\gamma^2 b^2 \Delta^2)^{-1}$. The equilibrium price is

$$P = C_H \left[a^{-1}\mu_0 - \gamma \bar{S} + (b^{-1} + \tau_p^H) X - (\tau_p^H \gamma b + \gamma) u \right]. \tag{63}$$

State-conditioned posterior means vary only through private signals, which gives $U_H = C_H$ and $D_H = C_H^2 / b$. Expected price is

$$\mathbb{E}[P] = \mu_0 - \gamma C_H \bar{S}.$$

Price learning consequently changes both posterior uncertainty and disagreement. With positions unrestricted, however, disagreement still has no independent positive effect on the clearing price.

A.3 Common Information

For two forecasts of the same fixed event, equation (27) implies

$$\frac{D_{t+1}}{U_{t+1}} < \frac{D_t}{U_t} \quad \Longleftrightarrow \quad \frac{\tau_{x,t+1}}{\tau_{c,t+1}} < \frac{\tau_{x,t}}{\tau_{c,t}}.$$

If private precision is fixed while common precision accumulates, uncertainty, disagreement, and their ratio all fall, with disagreement falling proportionally faster. This result motivates the survey-seasonality and market-pooling exercises without requiring survey respondents to learn from an asset price.

B Disclosure Precision and the Issuer's Choice

Suppose an issuer commits before news is realized to a public signal of precision $q \in [0, \bar{q}]$. Existing common precision is τ_c and private precision is τ_x. Bayesian plausibility leaves the expected posterior mean unchanged. Total posterior precision is $\tau_c + \tau_x + q$, so

$$U(q) = \frac{1}{\tau_c + \tau_x + q}, \qquad D(q) = \frac{\tau_x}{(\tau_c + \tau_x + q)^2}.$$

In the constrained Gaussian benchmark, the standardized ownership cutoff d remains fixed because $\sqrt{\tau_x}h(d) = \gamma S$ does not depend on q. Expected price is therefore

$$\mathbb{E}[p(q)] = \mu - \frac{d\sqrt{\tau_x}}{\tau_c + \tau_x + q}, \qquad \frac{d\mathbb{E}[p(q)]}{dq} = \frac{d\sqrt{\tau_x}}{(\tau_c + \tau_x + q)^2}.$$

Public precision raises expected price on the broad-ownership side, $d > 0$, and lowers it on the concentrated-ownership side, $d < 0$.

Let the issuer choose q to maximize $\mathbb{E}[p(q)] - C(q)$, where C is nondecreasing, convex, and differentiable. If $d \leq 0$, $q^* = 0$ is optimal. If $d > 0$, the price benefit is increasing and strictly concave in q, so the objective has a unique maximizer. An interior solution satisfies

$$C'(q^*) = \frac{d\sqrt{\tau_x}}{(\tau_c + \tau_x + q^*)^2}. \tag{64}$$

The solution is $q^* = 0$ when marginal cost at zero is at least as large as marginal benefit there, and $q^* = \bar{q}$ when marginal benefit at $\bar{q}$ is at least as large as marginal cost there. Otherwise the crossing in (64) is the unique interior optimum.

This is an ex ante information-design problem in the spirit of Kamenica and Gentzkow (2011). The firm selects a reporting technology before the news is realized, rather than selectively releasing favorable realizations. An intermediate level of precision can therefore be optimal before an initial public offering or another equity issuance even though more public information always reduces both posterior uncertainty and disagreement.

C Survey Evidence

The U.S. Survey of Professional Forecasters (SPF), administered by the Federal Reserve Bank of Philadelphia, permits uncertainty and disagreement to be measured from the same respondents for an identical outcome and forecast date. Point forecasts reveal dispersion in reported central expectations, while probability histograms reveal the width of each respondent's subjective distribution. Both moments are therefore observed without using an option-based uncertainty measure. Respondents naturally condition on public news and asset prices, but their forecasts are not weighted by the positions they take in a market-clearing allocation.

C.1 Point and Density Forecasts

Survey histograms and separately reported point forecasts are sometimes mutually inconsistent. Fitting a Normal distribution around an incoherent point forecast can then manufacture an artificially diffuse density. The benchmark construction removes this source of distortion before estimating subjective uncertainty.

For respondent j in quarter t, let π_{jtk} be the reported probability in bin $[a_k, b_k]$. After assigning finite supports to the otherwise open outer bins, the sharp set of means consistent with the histogram is

$$\mathcal{M}_{jt} = \left[\sum_k \pi_{jtk} a_k, \sum_k \pi_{jtk} b_k \right]. \tag{65}$$

A respondent is retained when the separately reported point forecast satisfies $f_{jt} \in \mathcal{M}_{jt}$, so that some distribution supported on the reported bins can possess the stated mean. The screen retains 79.7 percent of paired core consumer price index (CPI) reports, 77.9 percent of core personal consumption expenditures (PCE) reports, and 89.1 percent of unemployment reports.

Conditional on passing the screen, I fit $N(f_{jt}, \sigma_{jt}^2)$ by choosing $\sigma_{jt} > 0$ to minimize the squared distance between the ten reported bin probabilities and their Normal probabilities:

$$\widehat{\sigma}_{jt} = \arg\min_{\sigma > 0} \sum_k \left[\pi_{jtk} - \left\{ \Phi\left(\frac{b_k - f_{jt}}{\sigma} \right) - \Phi\left(\frac{a_k - f_{jt}}{\sigma} \right) \right\} \right]^2. \tag{66}$$

The two open tails remain open in this fit. Finite tail supports are used only for the consistency criterion in (65). Variable-quarter moments are computed on the identical retained respondent set:

$$D_t = \mathrm{Var}_j(f_{jt}), \qquad U_t = \mathbb{E}_j[\widehat{\sigma}_{jt}^2], \qquad G_t = \mathrm{Var}_j(\widehat{\sigma}_{jt}). \tag{67}$$

The third measure captures disagreement about uncertainty. Because G_t and U_t have the same units, G_t / U_t is directly comparable across variables.

In ordinary states, the empirical ratios are concentrated well below one, with only a small upper tail. Median D/U is 0.232 for core CPI, 0.218 for core PCE, and 0.061 for unemployment. Four inflation variable-quarters exceed one: core CPI in 2010Q1, 2011Q1, and 2023Q1, and core PCE in 2022Q1; the post-COVID ratios are 1.513 and 1.302. Because incoherent point-density pairs are excluded, these exceedances are not artifacts of centering a Normal density on an infeasible point forecast.

Table 9: SPF uncertainty and disagreement from fitted Normal distributions

Variable	T	Med. N	Med. D/U	P95	Max.	$D/U > 1$	Med. G/U
Core CPI	78	25	0.232	0.960	1.806	3	0.226
Core PCE	78	23	0.218	0.725	1.513	1	0.186
Unemployment	69	29	0.061	0.297	0.797	0	0.230

A respondent is retained only when the separately reported point forecast lies inside the sharp feasible mean interval implied by the reported histogram, using the stated finite support for otherwise open tails. For retained reports, a Normal density is centered at the point forecast and its standard deviation is chosen by minimum-distance fit to all ten reported bin probabilities. D is the population variance of point forecasts, U is the mean fitted posterior variance, and $G = \mathrm{Var}_i(\sqrt{U_i})$ measures disagreement about posterior uncertainty in the same squared units. The final column reports G/U.

All four exceedances occur in the first quarter, a timing pattern that is economically meaningful because the SPF questions concern the current calendar-year outcome. First-quarter forecasts span almost a full year, whereas fourth-quarter forecasts concern an event that is largely realized. Table 10 documents a monotone decline in the pooled mean ratio from 0.412 in Q1 to 0.069 in Q4. The median falls from 0.289 to 0.057.

Table 10: Within-year seasonality in SPF belief moments

Survey quarter	Variable-years	Mean U	Mean D	Mean D/U	Median D/U
Q1	57	0.300	0.136	0.412	0.289
Q2	58	0.416	0.163	0.289	0.224
Q3	55	0.292	0.049	0.180	0.128
Q4	55	0.210	0.012	0.069	0.057

Entries pool core CPI, core PCE, and unemployment current-year forecasts. Because every survey round targets the same calendar-year outcome, Q1 through Q4 trace a shrinking fixed-event horizon. The point and density construction is identical to Table 9.

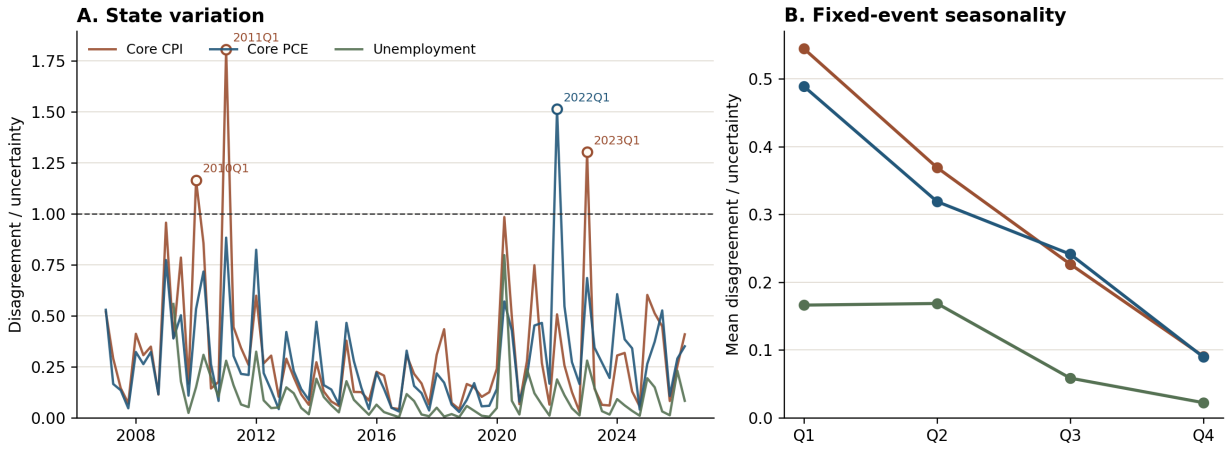


Figure 8: SPF disagreement relative to uncertainty using fitted Normal distributions. Panel A plots each variable-quarter ratio and marks $D/U = 1$, the statewise ceiling in the homoskedastic Gaussian noisy-information benchmark. Panel B averages the ratio by survey quarter for a fixed calendar-year target. Core CPI and core PCE begin in 2007Q1. Unemployment begins in 2009Q2.

C.2 Information Frictions

The fitted Normal evidence distinguishes several implications that an unconditional average would obscure. In the homoskedastic Gaussian noisy-information benchmark, posterior widths are common across forecasters, so $G = 0$, and the statewise ratio D/U cannot exceed one. Time variation in fundamental variance moves the ratio toward one in volatile periods without relaxing that ceiling. This is the dominant pattern in the data: nearly all observations lie below one, and the ratio rises markedly toward one in volatile quarters. The statewise evidence therefore places the data closer to noisy information with stochastic fundamental volatility than to its canonical sticky zero/infinite-precision limit, which imposes no statewise unit ceiling.

Taken literally, however, the homoskedastic noisy-information model is too stark because the median values of G/U, ranging from 0.186 to 0.230, show that posterior widths are not identical. Their modest magnitude points to limited heterogeneity in signal precision rather than wholesale separation into informed and uninformed groups. The two canonical models are nested in a heterogeneous-precision noisy-information family. Finite-precision specifications assign forecasters different but interior signal precisions; sticky updating is the limiting case in which each incremental signal has either zero or infinite precision, and information vintages record histories of those precision draws. Because any cross-sectional dispersion in posterior loadings removes the exact statewise ceiling, finite precision heterogeneity can itself generate rare exceedances after sufficiently large shocks.

Relative to this spectrum, the canonical unit ceiling respected almost everywhere and the modest posterior-width dispersion point to noisy information with stochastic fundamental volatility and limited precision heterogeneity, rather than the sticky limit as the central description.

The two-vintage sticky information benchmark provides a sharper comparison of the joint moments. If a share λ remains stale while the remainder observes the state,

$$\frac{G}{U} = 1 - \lambda \qquad \text{and} \qquad \frac{\mathbb{E}[D]}{U} = 1 - \lambda. \tag{68}$$

The inflation medians lie close to this one-factor equality, whereas unemployment displays more dispersion in posterior widths than in point forecasts. In addition, the time-series Spearman correlation between D/U and G/U is negative for all three variables, ruling out movements in a single stale share as an account of the joint distribution. The sticky limit is most useful for understanding the exceptional episodes. For a fixed stale share, a large gap between updated and stale forecasts raises D/U without changing G/U. The isolated post-COVID exceedances are therefore consistent with stronger precision heterogeneity within noisy information, especially the sticky zero/infinite-precision limit after a large shock. Heterogeneous priors can produce similar tail behavior. The central distribution supports noisy information with stochastic fundamental volatility and modest finite precision heterogeneity; the exceptional tail need not be attributed to a separate information friction.

The fixed-event seasonality provides a separate signature of common-information pooling. In the Gaussian projection underlying the Hellwig model,

$$\frac{D_t}{U_t} = \frac{\tau_{x,t}}{\tau_{c,t} + \tau_{x,t}}, \tag{69}$$

where τ_c is common precision and τ_x is respondent-specific precision. As releases about a fixed annual outcome accumulate, greater common precision reduces both uncertainty and disagreement, with a proportionately larger effect on disagreement. The same algebra describes information aggregated through an equilibrium price. It therefore provides a common explanation for the within-year decline in the SPF ratio and the much tighter ratios observed in financial markets, without implying that SPF respondents learn exclusively from prices.

C.3 Price Pooling and Market Beliefs

The survey-to-option ratios in Section 5 are considerably more compressed. For the Livingston Survey, the ratio of one-year return disagreement to constructed 365-day index SVIX variance averages 0.096 and never exceeds 0.225. For the CFO Survey conducted by Duke University and the Federal Reserve Banks of Richmond and Atlanta, the ratio to squared Cboe VIX1Y averages 0.057 and never exceeds 0.117. The option denominator is therefore survey specific rather than reconstructed from one common surface. All 74 survey-to-option observations lie below one, whereas the SPF distribution retains a visible upper tail and four statewise exceedances.

These market-based ratios are much more compressed than the within-survey SPF ratios. Price learning in noisy rational expectations models such as Grossman (1976) and Hellwig (1980) provides a natural explanation. Forecast dispersion is observed after substantial common market information has been produced, and the option surface incorporates information aggregated through prices. Common information raises posterior precision for every investor without creating new cross-sectional dispersion, thereby compressing D/U. SPF uncertainty is a subjective physical variance, whereas the option measures are risk neutral. For this comparison, I do not disambiguate expected physical variance from the variance-risk component; the pronounced compression of the market ratios is the central finding.

D Proof of the Common-Prior Projection Bound

Proof of Theorem 1. All expectations are taken on the investor–state product space. Center X without loss of generality and write $m = \mathbb{E}[X \mid \mathcal{H}]$ and $g = \mathbb{E}[m \mid \mathcal{G}]$. The law of total variance gives $U = V - \mathbb{E}[m^2]$. Since X is $\mathcal{G}$-measurable,

$$\mathbb{E}[(m - g)X] = 0.$$

Moreover, g is the L^2 projection of m onto the closed subspace generated by $\mathcal{G}$, and therefore

$$D = \mathbb{E}[(m - g)^2].$$

Because this subspace contains X, the residual from projecting m onto $L^2(\mathcal{G})$ cannot exceed the residual from projecting m onto the one-dimensional span of X. The latter projection coefficient is

$$\frac{\mathbb{E}[mX]}{\mathbb{E}[X^2]} = \frac{\mathbb{E}[m^2]}{V} = \frac{V-U}{V}.$$

Therefore

$$D \leq \mathbb{E}\left[\left(m - \frac{V-U}{V}X\right)^2\right] = (V-U) - \frac{(V-U)^2}{V} = U\left(1 - \frac{U}{V}\right).$$

For the nested closed subspaces $\text{span}\{X\} \subseteq L^2(\mathcal{G})$, the projection identity gives

$$\|m - P_{\text{span}\{X\}}m\|^2 = \|m - P_{L^2(\mathcal{G})}m\|^2 + \|P_{L^2(\mathcal{G})}m - P_{\text{span}\{X\}}m\|^2.$$

Consequently, equality holds if and only if the two projections coincide. If $\mathcal{G} = \sigma(X)$ and $\mathbb{E}[m \mid X] = ((V-U)/V)X$, that condition holds and the similar-triangle relation in Figure 2 is exact. The final inequality follows because $U \leq V$. The conditional formulation follows by applying the same argument under the regular conditional law given $\mathcal{C}_t$, with V, U, and D replaced by V_t, U_t, and D_t. □

E Dynamic Noisy and Sticky Information

The static benchmarks isolate the contemporaneous sources of belief dispersion. A dynamic environment adds persistence, allowing private signal noise and stale information to propagate across dates. Let

$$X_t = \rho X_{t-1} + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2), \qquad |\rho| < 1.$$

Noisy information. In the Gaussian noisy-information benchmark of Sims (2003), investor i observes a private signal

$$Y_{it} = X_t + \nu_{it}, \qquad \nu_{it} \sim N(0, \sigma_\nu^2),$$

with signal errors independent across investors and dates. Let A_t be the common prior variance immediately before the date-t signal, let U_t be the posterior variance, and let K_t be the Kalman gain. The variance recursion is

$$A_t = \rho^2 U_{t-1} + \sigma_\epsilon^2, \qquad K_t = \frac{A_t}{A_t + \sigma_\nu^2}, \qquad U_t = (1 - K_t)A_t = \sigma_\nu^2 K_t. \tag{70}$$

If $\widehat{X}_{it}$ denotes investor i's posterior mean, then

$$\widehat{X}_{it} = \rho(1 - K_t)\widehat{X}_{i,t-1} + K_t(X_t + \nu_{it}). \tag{71}$$

Consequently, conditional on the realized state history, cross-sectional disagreement satisfies

$$D_t = \rho^2(1 - K_t)^2 D_{t-1} + K_t^2 \sigma_\nu^2. \tag{72}$$

The inherited term in (72) records dispersion originating in earlier private signals. Omitting it would collapse the dynamic environment into an unrelated sequence of static problems.

In the stationary equilibrium, write K for the steady-state gain. Equations (70)–(72) give

$$U^N = \sigma_\nu^2 K, \qquad D^N = \frac{K^2 \sigma_\nu^2}{1 - \rho^2(1 - K)^2}, \qquad R^N := \frac{D^N}{U^N} = \frac{K}{1 - \rho^2(1 - K)^2}. \tag{73}$$

The statewise unit ceiling remains intact:

$$1 - R^N = \frac{(1 - K)\left[1 - \rho^2(1 - K)\right]}{1 - \rho^2(1 - K)^2} \geq 0. \tag{74}$$

Every investor has the same posterior variance, so $G_t^N = \mathrm{Var}_i(\sqrt{U_{it}}) = 0$.

The comparative statics are transparent in terms of K. Combining $A = \rho^2 U + \sigma_\epsilon^2$, $U = \sigma_\nu^2 K$, and $A = \sigma_\nu^2 K/(1 - K)$ yields

$$\frac{\sigma_\epsilon^2}{\sigma_\nu^2} = \frac{K}{1 - K} - \rho^2 K. \tag{75}$$

The right-hand side is strictly increasing in K. Moreover,

$$\frac{\partial R^N}{\partial K} = \frac{1 - \rho^2(1 - K^2)}{[1 - \rho^2(1 - K)^2]^2} > 0, \qquad \left.\frac{\partial R^N}{\partial \rho^2}\right|_K = \frac{K(1 - K)^2}{[1 - \rho^2(1 - K)^2]^2} > 0. \tag{76}$$

The steady-state ratio consequently rises with fundamental innovation variance and persistence, declines with signal-noise variance, and approaches one as fundamental volatility becomes large relative to signal noise or as signal noise vanishes. Higher innovation variance also raises both U^N and D^N. Slow variation in fundamental volatility can therefore move the local steady-state ratio toward its unit ceiling without requiring a departure from the Gaussian noisy-information structure.

Sticky information. Consider the standard geometric-vintage structure of Mankiw and Reis (2002). In this AR(1) benchmark it is exactly a limiting heterogeneous-precision noisy-information process: at each date the incremental signal has infinite precision for the fraction $1-\lambda$ that updates and zero precision for the fraction λ that does not. An age-s vintage is the history of s consecutive zero-precision draws since the most recent perfect signal. In the stationary cross section, the mass whose information is s periods old is

$$w_s = (1-\lambda)\lambda^s, \qquad s = 0, 1, 2, \ldots.$$

An age-s investor's posterior mean and variance for X_t are

$$m_{t,s} = \rho^s X_{t-s}, \qquad P_s = \sigma_\epsilon^2 \sum_{j=0}^{s-1} \rho^{2j} = \sigma_\epsilon^2 \frac{1-\rho^{2s}}{1-\rho^2}, \tag{77}$$

where $P_0 = 0$. State-conditioned disagreement is therefore

$$D_t^S = \sum_{s=0}^{\infty} w_s (\rho^s X_{t-s})^2 - \left(\sum_{s=0}^{\infty} w_s \rho^s X_{t-s} \right)^2. \tag{78}$$

Summing over vintages gives average posterior uncertainty and expected disagreement,

$$U^S = \sum_{s=0}^{\infty} w_s P_s = \frac{\lambda \sigma_\epsilon^2}{1-\lambda\rho^2}, \tag{79}$$

$$\bar{D}^S = \mathbb{E}[D_t^S] = \frac{\lambda(1-\lambda)\sigma_\epsilon^2}{(1-\lambda\rho^2)(1-\lambda^2\rho^2)}, \qquad \frac{\bar{D}^S}{U^S} = \frac{1-\lambda}{1-\lambda^2\rho^2} < 1. \tag{80}$$

The expected ratio respects the common-prior projection bound, but no analogous state-wise unit ceiling exists. Uncertainty U^S is pinned down by the parameters, whereas (78) is a quadratic function of the realized state history and is unbounded under Gaussian shocks.

The induced precision histories also generate disagreement about posterior widths. Let

$$\bar{\sigma}^S = \sum_{s=0}^{\infty} w_s \sqrt{P_s}.$$

Then

$$G^S = \sum_{s=0}^{\infty} w_s (\sqrt{P_s} - \bar{\sigma}^S)^2 > 0, \tag{81}$$

for nondegenerate parameters. Sticky information can exhibit a modest average D/U ratio while the resulting precision heterogeneity removes the statewise ceiling and generates cross-sectional dispersion in posterior widths.

Interpretation. The statewise evidence supports noisy information with stochastic fundamental volatility. The ratio remains below one in nearly all observations and rises toward the canonical noisy-information ceiling in volatile periods, as the benchmark predicts. The homoskedastic model cannot be taken literally because posterior widths differ across forecasters, but their dispersion is modest. Finite heterogeneity in signal precision provides natural intermediate cases, and sticky updating is the zero/infinite-precision limit of the same family. Information vintages are the histories generated by those precision draws rather than a distinct informational primitive. Cross-sectional precision heterogeneity can itself generate rare statewise exceedances after sufficiently large shocks; the sticky limit provides the sharpest version. Heterogeneous priors can generate similar tail behavior. These sparse episodes supplement, rather than overturn, the noisy-information interpretation. Panel histories and direct information on posterior widths help distinguish the mechanisms.

F Fixed Float with Information and Hedging Demand

The transformed-supply economy delivers a global closed-form equilibrium by matching the float schedule to constrained Gaussian demand. This appendix develops a complementary economy that holds physical float fixed and obtains noisy price revelation from rational hedging demand. The construction is related to Bai et al. (2006), who combine private payoff information with nontraded endowment exposure in a rational expectations model with short-sale constraints. Here the two motives are assigned to separate sectors, preventing fundamental disagreement from being mechanically contaminated by hedging endowments and preserving a transparent interpretation of uncertainty and disagreement.

The traded payoff is

$$V = X + \eta, \qquad X \sim N(\mu_0, \tau_0^{-1}), \qquad \eta \sim N(0, \sigma_\eta^2). \tag{82}$$

The learnable component X and residual payoff η are independent. A mass $\omega \in (0,1)$ of information traders observes

$$s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, \tau_x^{-1}), \tag{83}$$

and has no nontraded endowment. A mass $1-\omega$ of hedgers has no private fundamental signal and carries the nontraded exposure HN, where

$$H \sim N(0, \sigma_H^2), \qquad N \sim N(0, \sigma_N^2), \qquad \mathrm{Cov}(\eta, N) = \chi \in (0, \sigma_\eta \sigma_N). \tag{84}$$

The pair (η, N) is jointly Gaussian, and H is independent of (X, η, N) and all signal errors. The hedging sector observes H, whereas information traders do not. Everyone observes the equilibrium price, and physical float is the fixed constant $\bar{S} > 0$.

Begin with the unconstrained economy and conjecture a separating linear price

$$P = A + \beta_X X + \beta_H H = A + \beta_X (X + \kappa H), \qquad \kappa = \frac{\beta_H}{\beta_X}. \tag{85}$$

For information traders, observing price is equivalent to observing $Y = X + \kappa H$, with precision

$$\tau_p = \frac{1}{\kappa^2 \sigma_H^2}. \tag{86}$$

Define

$$C_X = (\tau_0 + \tau_x + \tau_p)^{-1}, \qquad U_I = C_X + \sigma_\eta^2, \qquad D_I = \tau_x C_X^2. \tag{87}$$

Here C_X is posterior uncertainty about the learnable component, U_I is total payoff uncertainty for an information trader, and D_I is cross-sectional dispersion in fundamental posterior means conditional on (X, H).

Information-trader demand is $(m_i - P)/(\gamma U_I)$, where

$$m_i = C_X(\tau_0 \mu_0 + \tau_x s_i + \tau_p Y). \tag{88}$$

On a separating branch, a hedger who observes (H, P) infers X. Hedger demand is

$$q^H = \frac{X - P - \gamma\chi H}{\gamma\sigma_\eta^2}. \tag{89}$$

Market clearing is

$$\omega \int_0^1 q_i^I \, di + (1-\omega) q^H = \bar{S}. \tag{90}$$

Proposition 8 (Separated sector linear equilibrium with fixed float)**.** *Define*

$$A_I = \frac{\omega}{U_I}, \qquad A_H = \frac{1-\omega}{\sigma_\eta^2}. \tag{91}$$

The unconstrained economy has a unique linear rational expectations equilibrium. Its price-noise loading is the unique negative solution of

$$\boxed{\kappa = -\frac{A_H \gamma\chi}{A_H + A_I C_X \tau_x}, \qquad C_X = \left(\tau_0 + \tau_x + \frac{1}{\kappa^2 \sigma_H^2}\right)^{-1}.} \tag{92}$$

The price coefficients are

$$A = \frac{A_I C_X \tau_0 \mu_0 - \gamma\bar{S}}{A_I + A_H}, \tag{93}$$

$$\beta_X = \frac{A_I C_X(\tau_x + \tau_p) + A_H}{A_I + A_H}, \qquad \beta_H = \frac{A_I C_X \tau_p \kappa - A_H \gamma\chi}{A_I + A_H}. \tag{94}$$

Expected price is

$$\mathbb{E}[P] = \mu_0 - \frac{\gamma\bar{S}}{A_I + A_H}. \tag{95}$$

Proof. Aggregate information-trader demand uses the cross-sectional mean

$$\bar{m}_I = C_X[\tau_0\mu_0 + (\tau_x + \tau_p)X + \tau_p\kappa H].$$

Multiplying market clearing by γ gives

$$A_I(\bar{m}_I - P) + A_H(X - P - \gamma\chi H) = \gamma\bar{S}.$$

Solving this equation for price gives (93) and (94). Imposing $\kappa = \beta_H/\beta_X$ cancels the common price-signal term and yields (92).

For uniqueness, write $k = -\kappa > 0$. The right side of

$$k = \frac{A_H \gamma \chi}{A_H + A_I C_X(k) \tau_x}$$

is continuous and strictly decreasing in k, because $C_X(k)$ is increasing and $A_I C_X = \omega C_X/(C_X + \sigma_\eta^2)$ is increasing. The left side is strictly increasing. At $k = 0$ the right side approaches $\gamma\chi$, whereas at $k = \gamma\chi$ it is strictly smaller than the left side. There is therefore a unique solution in $(0, \gamma\chi)$. □

The equilibrium has $\kappa < 0$ because positive hedging exposure induces the hedging sector to sell, contaminating an otherwise favorable fundamental signal in price. The extent of this contamination is endogenous and, unlike in the canonical Hellwig model, depends on the relative risk-bearing elasticities of informed traders and hedgers.

Let $v_X = \tau_0^{-1}$ denote learnable fundamental variance. Since $|\kappa|$ falls with C_X, price precision $\tau_p(C_X)$ rises with C_X. Implicit differentiation gives

$$\frac{dC_X}{dv_X} = \frac{v_X^{-2}}{C_X^{-2} + \tau_p'(C_X)} > 0. \tag{96}$$

Hence

$$\frac{dU_I}{dv_X} > 0, \qquad \frac{dD_I}{dv_X} = 2\tau_x C_X \frac{dC_X}{dv_X} > 0. \tag{97}$$

These comparative-static signs coincide with the Hellwig benchmark. The denominator in (96) captures an additional equilibrium response: greater learnable volatility increases the fundamental loading of informed demand relative to hedging pressure, thereby improving price informativeness and moderating the rise in both belief moments.

Unlearnable residual volatility σ_η^2 operates differently. Holding the information structure fixed, it raises U_I one for one without directly increasing D_I. In the full equilibrium it also alters the demand elasticities of both sectors and, through them, the informativeness of price. This indirect effect is not signed without additional parameter restrictions. The distinction has a direct empirical implication:

$$\text{learnable volatility: } U \uparrow,\ D \uparrow; \qquad \text{unlearnable volatility: } U \uparrow \text{ directly, no direct } D \uparrow . \tag{98}$$

Remark 1 (Survey dispersion and the Bayesian decomposition)**.** *The state-conditioned survey measure in* (87) *is* $D_I = \tau_x C_X^2$*. The disagreement term in the exact conditional law of total variance is different. If*

$$C_P = \mathrm{Var}(X \mid P) = (\tau_0 + \tau_p)^{-1},$$

then

$$D^B = \mathrm{Var}(\mathbb{E}[V \mid P, s_i] \mid P) = C_P - C_X = \tau_x C_P C_X, \qquad \mathrm{Var}(V \mid P) = U_I + D^B. \tag{99}$$

Both D_I *and* D^B *rise with learnable fundamental variance, but they should not be treated as numerically identical. Analyst dispersion maps more naturally to* D_I*. Equation* (99) *is the exact Bayesian accounting identity.*

Impose the informed-sector short floor $q_i^I \geq -\ell$ while keeping hedgers interior. Let $\pi(X, H)$ denote a differentiable separating branch of the constrained price rule. Information trader i forms the exact Bayesian posterior implied by that same rule, conditional on (s_i, P), and let $a_i(P)$ denote the trader's unique unconstrained CARA desired position,

$$a_i(P) \in \arg\max_{a \in \mathbb{R}} \mathbb{E}[-\exp\{-\gamma a(V - P)\} \mid s_i, P; \pi]\,. \tag{100}$$

Assume the required conditional exponential moments are finite. Equivalently, a_i uniquely minimizes

$$\Psi_i(a) = \gamma a P + \frac{\gamma^2 \sigma_\eta^2 a^2}{2} + \log \mathbb{E}[e^{-\gamma a X} \mid s_i, P; \pi]. \tag{101}$$

The objective is coercive and strictly convex because

$$\Psi_i''(a) = \gamma^2 \left[\sigma_\eta^2 + \mathrm{Var}_i^{(a)}(X \mid s_i, P; \pi)\right] > 0, \tag{102}$$

where $\mathrm{Var}_i^{(a)}$ is variance under the exponential tilt proportional to $e^{-\gamma a X}$. Its first-order condition is

$$P + \gamma \sigma_\eta^2 a_i = \frac{\mathbb{E}[X e^{-\gamma a_i X} \mid s_i, P; \pi]}{\mathbb{E}[e^{-\gamma a_i X} \mid s_i, P; \pi]}. \tag{103}$$

The short floor maps desired positions into actual positions through

$$T_\ell(a) := \max\{a, -\ell\}, \qquad q_i^I = T_\ell(a_i(P)). \tag{104}$$

Conditional on the aggregate state, aggregate informed demand is

$$Q_I(P; X, H) = \mathbb{E}_i[T_\ell(a_i(P)) \mid X, H]\,. \tag{105}$$

Clearing is

$$\omega Q_I(P; X, H) + (1-\omega)\frac{X - P - \gamma\chi H}{\gamma\sigma_\eta^2} = \bar{S}. \tag{106}$$

On a stable branch define the aggregate price elasticity and interior informed mass by

$$\mathcal{K} := -\frac{\partial}{\partial P}\left[\omega Q_I(P; X, H) + (1-\omega)q^H(P; X, H)\right] > 0, \qquad \Omega := \Pr_i(a_i(P) > -\ell \mid X, H). \tag{107}$$

The derivative in $\mathcal{K}$ includes the response of exact Bayesian desired demand to the branch price.

Proposition 9 (Selection under exact Bayesian demand with fixed float)**.** *On any stable differentiable separating branch, assume the conditional desired-position distribution has no atom at the short floor. Then:*

(i) Holding the branch's inference rule and unconstrained demand schedule fixed,

$$\frac{\partial P}{\partial \bar{S}} = -\frac{1}{\mathcal{K}} < 0, \qquad \frac{\partial P}{\partial \ell} = -\frac{\omega(1-\Omega)}{\mathcal{K}} \leq 0. \tag{108}$$

The short-capacity inequality is strict whenever a positive mass of informed traders is constrained at the floor.

(ii) At the old clearing price, suppose the conditional cross-sectional distribution of unconstrained desired positions is replaced by a mean-preserving spread, $a_i' \succeq_{\text{cx}} a_i$, while the common state, branch inference rule, and hedger demand schedule are unchanged. If the old and new clearing points are connected by a branch on which $\mathcal{K} > 0$, the clearing price weakly rises. The increase is strict whenever the mean-preserving spread raises expected actual demand:

$$\mathbb{E}_i[T_\ell(a_i') \mid X, H] > \mathbb{E}_i[T_\ell(a_i) \mid X, H]. \tag{109}$$

(iii) If desired positions form a differentiable location–scale family

$$a_i = \bar{a} + \sigma_a Z_i, \qquad \mathbb{E}[Z_i] = 0, \tag{110}$$

with a fixed continuous distribution, then

$$\frac{\partial P}{\partial \sigma_a} = \frac{\omega}{\mathcal{K}}\mathbb{E}_i[Z_i \mathbf{1}\{\bar{a} + \sigma_a Z_i > -\ell\} \mid X, H] \geq 0. \tag{111}$$

The inequality is strict whenever the floor cuts through the support.

Proof. The truncation map $T_\ell(a) = \max\{a, -\ell\}$ is increasing and convex. Differentiating (106) gives the float derivative. Because $\partial T_\ell(a)/\partial \ell = -\mathbf{1}\{a < -\ell\}$ away from a mass point at the floor,

$$\frac{\partial Q_I}{\partial \ell} = -(1 - \Omega),$$

which gives the short-capacity derivative. If $a_i' \succeq_{\text{cx}} a_i$, the definition of convex order and convexity of T_ℓ imply

$$\mathbb{E}_i[T_\ell(a_i') \mid X, H] \geq \mathbb{E}_i[T_\ell(a_i) \mid X, H].$$

At the old price, informed demand therefore weakly rises. If excess demand remains strictly decreasing in price along the branch connecting the two clearing points, the new clearing price weakly rises. For the location–scale family, differentiation under the expectation gives

$$\frac{\partial Q_I}{\partial \sigma_a} = \mathbb{E}_i[Z_i \mathbf{1}\{\bar{a} + \sigma_a Z_i > -\ell\} \mid X, H].$$

Let $c = (-\ell - \bar{a})/\sigma_a$. If $c \geq 0$, the displayed expectation integrates only nonnegative values of Z_i. If $c < 0$, mean zero gives $\mathbb{E}[Z_i \mathbf{1}\{Z_i > c\}] = -\mathbb{E}[Z_i \mathbf{1}\{Z_i \leq c\}] \geq 0$. The implicit-function theorem gives (111). □

The proposition does not rely on a particular cross-sectional distribution. It neither requires posterior beliefs generated by a nonlinear price to remain Gaussian nor reduces exact CARA demand to posterior means and variances. Ownership selection is governed by the cross-sectional distribution of unconstrained desired positions. In a Gaussian common-variance benchmark, that distribution is an affine transformation of posterior means, and the familiar disagreement result follows as a special case. More generally, the proposition identifies the demand comparison that remains valid when posterior variance is insufficient to summarize exact CARA demand. Posterior-mean variance D_I becomes sufficient only under the Gaussian common-variance specialization.

Corollary 2 (Gaussian constrained-demand specialization)**.** *Suppose, for a fixed-belief clearing perturbation, posterior valuations satisfy $m_i \mid X, H \sim N(\bar{m}_I, D_I)$ with common posterior payoff variance U_I, so that $a_i = (m_i - P)/(\gamma U_I)$. Then*

$$Q_I = \frac{\sqrt{D_I}}{\gamma U_I} h(z) - \ell, \qquad z = \frac{\bar{m}_I - P + \gamma U_I \ell}{\sqrt{D_I}}, \qquad \Omega = \Phi(z). \tag{112}$$

For the price elasticity, hold the belief moments fixed and define

$$\mathcal{K}_G = \frac{\omega \Phi(z)}{\gamma U_I} + \frac{1 - \omega}{\gamma \sigma_\eta^2} > 0. \tag{113}$$

Holding the common posterior center and all unvaried moments fixed, the clearing derivatives are

$$\frac{\partial P}{\partial \sqrt{D_I}} = \frac{\omega \phi(z)}{\gamma U_I \mathcal{K}_G} > 0, \qquad \frac{\partial P}{\partial \bar{S}} = -\frac{1}{\mathcal{K}_G} < 0, \tag{114}$$

$$\frac{\partial P}{\partial \ell} = -\frac{\omega[1-\Phi(z)]}{\mathcal{K}_G} < 0, \qquad \frac{\partial P}{\partial U_I} = -\frac{\omega[Q_I + \ell(1-\Phi(z))]}{U_I \mathcal{K}_G}. \tag{115}$$

The uncertainty derivative is negative whenever the aggregate unconstrained position of interior informed traders is positive.

The Gaussian formulas are thus a specialization of the exact Bayesian selection theorem, not an assumption imposed on posterior beliefs under an arbitrary nonlinear price. The hedging sector attenuates the magnitude of the selection effect through its contribution to price elasticity, but it does not overturn the sign. Lower float raises the clearing price, whereas greater shorting capacity admits additional pessimistic positions and weakly lowers it, strictly so whenever the floor binds for a positive mass of informed traders. The exact one-half threshold need not survive because hedger elasticity, the posterior center, and endogenous price informativeness may all respond to a primitive volatility shock.

Once the short floor binds, the global price rule is generally nonlinear, and rational posteriors must be recovered from that same nonlinear rule. The linear precision in (86) therefore cannot be held fixed after the constraint is imposed. Proposition 9 is exact along a stable branch without requiring Gaussian posterior closure. It establishes the ownership-selection result, although it does not deliver a closed-form global price surface. The two economies are consequently complementary: transformed supply provides a global closed-form constrained equilibrium, while the separated-sector economy provides fixed float, exact Bayesian demand, and a clean separation between information and hedging. Bounded-signal variants can additionally generate complete withdrawal by informed traders and the pooling regions emphasized by Bai et al. (2006), in which prices lose information and perceived uncertainty rises discontinuously.

G Empirical Construction

This appendix documents the stock-level measures and estimators used in Section 5. Every predictor is available by the end of month t, and the outcome is the total stock return realized in month $t+1$.

G.1 Sample Construction and Security Identifiers

The main panel links I/B/E/S analyst forecasts, OptionMetrics option surfaces and implied borrowing rates, CRSP returns and market equity, lagged S12 holdings, and the public firm characteristics associated with Gu et al. (2020). Security links are applied using only information available at each date. The CRSP sample contains regular-way common equity issued by U.S.-incorporated firms, and total returns include cash and delisting distributions.

The option–analyst panel spans February 1996 through December 2021 and contains 311 formation months and 595,181 stock-month observations. Market equity and price are measured at the end of month t, while accounting and trading characteristics are the latest values publicly available at formation. Cross-sectional ranks and standardizations are calculated after the point-in-time panel has been assembled.

G.2 Idiosyncratic Volatility

Idiosyncratic volatility is the `idiovol` characteristic of Gu et al. (2020). The broader-sample benchmark contains 1,443,616 stock-month observations over the common period. In the baseline, composition, continuous market structure, conditional sign, and IVOL-slope decomposition regressions, positive values are logged, winsorized at the monthly 1st and 99th percentiles, and standardized within month. The ownership breadth quintile regressions instead use the standardized within-month percentile rank of IVOL.

For the within-stock exercise, current log IVOL is standardized using only the stock's preceding 60 monthly observations, with at least 36 observations required:

$$Z_{it}^{IVOL} = \frac{\log IVOL_{it} - \overline{\log IVOL}_{i,t-60:t-1}}{s(\log IVOL)_{i,t-60:t-1}}. \tag{116}$$

The current observation and all future observations are excluded from the normalizing history. The reported states are $z \leq -1$, $(-1, 0]$, $(0, 1]$, $(1, 2]$, $(2, 3]$, and $z > 3$. The tail contrast subtracts the return in $(-1, 0]$ from the return in $z > 3$.

G.3 Analyst Earnings Forecast Disagreement

The disagreement measure uses the dispersion of annual EPS forecasts. I retain U.S.-dollar forecasts for U.S. firms, require at least three analysts, and require the fiscal endpoint to

lie 180 to 550 days after the forecast date. When two forward fiscal years coexist, I select the endpoint nearest one year. Let $\mathrm{SD}(EPS)_{it}$ be the cross-analyst standard deviation and $P_{it} > 0$ the CRSP formation price. The measure is

$$D_{it}^{EPS} = \left(\frac{\mathrm{SD}(EPS)_{it}}{P_{it}} \right)^2. \tag{117}$$

Price scaling places earnings dispersion in approximate equity-return-variance units. Firm-month observations with zero dispersion are excluded. The principal pricing regressions, continuous market structure tests, IVOL-slope decompositions, and poststratification exercises use $\log D_{it}^{EPS}$, winsorized at the monthly 1st and 99th percentiles and standardized within month. The component-loading table, ownership breadth quintiles, and conditional sign maps instead use the within-month percentile rank of D_{it}^{EPS}, standardized to unit variance. This common rank scale makes the relative loadings of uncertainty and disagreement directly interpretable and limits sensitivity to the right tail. These rank-based exercises recompute the signals among the stocks in the indicated sample.

Analyst dispersion concerns earnings, whereas option variance concerns the value of equity. The two measures therefore enter the stock-level analysis separately. The same-payoff D/U comparison is reserved for the survey exercises.

G.4 Option-Implied Firm Uncertainty

Firm uncertainty is measured from OptionMetrics surfaces at a constant 30-day horizon. For each firm-month, I use the final valid surface date in month t. The forward price is inferred from put–call parity. Out-of-the-money puts below the forward and calls above it are integrated across strikes, with the interpolated put and call averaged at the forward. At least eight fitted strike nodes are required. If S_{it} is spot price, $R_{f,t,h}$ is the gross risk-free return, $Q_{it}(K_k)$ is the fitted option premium, and ΔK_k is the midpoint strike width, total simple variance is

$$U_{it,h}^{total} = \frac{2R_{f,t,h}}{S_{it}^2} \sum_{k=1}^{K_{it,h}} \Delta K_k Q_{it}(K_k). \tag{118}$$

The 30-day quantity is annualized by dividing by $30/365$. Except in the Martin–Wagner benchmark described below, the stock-level regressions use its within-month percentile rank, standardized to unit variance. Restricted composition and conditional sign exercises recompute the rank among the stocks in the indicated sample.

The market-adjusted measure subtracts index variance at the same horizon:

$$U_{it,h}^{idio} = \max\left\{U_{it,h}^{total} - \beta_{it}^2 U_{t,h}^{mkt}, 0\right\}. \tag{119}$$

Here β_{it} is measured before formation and $U_{t,h}^{mkt}$ is S&P 500 option variance at the same horizon. The adjustment removes a beta-scaled market component.

For the Martin–Wagner comparison, the regressor is 30-day firm SVIX variance net of the value-weighted average firm SVIX variance in the sample. The dependent variable is the stock's next-month return net of the contemporaneous value-weighted sample return. The resulting gamma is expressed in variance units and is therefore reported separately from the ranked signal coefficients.

G.5 Transformation Robustness

The signal-scale exercise holds the sample, 30-day horizon, next-month return, weighting scheme, and monthly regression design fixed. It therefore isolates the consequences of the signal representation. For each signal, the alternatives are the monthly 1st/99th-percentile winsorized level, the winsorized square root, the inverse hyperbolic sine of the level scaled by its monthly median, the logarithm, the winsorized logarithm, the within-month percentile rank, and the inverse-normal rank. Every variable is standardized within month after transformation. The level-based definitions retain cardinal information; the square-root, inverse-hyperbolic-sine, and logarithmic definitions preserve magnitude while progressively limiting tail leverage; and the two rank definitions retain only cross-sectional ordering. Unwinsorized levels are excluded because the extreme right tail of price-scaled EPS dispersion would allow a few observations to dominate the estimates. The reported grid contains all 49 pairings of the seven uncertainty and seven disagreement definitions. A second exercise applies the same asinh, log, percentile-rank, or inverse-normal rank transformation to both signals and adds their interaction. These regressions contain no additional controls.

G.6 Return Regressions and Inference

The baseline pricing regressions contain uncertainty and disagreement without additional controls. A second set adds standardized monthly values of log market equity, book-to-market, 12-month momentum, one-month reversal, operating profitability, asset growth, and log one plus analyst coverage. IVOL, realized volatility, MAX, and realized skewness

are excluded from this set because including them would change the estimand attached to firm SVIX variance.

Unless a table states otherwise, one cross-sectional regression is estimated each month and the reported coefficient is the time-series mean of the monthly slopes. Six-lag Newey–West standard errors allow for serial correlation in these estimates. Main monthly regressions require at least 100 firms; subgroup regressions require at least 30. Equal weighting assigns one unit to each firm. Capped value weighting clips positive formation market equity at the monthly 99th percentile and then normalizes the weights within the monthly sample.

The Martin–Wagner regressions are pooled panels with equal observation weights and standard errors clustered by stock and month. Their raw-variance coefficient has different units from the Fama–MacBeth slopes.

G.7 Market Structure and Interaction Regressions

Composite investability averages monthly ranks of size, liquidity, and ease of borrowing, requiring all three components, and then ranks and standardizes the result within month. Liquidity combines the favorable orientations of bid–ask spread, Amihud illiquidity, and zero trading. Higher investability denotes broader and easier participation.

The index of weak pessimistic trading capacity averages the unfavorable ranks of small size, bid–ask spread, Amihud illiquidity, and zero trading:

$$C_{it} = \frac{1}{4}\left[\text{rank}_t(-ME_{it}) + \text{rank}_t(BA_{it}) + \text{rank}_t(ILLIQ_{it}) + \text{rank}_t(ZERO_{it})\right]. \tag{120}$$

The borrowing measure is the OptionMetrics implied rate for the maturity nearest 365 days among contracts with 30 to 550 days remaining on the final valid date in month t. Borrowing cost is ranked and standardized separately from the capacity index.

For a market characteristic M_{it}, the pricing regressions take the form

$$r_{i,t+1} = a_t + \beta_U U_{it} + \beta_D D_{it} + \beta_M M_{it} + \delta_U U_{it} M_{it} + \delta_D D_{it} M_{it} + \Gamma' X_{it} + \varepsilon_{i,t+1}, \tag{121}$$

with an analogous regression replacing U and D by IVOL. The composition regressions replace returns with IVOL and interact both belief measures with M. Implied curves are obtained by evaluating each month's base coefficient plus its interaction coefficient at a

common moderator value and then applying Newey–West inference to that monthly linear combination.

The triple-interaction specification contains U, D, the capacity index, all pairwise interactions, and $U \times D \times C$. Its final coefficient tests whether the joint high-uncertainty, high-disagreement return relation becomes more adverse as pessimistic trading capacity weakens.

G.8 Institutional Ownership Breadth and the IVOL Decomposition

Ownership breadth follows Chen et al. (2002). For each quarter, it is the number of reporting S12 institutions that hold the stock divided by the total number of reporting institutions. The quarterly measure is lagged before it is assigned to monthly returns. The raw measure is ranked within month; the size-neutral version is ranked separately within each month and size decile. Raw-breadth interaction regressions include size and the corresponding signal-by-size interactions.

The IVOL-slope decomposition uses standardized log IVOL, ranked uncertainty, and standardized log disagreement and is calculated separately in each month and breadth quintile. After weighted demeaning, let P_S denote the cross-sectional projection onto a subset S of $\{U, D\}$. Define

$$v(S) = \frac{\mathrm{Cov}(P_S IVOL, P_S r)}{\mathrm{Var}(IVOL)}. \tag{122}$$

The contribution of uncertainty averages its incremental value when it enters first and when it enters after disagreement:

$$\phi_U = \frac{1}{2}\{v(\{U\}) - v(\varnothing) + v(\{U, D\}) - v(\{D\})\}, \tag{123}$$

and ϕ_D is defined symmetrically. The residual equals the IVOL slope before decomposition less $v(\{U, D\})$. Hence $\phi_U + \phi_D +$ residual equals that IVOL slope in every month. For the control exercise, IVOL and returns are first residualized within month and breadth group on the seven non-volatility return controls.

G.9 Reweighting to the Broader Stock Sample

The broader stock sample supplies monthly size and liquidity strata. In the original five-by-five grid, cells accounting for 13.1 percent of assignable observations contain no linked option–analyst records. These cells are concentrated among the smallest and least-liquid firms. The main reweighting therefore uses a two-by-two grid, the finest tested rectangular partition for which every monthly size–liquidity cell is represented in the linked sample. As a sensitivity exercise, I also use the finer five-by-five grid after excluding bins that contain no linked observations.

A linked observation in size–liquidity group c receives weight

$$w_{it} = \frac{N_{c,t}^{broad}}{N_{c,t}^{link}}, \tag{124}$$

normalized to monthly mean one. Untrimmed weights and versions capped at the monthly 99th percentile are reported. The exercise reproduces the observed size–liquidity composition of the broader sample within the linked data; it does not create uncertainty or disagreement measures for firms without option and analyst observations.

G.10 Investability Portfolios and Tradability Screens

The broader-sample benchmark sorts the monthly stock sample into IVOL deciles and reports the high-minus-low next-month return. Repeating the exercise in the common option–analyst sample makes the coverage difference explicit before uncertainty and disagreement are introduced.

For the mutually exclusive investability exercise, stocks are sorted into investability deciles each month and then into IVOL quintiles within each decile. The reported spread is the return on the highest-IVOL quintile less the return on the lowest. Additional samples are defined before examining returns: the historical S&P 500, top market-cap groups, size and liquidity quantiles, borrowing-cost groups, and combinations of these screens. The megacap tests report the complete sequence of adjacent top-N and top-market-share cutoffs. IVOL breakpoints and weights are formed anew within each monthly sample.

G.11 Survey Measures of Uncertainty and Disagreement

The aggregate exercise compares one-year S&P 500 return forecast disagreement with a survey-specific one-year option variance denominator. For the Livingston Survey, administered by the Federal Reserve Bank of Philadelphia, disagreement is the cross-respondent variance of log return forecasts. Each observation is paired within four calendar days with the nearest prior value of `svix1y_rp` in `spf_h.txt`, the constructed 365-day index SVIX variance series. For the CFO Survey conducted by Duke University and the Federal Reserve Banks of Richmond and Atlanta, I clean respondent-level most-likely one-year return forecasts and calculate their cross-sectional variance. Each survey close date is paired with `vix1y` in `richmond_mean_q.txt`; Cboe VIX1Y is squared to express the denominator in variance units. The forecast construction follows Gandhi et al. (2026). The two exercises therefore retain their separate survey and option data structures rather than reconstructing a common one-year option surface in this package.

The U.S. Survey of Professional Forecasters exercise is entirely survey based. Point and density forecasts are merged by respondent, variable, and quarter. For reported histogram probabilities π_{jtk} on bins $[a_k, b_k]$, the feasible mean interval is

$$\left[\sum_k \pi_{jtk} a_k, \sum_k \pi_{jtk} b_k\right].$$

A point-density pair is excluded when the separately stated point forecast lies outside this interval. For each retained report, I center a Normal density at the point forecast and choose its standard deviation by minimum-distance fit to all reported histogram probabilities. If f_{jt} is the point forecast and $\widehat{\sigma}_{jt}$ the fitted standard deviation, I calculate

$$D_t = \mathrm{Var}_j(f_{jt}), \qquad U_t = \mathbb{E}_j[\widehat{\sigma}_{jt}^2], \qquad G_t = \mathrm{Var}_j(\widehat{\sigma}_{jt}).$$

All three quantities use the same retained respondent set. Appendix C gives the fit, retained shares, variable-quarter results, and fixed-event seasonality.

H Additional Empirical Results

The tables in this appendix report the specifications summarized in Sections 5 and 6. Samples, horizons, and timing are held fixed unless a table states otherwise.

H.1 Cardinal and Ordinal Signal Definitions

Table 11 varies the empirical scale of both belief measures while retaining the common sample, horizon, and return specification. Panel A summarizes the complete transformation grid. Panel B reports the interaction estimate when the two signals use the same transformed cardinal or ordinal definition. The negative disagreement ordering and the negative uncertainty–disagreement interaction do not depend on a single treatment of signal distances.

Table 11: Cardinal and ordinal representations of uncertainty and disagreement

Panel A: Sign stability across the 49 transformation pairs

Weighting	$D < 0$	Median D	$U > 0$	Median U
Equal weight	49/49	−0.168	0/49	−0.096
Capped value weight	49/49	−0.120	44/49	0.086

Panel B: The uncertainty–disagreement interaction under matched transformations

Transformation of U and D	Equal weight	Capped value weight
Inverse hyperbolic sine	−0.146 **(**−3.46**)**	−0.144 **(**−2.25**)**
Log, winsorized	−0.170 **(**−3.72**)**	−0.131 **(**−2.80**)**
Percentile rank	−0.182 **(**−3.56**)**	−0.152 **(**−3.08**)**
Inverse-normal rank	−0.177 **(**−3.99**)**	−0.095 **(**−2.12**)**

Monthly Fama–MacBeth coefficients are in percentage points per month, with six-lag Newey–West t-statistics in parentheses in Panel B. The outcome is the next-month total return, and all transformed signals are standardized within month. Panel A jointly estimates U and D for every pairing of seven transformations: monthly 1st/99th-percentile winsorized levels, winsorized square roots, median-scaled inverse hyperbolic sines, logs, winsorized logs, percentile ranks, and inverse-normal ranks. Panel B adds $U \times D$ and uses the indicated common transformation for both signals. The regressions contain no additional controls. Capped value weights clip market capitalization at its monthly 99th percentile. The sample contains 595,181 stock-months over 311 months.

H.2 Return Controls

Table 12 presents the complete control sequence. Panel A contains uncertainty and disagreement, Panel B adds their interaction, and Panel C uses market-adjusted firm SVIX variance. Standard return predictors strengthen the conditional disagreement estimate and leave the broad uncertainty coefficient close to zero. The uncertainty–disagreement interaction remains negative throughout the principal specifications.

H.3 Market Structure

Table 13 reports the full interaction results, including non-volatility-control specifications. Investability raises the uncertainty and IVOL coefficients under both weighting schemes.

Table 12: Uncertainty and disagreement with conventional characteristics

Panel	Controls	Weighting	U	D	$U \times D$
A: Total 30-day U and D	No controls	Equal weight	-0.055 (-0.32)	-0.181 (-1.77)	–
A: Total 30-day U and D	No controls	Capped value weight	0.093 (0.51)	-0.112 (-1.03)	–
A: Total 30-day U and D	Non-volatility controls	Equal weight	-0.144 (-0.86)	**-0.185 (-2.24)**	–
A: Total 30-day U and D	Non-volatility controls	Capped value weight	-0.028 (-0.15)	-0.065 (-0.75)	–
B: Total 30-day U, D, interaction	No controls	Equal weight	-0.068 (-0.40)	-0.126 (-1.28)	**-0.173 (-3.52)**
B: Total 30-day U, D, interaction	No controls	Capped value weight	0.056 (0.33)	-0.232 (-1.90)	**-0.147 (-3.04)**
B: Total 30-day U, D, interaction	Non-volatility controls	Equal weight	-0.169 (-1.00)	-0.145 (-1.82)	**-0.144 (-3.45)**
B: Total 30-day U, D, interaction	Non-volatility controls	Capped value weight	-0.049 (-0.27)	-0.134 (-1.44)	**-0.088 (-2.05)**
C: Market-adjusted 30-day U	No controls	Equal weight	-0.096 (-0.73)	-0.083 (-0.69)	**-0.214 (-4.82)**
C: Market-adjusted 30-day U	No controls	Capped value weight	-0.004 (-0.03)	-0.164 (-1.26)	**-0.144 (-2.88)**
C: Market-adjusted 30-day U	Non-volatility controls	Equal weight	-0.143 (-1.24)	-0.142 (-1.49)	**-0.164 (-4.37)**
C: Market-adjusted 30-day U	Non-volatility controls	Capped value weight	-0.107 (-0.92)	-0.106 (-0.95)	-0.083 (-1.87)

Monthly Fama–MacBeth coefficients in percentage points per month; six-lag Newey–West t-statistics are in parentheses. U is ranked total 30-day firm SVIX variance, D is standardized log EPS disagreement, and the non-volatility controls are size, book-to-market, momentum, reversal, profitability, investment, and log analyst coverage. IVOL, realized volatility, MAX, and skewness are excluded. Capped value weights use the monthly 99th percentile of market capitalization.

Under equal weights, the disagreement coefficient becomes less negative as investability rises and more negative as trading capacity weakens or borrowing costs increase. Ownership breadth produces its clearest pricing movement in the uncertainty coefficient.

Table 14 gives the corresponding IVOL loading regressions. Investability raises the uncertainty loading and lowers the disagreement loading. Weak trading capacity and high implied borrowing costs reverse both movements. Breadth consistently reduces the disagreement loading, but does not reproduce the investability movement in the uncertainty loading once size is included.

Table 15 evaluates the implied loadings at the 10th, 50th, and 90th percentiles of each market characteristic.

H.4 The Joint Belief State and Trading Capacity

Table 16 allows the uncertainty–disagreement coefficient to vary with the trading-capacity index. Under equal weights, the joint high-uncertainty, high-disagreement return relation becomes more negative as pessimistic trading capacity weakens. The estimate under capped value weights is close to zero.

Table 13: Return coefficient gradients across market structure

Market-structure variable	Controls	Weighting	$U \times M$	$D \times M$	$IVOL \times M$
Composite investability	No controls	Equal weight	**0.108 (2.73)**	**0.103 (2.17)**	**0.153 (4.17)**
Composite investability	No controls	Capped value weight	**0.131 (2.52)**	-0.048 (-0.82)	**0.131 (3.01)**
Composite investability	Non-volatility controls	Equal weight	0.055 (1.66)	**0.107 (2.74)**	**0.102 (3.06)**
Composite investability	Non-volatility controls	Capped value weight	0.050 (1.02)	-0.044 (-0.86)	0.076 (1.90)
Raw S12 breadth	No controls	Equal weight	**0.240 (2.94)**	0.064 (0.98)	0.080 (1.15)
Raw S12 breadth	No controls	Capped value weight	0.064 (0.70)	0.044 (0.71)	-0.033 (-0.46)
Raw S12 breadth	Non-volatility controls	Equal weight	**0.251 (3.43)**	0.061 (1.04)	**0.130 (1.97)**
Raw S12 breadth	Non-volatility controls	Capped value weight	0.106 (1.34)	0.041 (0.80)	-0.006 (-0.09)
Size-neutral S12 breadth	No controls	Equal weight	**0.114 (2.33)**	0.057 (1.78)	**0.075 (2.09)**
Size-neutral S12 breadth	No controls	Capped value weight	0.068 (1.35)	-0.005 (-0.15)	0.031 (0.63)
Size-neutral S12 breadth	Non-volatility controls	Equal weight	**0.133 (3.35)**	0.048 (1.64)	**0.109 (3.72)**
Size-neutral S12 breadth	Non-volatility controls	Capped value weight	0.066 (1.66)	0.001 (0.04)	0.027 (0.65)
Weak trading capacity	No controls	Equal weight	**-0.087 (-2.08)**	**-0.100 (-2.15)**	–
Weak trading capacity	No controls	Capped value weight	**-0.138 (-2.38)**	0.050 (0.71)	–
Weak trading capacity	Non-volatility controls	Equal weight	-0.038 (-1.09)	**-0.097 (-2.54)**	–
Weak trading capacity	Non-volatility controls	Capped value weight	-0.050 (-0.92)	0.053 (0.89)	–
Implied borrowing cost	No controls	Equal weight	**-0.090 (-3.40)**	**-0.094 (-2.20)**	–
Implied borrowing cost	No controls	Capped value weight	**-0.091 (-2.14)**	0.036 (0.85)	–
Implied borrowing cost	Non-volatility controls	Equal weight	**-0.071 (-3.05)**	**-0.092 (-2.60)**	–
Implied borrowing cost	Non-volatility controls	Capped value weight	-0.044 (-1.15)	0.026 (0.72)	–

Each row is a single monthly Fama–MacBeth interaction specification. U is the standardized within-month rank of total 30-day firm SVIX variance, D is standardized log EPS disagreement, and IVOL is standardized winsorized log `idiovol`. Greater investability and breadth should make the U interaction positive and the D interaction positive; larger constraint and borrowing-cost values reverse those predictions. Raw breadth specifications also include size and the corresponding signal-by-size interactions. The IVOL interaction is reported only for investability and breadth designs.

H.5 Extreme Within-Stock IVOL States

The main cross-sectional tests compare firms at a given point in time. This complementary exercise asks what follows when a stock's own IVOL becomes unusually high relative to its prior history. Current log IVOL is standardized using only the preceding 60 months for that stock, with at least 36 historical observations required. Neither the current observation nor any future information enters the standardization.

Table 17 shows a pronounced nonlinearity. In the full sample, average next-month returns decline from 0.81 percent in the $(-1, 0]$ state to -0.12 percent when the past-only IVOL score exceeds three. The temporal tail spread is -0.94 percentage points ($t = -4.94$). The decline is present among both small and large stocks and in groups with stronger and weaker shorting ability. Figure 9 displays the far-tail deterioration.

Table 14: IVOL loadings across market structure

Market-structure variable	Controls	Weighting	$U \times M$	$D \times M$
Composite investability	No controls	Equal weight	**0.039 (6.05)**	**-0.043 (-6.20)**
Composite investability	No controls	Capped value weight	**0.058 (4.42)**	**-0.047 (-6.40)**
Composite investability	Non-volatility controls	Equal weight	0.005 (1.20)	**-0.019 (-3.74)**
Composite investability	Non-volatility controls	Capped value weight	**0.027 (2.49)**	**-0.031 (-5.10)**
Raw S12 breadth	No controls	Equal weight	**-0.062 (-4.89)**	**-0.044 (-6.99)**
Raw S12 breadth	No controls	Capped value weight	**-0.046 (-3.61)**	**-0.022 (-3.62)**
Raw S12 breadth	Non-volatility controls	Equal weight	-0.001 (-0.22)	**-0.023 (-4.45)**
Raw S12 breadth	Non-volatility controls	Capped value weight	-0.015 (-1.72)	**-0.011 (-2.00)**
Size-neutral S12 breadth	No controls	Equal weight	**-0.042 (-8.13)**	**-0.020 (-5.82)**
Size-neutral S12 breadth	No controls	Capped value weight	-0.012 (-1.75)	**-0.012 (-2.93)**
Size-neutral S12 breadth	Non-volatility controls	Equal weight	**-0.025 (-9.17)**	**-0.014 (-5.41)**
Size-neutral S12 breadth	Non-volatility controls	Capped value weight	0.000 (0.01)	**-0.013 (-3.59)**
Weak trading capacity	No controls	Equal weight	**-0.034 (-5.25)**	**0.034 (5.14)**
Weak trading capacity	No controls	Capped value weight	-0.014 (-1.08)	**0.037 (4.44)**
Weak trading capacity	Non-volatility controls	Equal weight	0.002 (0.39)	**0.013 (2.45)**
Weak trading capacity	Non-volatility controls	Capped value weight	0.010 (1.00)	**0.018 (2.78)**
Implied borrowing cost	No controls	Equal weight	**-0.031 (-7.51)**	**0.045 (7.94)**
Implied borrowing cost	No controls	Capped value weight	**-0.061 (-6.50)**	**0.040 (8.60)**
Implied borrowing cost	Non-volatility controls	Equal weight	**-0.011 (-4.01)**	**0.028 (7.34)**
Implied borrowing cost	Non-volatility controls	Capped value weight	**-0.040 (-5.26)**	**0.033 (7.57)**

The dependent variable is standardized winsorized log GKX `idiovol`. U is the standardized within-month rank of total 30-day firm SVIX variance, and D is standardized log EPS disagreement. Market-structure variables are standardized within month. Positive $U \times M$ and negative $D \times M$ coefficients indicate that IVOL becomes more uncertainty-heavy as investability or breadth rises. Raw breadth specifications include size and both signal-by-size interactions.

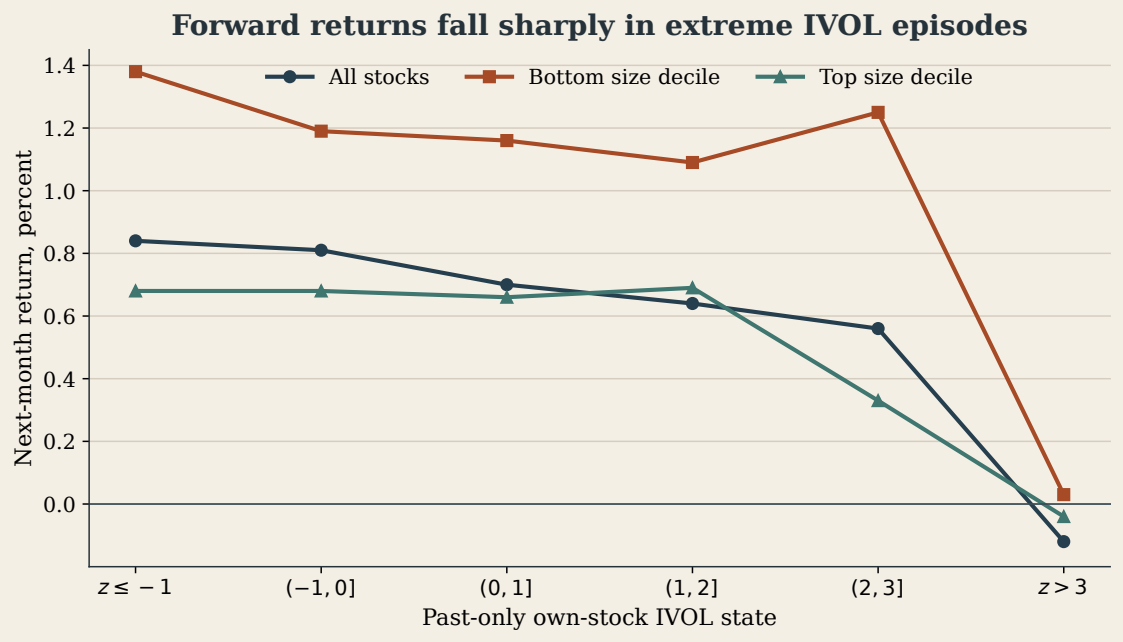


Figure 9: Next-month returns following extreme own-stock IVOL states. The state variable is current log IVOL standardized using only the stock's preceding 60 months. The figure therefore describes a within-stock time-series relation rather than a cross-sectional IVOL portfolio.

Table 15: IVOL loadings across the market-structure distribution

Market-structure variable	Weighting	Loading	P10	Median	P90	P90 minus P10
Composite investability	Equal weight	U loading	**0.614 (24.28)**	**0.668 (32.05)**	**0.723 (36.46)**	**0.109 (6.05)**
Composite investability	Equal weight	D loading	**0.195 (7.24)**	**0.134 (6.11)**	**0.074 (3.57)**	**-0.120 (-6.20)**
Composite investability	Capped value weight	U loading	**0.636 (21.19)**	**0.717 (48.13)**	**0.798 (55.55)**	**0.161 (4.42)**
Composite investability	Capped value weight	D loading	**0.216 (7.39)**	**0.151 (6.75)**	**0.086 (4.56)**	**-0.131 (-6.40)**
Raw S12 breadth	Equal weight	U loading	**0.690 (29.11)**	**0.603 (24.63)**	**0.517 (14.56)**	**-0.173 (-4.89)**
Raw S12 breadth	Equal weight	D loading	**0.195 (10.52)**	**0.135 (6.05)**	**0.074 (2.63)**	**-0.121 (-6.99)**
Raw S12 breadth	Capped value weight	U loading	**0.729 (25.92)**	**0.665 (25.97)**	**0.602 (17.80)**	**-0.127 (-3.61)**
Raw S12 breadth	Capped value weight	D loading	**0.172 (6.93)**	**0.141 (5.99)**	**0.111 (4.38)**	**-0.061 (-3.62)**
Size-neutral S12 breadth	Equal weight	U loading	**0.732 (36.29)**	**0.674 (31.50)**	**0.617 (24.98)**	**-0.115 (-8.13)**
Size-neutral S12 breadth	Equal weight	D loading	**0.182 (8.68)**	**0.154 (6.93)**	**0.126 (5.16)**	**-0.056 (-5.82)**
Size-neutral S12 breadth	Capped value weight	U loading	**0.747 (42.16)**	**0.730 (54.64)**	**0.713 (47.15)**	-0.034 (-1.75)
Size-neutral S12 breadth	Capped value weight	D loading	**0.129 (6.11)**	**0.113 (5.72)**	**0.097 (4.86)**	**-0.033 (-2.93)**
Weak trading capacity	Equal weight	U loading	**0.705 (33.40)**	**0.664 (29.62)**	**0.615 (22.70)**	**-0.090 (-5.25)**
Weak trading capacity	Equal weight	D loading	**0.098 (4.37)**	**0.140 (6.17)**	**0.191 (7.20)**	**0.092 (5.14)**
Weak trading capacity	Capped value weight	U loading	**0.781 (56.37)**	**0.764 (50.65)**	**0.744 (24.62)**	-0.037 (-1.08)
Weak trading capacity	Capped value weight	D loading	**0.092 (4.75)**	**0.136 (6.68)**	**0.190 (7.02)**	**0.098 (4.44)**
Implied borrowing cost	Equal weight	U loading	**0.738 (41.32)**	**0.694 (36.02)**	**0.651 (29.37)**	**-0.087 (-7.51)**
Implied borrowing cost	Equal weight	D loading	**0.076 (4.23)**	**0.139 (6.43)**	**0.202 (7.44)**	**0.125 (7.94)**
Implied borrowing cost	Capped value weight	U loading	**0.810 (62.89)**	**0.726 (63.67)**	**0.641 (30.89)**	**-0.169 (-6.50)**
Implied borrowing cost	Capped value weight	D loading	**0.070 (3.90)**	**0.126 (5.99)**	**0.182 (7.17)**	**0.111 (8.60)**

The entries are implied loadings from the no-control interaction regressions in Table 14 and inherit that table's signal transformations. P10, the median, and P90 refer to the pooled distribution of the standardized market-structure variable. The final column is a direct six-lag Newey–West test of the P90-minus-P10 difference. For constraint and borrowing-cost measures, moving from P10 to P90 means that pessimistic trading becomes harder.

Table 16: Trading capacity and the uncertainty–disagreement interaction

Controls	Weighting	$U \times D$	$D \times C$	$U \times D \times C$
No controls	Equal weight	**-0.139 (-3.01)**	0.011 (0.24)	**-0.060 (-1.98)**
No controls	Capped value weight	**-0.145 (-2.48)**	0.172 (1.74)	0.010 (0.23)
Non-volatility controls	Equal weight	**-0.117 (-2.65)**	0.002 (0.06)	**-0.074 (-2.58)**
Non-volatility controls	Capped value weight	**-0.125 (-2.29)**	0.140 (1.81)	-0.022 (-0.56)

U is the standardized within-month rank of total 30-day firm SVIX variance, D is standardized log EPS disagreement, and C is standardized within month. C rises as size falls and bid–ask spreads, Amihud illiquidity, and zero trading increase. The triple coefficient summarizes whether the negative joint high-U, high-D return relation becomes more pronounced as pessimistic trading capacity deteriorates.

Table 17: Returns after extreme own-stock IVOL states

Group	$z \leq -1$	$(-1, 0]$	$(0, 1]$	$(1, 2]$	$(2, 3]$	$z > 3$	Tail spread
All	0.84	0.81	0.70	0.64	0.56	-0.12	**-0.94 (-4.94)**
Small half	0.88	0.86	0.71	0.65	0.60	-0.11	**-0.93 (-3.57)**
Large half	0.78	0.78	0.72	0.72	0.59	-0.07	**-0.83 (-3.94)**
Bottom size decile	1.38	1.19	1.16	1.09	1.25	0.03	**-1.20 (-2.49)**
Top size decile	0.68	0.68	0.66	0.69	0.33	-0.04	**-0.77 (-2.61)**
Weakest shorting-ability tercile	0.61	0.65	0.52	0.55	0.39	0.02	**-0.61 (-2.22)**
Strongest shorting-ability tercile	1.00	1.00	0.84	0.75	0.71	-0.03	**-1.01 (-3.78)**

Next-month returns under equal weights, in percent. IVOL is the GKX `idiovol` characteristic. The state is current log IVOL standardized using only the preceding 60 monthly observations for that stock, with at least 36 observations. The current and all future observations are excluded. Returns are winsorized within month at the 1st and 99th percentiles. The final column is the temporal tail spread: the $z > 3$ return minus the $(-1, 0]$ return for the same stock group. It is a within-stock temporal contrast rather than a cross-sectional IVOL portfolio. Parentheses contain six-lag Newey–West t-statistics. The broad far-tail decline establishes nonlinear state dependence across groups; its presence in both stronger- and weaker-shorting groups means that it is not a direct ownership-frontier test.

Because the far-tail decline is present in both stronger- and weaker-shorting groups, this exercise is not a direct test of the ownership frontier. It instead documents nonlinear state dependence. Noisy information with stochastic fundamental volatility can raise uncertainty and disagreement together in unusually volatile episodes. Stronger precision heterogeneity within the same family, including the sticky zero/infinite-precision limit, or heterogeneous priors can generate additional rare belief spikes, as in the tail of the SPF distribution. Reversal and lottery demand remain alternative explanations for the return decline.

H.6 Reweighting toward the Broader Stock Sample

Table 18 reweights the linked option–analyst sample toward the monthly size–liquidity distribution of the broader stock sample. The principal design uses the finest rectangular partition for which every monthly cell is represented in the linked sample. The finer five-by-five calculation uses the cells represented in the linked data. Both approaches make the disagreement coefficient more negative and preserve the uncertainty–disagreement interaction.

Table 18: Reweighting the linked sample toward the broader stock sample

Design	Weighting scheme	IVOL	U	D	$U \times D$
IVOL on U and D	Unweighted	–	**0.684 (33.06)**	**0.151 (6.63)**	–
IVOL on U and D	2×2 grid, all bins represented	–	**0.667 (32.89)**	**0.182 (8.10)**	–
IVOL on U and D	2×2 grid, all bins, capped weights	–	**0.667 (32.87)**	**0.181 (8.02)**	–
IVOL on U and D	5×5 grid, represented bins only	–	**0.676 (36.32)**	**0.184 (7.91)**	–
IVOL on U and D	5×5 grid, represented bins, capped weights	–	**0.673 (33.96)**	**0.177 (7.67)**	–
Return on U and D	Unweighted	–	-0.055 (-0.32)	-0.181 (-1.77)	–
Return on U and D	2×2 grid, all bins represented	–	0.007 (0.04)	**-0.333 (-2.55)**	–
Return on U and D	2×2 grid, all bins, capped weights	–	-0.000 (-0.00)	**-0.364 (-2.79)**	–
Return on U and D	5×5 grid, represented bins only	–	-0.092 (-0.49)	**-0.394 (-2.93)**	–
Return on U and D	5×5 grid, represented bins, capped weights	–	-0.082 (-0.46)	**-0.325 (-2.76)**	–
Return on U, D, and U x D	Unweighted	–	-0.068 (-0.40)	-0.126 (-1.28)	**-0.173 (-3.52)**
Return on U, D, and U x D	2×2 grid, all bins represented	–	0.004 (0.02)	-0.216 (-1.82)	**-0.147 (-2.23)**
Return on U, D, and U x D	2×2 grid, all bins, capped weights	–	0.001 (0.00)	-0.221 (-1.86)	**-0.173 (-2.77)**
Return on U, D, and U x D	5×5 grid, represented bins only	–	-0.096 (-0.52)	-0.139 (-1.08)	**-0.296 (-4.33)**
Return on U, D, and U x D	5×5 grid, represented bins, capped weights	–	-0.081 (-0.46)	-0.158 (-1.43)	**-0.216 (-3.92)**
Return on IVOL, U, and D	Unweighted	-0.087 (-0.73)	0.009 (0.07)	-0.160 (-1.58)	–
Return on IVOL, U, and D	2×2 grid, all bins represented	-0.144 (-0.89)	0.109 (0.63)	**-0.294 (-2.29)**	–
Return on IVOL, U, and D	2×2 grid, all bins, capped weights	-0.088 (-0.59)	0.059 (0.36)	**-0.326 (-2.55)**	–
Return on IVOL, U, and D	5×5 grid, represented bins only	-0.259 (-1.63)	0.095 (0.56)	**-0.331 (-2.48)**	–
Return on IVOL, U, and D	5×5 grid, represented bins, capped weights	-0.140 (-1.12)	0.016 (0.11)	**-0.286 (-2.52)**	–

IVOL is standardized winsorized log `idiovol`, U is the standardized within-month rank of total 30-day firm SVIX variance, and D is standardized log EPS disagreement in every design. In the five-by-five grid, cells represented in the linked sample account for 86.9 percent of observations in the broader stock sample. Unrepresented cells are concentrated among the smallest, least-liquid firms. The primary two-by-two grid provides complete monthly overlap; the five-by-five estimates use the represented cells. Each linked observation receives the ratio of the broader-sample cell count to the linked-sample cell count, normalized to monthly mean one. Capped versions trim weights at the monthly 99th percentile. Return coefficients are monthly percentage points and IVOL loading coefficients are standardized units.

H.7 Ownership Breadth

Table 19 reports the breadth-quintile uncertainty and disagreement return slopes together with IVOL's loadings on each signal. From the lowest breadth quintile to the highest, the uncertainty slope becomes more positive, the disagreement slope becomes less negative, and the disagreement loading declines sharply. The main-text decomposition allocates the accompanying movement in the IVOL return slope.

Table 20 first residualizes IVOL and returns on the non-volatility return controls. Tables 21 and 22 repeat the decomposition using size-neutral ownership breadth.

Table 19: Uncertainty, disagreement, and IVOL composition across ownership breadth

Breadth quintile	U return slope	D return slope	IVOL loading on U	IVOL loading on D
1	-0.282 (-1.51)	**-0.321 (-2.72)**	**0.465 (27.08)**	**0.219 (9.38)**
2	-0.020 (-0.12)	-0.094 (-0.96)	**0.538 (17.82)**	**0.160 (6.25)**
3	-0.082 (-0.51)	-0.025 (-0.25)	**0.591 (19.97)**	**0.125 (5.61)**
4	0.060 (0.38)	-0.031 (-0.35)	**0.637 (23.27)**	**0.091 (3.97)**
5	0.070 (0.50)	-0.071 (-0.79)	**0.672 (36.04)**	**0.111 (5.70)**

Quintiles are formed monthly on Chen–Hong–Stein S12 ownership breadth, the number of reporting funds holding a stock divided by all reporting funds. Quintile 1 has the lowest breadth. All entries are Fama–MacBeth coefficients under equal weights from monthly regressions without controls. Parentheses contain six-lag Newey–West t-statistics. IVOL, total 30-day firm SVIX variance, and EPS disagreement are each converted to a standardized within-month rank among the stocks in the indicated breadth quintile. Return slopes are in monthly percentage points per signal standard deviation. The loading columns regress ranked IVOL jointly on ranked uncertainty and disagreement.

Table 20: The Shapley decomposition after residualizing conventional characteristics

Breadth quintile	Weighting	IVOL slope	U contribution	D contribution	Residual	Sum
1	Equal weight	-0.207 (-1.29)	-0.083 (-1.51)	**-0.062 (-2.82)**	-0.062 (-0.56)	-0.207 (-1.29)
1	Capped value weight	-0.237 (-1.38)	-0.112 (-1.65)	**-0.052 (-2.15)**	-0.073 (-0.61)	-0.237 (-1.38)
2	Equal weight	-0.152 (-0.86)	-0.033 (-0.51)	-0.036 (-1.43)	-0.083 (-0.74)	-0.152 (-0.86)
2	Capped value weight	-0.125 (-0.71)	-0.059 (-0.83)	-0.027 (-0.91)	-0.039 (-0.37)	-0.125 (-0.71)
3	Equal weight	-0.154 (-0.90)	-0.039 (-0.65)	-0.050 (-1.89)	-0.065 (-0.60)	-0.154 (-0.90)
3	Capped value weight	-0.113 (-0.66)	-0.037 (-0.52)	-0.025 (-0.79)	-0.052 (-0.52)	-0.113 (-0.66)
4	Equal weight	-0.108 (-0.66)	0.015 (0.25)	-0.026 (-1.04)	-0.097 (-0.94)	-0.108 (-0.66)
4	Capped value weight	-0.030 (-0.18)	0.023 (0.37)	-0.011 (-0.35)	-0.042 (-0.41)	-0.030 (-0.18)
5	Equal weight	-0.085 (-0.53)	0.010 (0.17)	-0.039 (-1.40)	-0.056 (-0.60)	-0.085 (-0.53)
5	Capped value weight	0.037 (0.22)	0.009 (0.15)	-0.027 (-0.92)	0.055 (0.54)	0.037 (0.22)
5 minus 1	Equal weight	0.122 (1.18)	**0.093 (2.28)**	0.023 (1.04)	0.006 (0.07)	0.122 (1.18)
5 minus 1	Capped value weight	**0.274 (2.24)**	**0.121 (2.45)**	0.025 (0.90)	0.127 (1.17)	**0.274 (2.24)**

IVOL is standardized winsorized log `idiovol`, U is ranked total 30-day firm SVIX variance, and D is standardized log EPS disagreement. The decomposition is computed separately in every month and breadth quintile. For each subset of $\{U, D\}$, the value is the covariance between the projection of IVOL and the projection of returns onto that subset, divided by the variance of IVOL. The Shapley contributions are order invariant. The residual equals the IVOL slope before decomposition minus the joint projected value, so the three components sum to that slope in every month by construction. Before the decomposition, IVOL and returns are residualized cross-sectionally on the seven non-volatility controls. Entries are monthly percentage points with six-lag Newey–West t-statistics.

Table 21: The Shapley decomposition using size-neutral ownership breadth

Breadth quintile	Weighting	IVOL slope	U contribution	D contribution	Residual	Sum
1	Equal weight	-0.213 (-1.27)	-0.183 (-1.45)	**-0.111 (-2.98)**	0.081 (1.75)	-0.213 (-1.27)
1	Capped value weight	0.029 (0.15)	-0.044 (-0.29)	-0.023 (-0.84)	0.096 (1.59)	0.029 (0.15)
2	Equal weight	-0.197 (-1.07)	-0.129 (-0.97)	-0.070 (-1.90)	0.002 (0.03)	-0.197 (-1.07)
2	Capped value weight	0.016 (0.08)	0.004 (0.03)	-0.009 (-0.35)	0.021 (0.35)	0.016 (0.08)
3	Equal weight	-0.136 (-0.77)	-0.046 (-0.36)	-0.044 (-1.16)	-0.046 (-0.85)	-0.136 (-0.77)
3	Capped value weight	-0.046 (-0.27)	-0.036 (-0.28)	-0.035 (-1.21)	0.026 (0.33)	-0.046 (-0.27)
4	Equal weight	-0.066 (-0.36)	-0.033 (-0.25)	-0.025 (-0.57)	-0.008 (-0.14)	-0.066 (-0.36)
4	Capped value weight	0.075 (0.36)	0.096 (0.62)	-0.062 (-1.54)	0.041 (0.53)	0.075 (0.36)
5	Equal weight	-0.052 (-0.28)	0.027 (0.22)	-0.036 (-0.68)	-0.042 (-0.83)	-0.052 (-0.28)
5	Capped value weight	0.081 (0.44)	-0.002 (-0.01)	-0.019 (-0.46)	0.102 (1.18)	0.081 (0.44)
5 minus 1	Equal weight	0.161 (1.51)	**0.209 (2.53)**	**0.074 (2.46)**	**-0.123 (-2.97)**	0.161 (1.51)
5 minus 1	Capped value weight	0.052 (0.37)	0.042 (0.44)	0.005 (0.15)	0.006 (0.08)	0.052 (0.37)

IVOL is standardized winsorized log `idiovol`, U is ranked total 30-day firm SVIX variance, and D is standardized log EPS disagreement. The decomposition is computed separately in every month and breadth quintile. For each subset of $\{U, D\}$, the value is the covariance between the projection of IVOL and the projection of returns onto that subset, divided by the variance of IVOL. The Shapley contributions are order invariant. The residual equals the IVOL slope before decomposition minus the joint projected value, so the three components sum to that slope in every month by construction. Breadth quintiles use the size-neutral S12 measure. Entries are monthly percentage points with six-lag Newey–West t-statistics.

Table 22: Size-neutral breadth after residualizing conventional characteristics

Breadth quintile	Weighting	IVOL slope	U contribution	D contribution	Residual	Sum
1	Equal weight	-0.134 (-0.91)	-0.061 (-1.10)	**-0.044 (-2.40)**	-0.029 (-0.30)	-0.134 (-0.91)
1	Capped value weight	-0.014 (-0.08)	-0.041 (-0.66)	-0.008 (-0.24)	0.035 (0.35)	-0.014 (-0.08)
2	Equal weight	-0.178 (-1.03)	-0.038 (-0.84)	-0.031 (-1.38)	-0.109 (-0.88)	-0.178 (-1.03)
2	Capped value weight	-0.043 (-0.25)	-0.011 (-0.20)	-0.003 (-0.08)	-0.029 (-0.28)	-0.043 (-0.25)
3	Equal weight	-0.217 (-1.26)	-0.021 (-0.52)	**-0.055 (-2.57)**	-0.140 (-1.11)	-0.217 (-1.26)
3	Capped value weight	-0.115 (-0.69)	-0.039 (-0.76)	-0.041 (-1.40)	-0.035 (-0.32)	-0.115 (-0.69)
4	Equal weight	-0.081 (-0.44)	-0.000 (-0.01)	**-0.047 (-2.08)**	-0.033 (-0.25)	-0.081 (-0.44)
4	Capped value weight	0.092 (0.55)	0.067 (1.27)	-0.051 (-1.50)	0.076 (0.67)	0.092 (0.55)
5	Equal weight	-0.173 (-1.17)	0.003 (0.07)	**-0.051 (-3.06)**	-0.125 (-1.09)	-0.173 (-1.17)
5	Capped value weight	-0.051 (-0.26)	-0.047 (-0.86)	-0.042 (-1.82)	0.038 (0.27)	-0.051 (-0.26)
5 minus 1	Equal weight	-0.039 (-0.45)	0.063 (1.59)	-0.007 (-0.47)	-0.096 (-1.27)	-0.039 (-0.45)
5 minus 1	Capped value weight	-0.037 (-0.25)	-0.005 (-0.14)	-0.033 (-1.10)	0.002 (0.02)	-0.037 (-0.25)

IVOL is standardized winsorized log `idiovol`, U is ranked total 30-day firm SVIX variance, and D is standardized log EPS disagreement. The decomposition is computed separately in every month and breadth quintile. For each subset of $\{U, D\}$, the value is the covariance between the projection of IVOL and the projection of returns onto that subset, divided by the variance of IVOL. The Shapley contributions are order invariant. The residual equals the IVOL slope before decomposition minus the joint projected value, so the three components sum to that slope in every month by construction. Before the decomposition, IVOL and returns are residualized cross-sectionally on the seven non-volatility controls. Breadth quintiles use the size-neutral S12 measure. Entries are monthly percentage points with six-lag Newey–West t-statistics.

H.8 Alternative Tradability Screens and the IVOL Sign

Table 23 holds one return specification fixed across seven pre-specified samples defined by tradability screens. The uncertainty coefficient is positive in every row and significant among historical S&P 500 constituents, the most liquid quintile, and the highest-investability quintile.

Table 23: The priced uncertainty channel and the IVOL sign among readily tradable stocks

Sample	$IVOL \mid U, D$	$U \mid IVOL, D$	$D \mid IVOL, U$	EW IVOL Q5–Q1	Capped-VW IVOL Q5–Q1
Historical S&P 500	-0.078 (-0.83)	**0.197 (2.04)**	-0.034 (-0.40)	0.133 (0.37)	0.279 (0.73)
Most liquid quintile	-0.175 (-1.53)	**0.215 (2.05)**	-0.010 (-0.11)	-0.057 (-0.14)	0.458 (0.98)
Top investability quintile	-0.123 (-1.03)	**0.198 (1.99)**	-0.055 (-0.65)	-0.010 (-0.02)	0.292 (0.66)
Easy-to-borrow tercile	-0.129 (-0.99)	0.159 (1.35)	-0.022 (-0.22)	0.026 (0.06)	0.320 (0.63)
Top 500 by market cap	-0.072 (-0.59)	0.147 (1.49)	-0.092 (-1.13)	-0.013 (-0.03)	0.294 (0.66)
Large and easy-to-borrow	-0.094 (-0.79)	0.137 (1.41)	-0.050 (-0.62)	0.022 (0.05)	0.343 (0.80)
Largest size quintile	-0.029 (-0.24)	0.104 (1.06)	-0.093 (-1.16)	0.057 (0.13)	0.286 (0.67)

The first three columns are full-period Fama–MacBeth coefficients under equal weights from a monthly regression of next-month returns on standardized log GKX `idiovol`, ranked total 30-day firm SVIX variance, and ranked EPS disagreement. All three signals are recomputed within each indicated sample and month. The last two columns are high-minus-low IVOL quintile returns formed within the same indicated sample. Entries are monthly percentage points with six-lag Newey–West t-statistics in parentheses. The positive capped-VW portfolio spread in every row is economically consistent with uncertainty becoming the dominant priced IVOL component, but none of the individual spread estimates is statistically distinguishable from zero.

Table 24 extends the same three-variable specification to the broader set of tradability screens.

Table 24: Uncertainty, disagreement, and IVOL across tradability screens

Sample	IVOL only	IVOL with U, D	U with IVOL,D	D with IVOL,U
Entire linked sample	-0.166 (-0.95)	-0.099 (-0.82)	-0.004 (-0.03)	-0.119 (-1.23)
Historical S&P 500	0.053 (0.40)	-0.078 (-0.83)	**0.197 (2.04)**	-0.034 (-0.40)
Largest size quintile	0.015 (0.09)	-0.029 (-0.24)	0.104 (1.06)	-0.093 (-1.16)
Most liquid quintile	-0.023 (-0.15)	-0.175 (-1.53)	**0.215 (2.05)**	-0.010 (-0.11)
Easy-to-borrow tercile	-0.018 (-0.11)	-0.129 (-0.99)	0.159 (1.35)	-0.022 (-0.22)
Top investability quintile	-0.002 (-0.01)	-0.123 (-1.03)	**0.198 (1.99)**	-0.055 (-0.65)
Bottom risk-direction quintile	**-0.356 (-2.08)**	-0.142 (-1.47)	-0.174 (-1.14)	-0.132 (-1.13)
Top risk-direction quintile	-0.058 (-0.32)	-0.084 (-0.57)	0.077 (0.61)	-0.091 (-1.10)

Monthly Fama–MacBeth coefficients under equal weights, in percentage points per month, with six-lag Newey–West t-statistics in parentheses. IVOL is standardized log GKX `idiovol`. U and D are standardized within-sample ranks of total 30-day firm SVIX variance and EPS disagreement, respectively. All three signals are recomputed within each indicated sample and month. All specifications use the same complete observations within each row and contain no controls. Investability combines size, liquidity, and ease of borrowing. The risk-direction score combines investability with the within-month rank of U dominance over D.

Tables 25 and 26 report the portfolio sign maps. Positive estimates under capped value weights cluster in samples of readily tradable and megacapitalization stocks, while the confidence intervals remain wide.

Table 25: Positive IVOL estimates under capped value weights across tradability screens

Pre-specified monthly sample	EW Q5–Q1	Capped-VW Q5–Q1	Mean stocks/month
Historical S&P 500	0.133 (0.37)	0.279 (0.73)	462
Top 500 by market cap	-0.013 (-0.03)	0.294 (0.66)	500
Largest size quintile	0.057 (0.13)	0.286 (0.67)	384
Largest size decile	0.036 (0.08)	0.226 (0.56)	192
Most liquid quintile	-0.057 (-0.14)	0.458 (0.98)	384
Most liquid decile	0.013 (0.03)	0.613 (1.21)	192
Easy-to-borrow tercile	0.026 (0.06)	0.320 (0.63)	630
Easiest-to-borrow decile	-0.076 (-0.19)	0.389 (0.77)	189
Top investability quintile	-0.010 (-0.02)	0.292 (0.66)	379
Top investability decile	0.008 (0.02)	0.453 (0.99)	190
Large and easy-to-borrow	0.022 (0.05)	0.343 (0.80)	229
Historical S&P, U-dominant tercile	-0.230 (-0.49)	0.381 (0.75)	66
Investable quintile, U-dominant tercile	-0.259 (-0.44)	0.356 (0.57)	67
Top risk-direction quintile	-0.117 (-0.22)	0.353 (0.64)	378
Top risk-direction decile	-0.361 (-0.72)	0.172 (0.32)	189

Entries are high-minus-low IVOL quintile returns in monthly percentage points over February 1996–December 2021. IVOL bins and weights are recomputed within each indicated sample. Returns are winsorized monthly at the 1st and 99th percentiles. Parentheses contain six-lag Newey–West t-statistics. Positive estimates point opposite to the usual IVOL anomaly. Fifteen of the 17 pre-specified tradability screens produce positive estimates under capped value weights; the entire linked sample and the bottom risk-direction quintile are the two exceptions. Confidence intervals for the individual screens include zero.

Table 26: Adjacent megacap cutoffs and the positive IVOL spread

Monthly megacap sample	EW Q5–Q1	Capped-VW Q5–Q1	Average stocks in smallest IVOL quintile
Top 25 by market cap	0.326 (0.80)	0.311 (0.82)	5.0
Top 50 by market cap	0.270 (0.70)	0.355 (1.03)	10.0
Top 75 by market cap	0.277 (0.74)	0.398 (1.10)	15.0
Top 100 by market cap	0.095 (0.24)	0.290 (0.76)	20.0
Top 150 by market cap	0.040 (0.10)	0.191 (0.48)	30.0
Top 200 by market cap	0.089 (0.21)	0.249 (0.61)	40.0
Top 300 by market cap	0.139 (0.31)	0.307 (0.72)	60.0
Top 500 by market cap	-0.014 (-0.03)	0.286 (0.64)	100.0
Top 750 by market cap	-0.125 (-0.28)	0.251 (0.55)	150.0
Top 1% by market cap	0.306 (0.85)	0.237 (0.68)	3.4
Top 2% by market cap	0.363 (0.95)	0.386 (1.07)	7.1
Top 5% by market cap	0.169 (0.44)	0.308 (0.81)	18.7
Top 10% by market cap	0.047 (0.11)	0.225 (0.56)	37.9

Each month, the sample is selected using lagged market capitalization and then divided into five IVOL portfolios. Entries are high-minus-low next-month returns in percentage points over the full February 1996–December 2021 period, with six-lag Newey–West t-statistics in parentheses. The sequence is reported to show stability across adjacent cutoffs, not to select a preferred breakpoint. Capped-VW point estimates are positive at every reported cutoff, but no individual estimate reaches conventional two-sided statistical significance.

I Price Information and Belief Dispersion

The main text solves information aggregation and constrained market clearing within a single equilibrium. This appendix records the implications of that equilibrium for the disagreement–uncertainty ratio and for the amount of private information absorbed by the price.

Corollary 3 (Structure-specific market ceiling)**.** *In the full rational expectations equilibrium of Proposition 5,*

$$\frac{D}{U} = \frac{\tau_x}{\tau_0 + \tau_x + \tau_p} < 1 \tag{125}$$

whenever prior or price information is nondegenerate. The ratio is strictly decreasing in price precision τ_p and converges to zero as price becomes fully revealing.

Although the universal projection theorem permits a ratio as high as one, the market structure generally implies a substantially lower ceiling because investors condition on a common price signal before choosing demand. A low empirical ratio should therefore not be read as evidence that disagreement is economically irrelevant. It may instead indicate that prices have already pooled a large share of private information.

Let the corresponding survey measures, formed without conditioning on price, be

$$U^S = \frac{1}{\tau_0 + \tau_x}, \qquad D^S = \frac{\tau_x}{(\tau_0 + \tau_x)^2}. \tag{126}$$

The market counterparts are

$$U^P = \frac{1}{\tau_0 + \tau_x + \tau_p}, \qquad D^P = \frac{\tau_x}{(\tau_0 + \tau_x + \tau_p)^2}. \tag{127}$$

Corollary 4 (Information-pooling wedge)**.** *Define the share of the survey disagreement ratio removed by price learning as*

$$\Pi = 1 - \frac{D^P/U^P}{D^S/U^S}. \tag{128}$$

Then

$$\Pi = \frac{\tau_p}{\tau_0 + \tau_x + \tau_p}. \tag{129}$$

Moreover, $D^P < D^S$ whenever $\tau_p > 0$.

The pooling wedge is endogenous to the same constrained equilibrium that generates the frontier. It is not an external correction appended to a separate information block. Price

precision also enters the mean-supply threshold in (60). More informative prices therefore reduce residual disagreement and, simultaneously, contract the region in which ownership selection dominates risk bearing.

Empirically, one-year index option variance supplies the market uncertainty scale and one-year index-return forecasts supply disagreement. Because option variance is risk neutral whereas the theorem concerns physical posterior moments, the aggregate exercise is a comparison of scale rather than a direct estimate of the projection inequality. At the stock level, option variance serves as a forward-looking proxy for uncertainty. The analysis does not separately identify expected physical variance and the variance-risk component. It asks whether total uncertainty priced in options and analyst disagreement have distinct relations with IVOL and subsequent returns.

References

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance* 61(1), 259–299.

Bai, Y., E. C. Chang, and J. Wang (2006). Asset prices under short-sale constraints. Working paper, University of Hong Kong and MIT.

Bali, T. G., N. Cakici, and R. F. Whitelaw (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics* 99(2), 427–446.

Banerjee, S. (2011). Learning from prices and the dispersion in beliefs. *Review of Financial Studies* 24(9), 3025–3068.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance* 56(1), 1–43.

Chen, J., H. Hong, and J. C. Stein (2002). Breadth of ownership and stock returns. *Journal of Financial Economics* 66(2–3), 171–205.

Cohen, L., K. B. Diether, and C. J. Malloy (2007). Supply and demand shifts in the shorting market. *Journal of Finance* 62(5), 2061–2096.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105(8), 2644–2678.

Daniel, K., A. Klos, and S. Rottke (2023). The dynamics of disagreement. *Review of Financial Studies* 36(6), 2431–2467.

Diamond, D. W. and R. E. Verrecchia (1987). Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics* 18(2), 277–311.

Diether, K., C. Malloy, and A. Scherbina (2002). Differences of opinion and the cross section of stock returns. *Journal of Finance* 57(5), 2113–2141.

Fu, F. (2009). Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics* 91(1), 24–37.

Gandhi, M., N. J. Gormsen, and E. Lazarus (2026). Forward return expectations. NBER Working Paper 31687, revised April 2026.

Goulding, C. L., C. R. Harvey, and H. Kurtović (2026). Disagreement of disagreement. NBER Working Paper 35049.

Grossman, S. J. (1976). On the efficiency of competitive stock markets where trades have diverse information. *Journal of Finance* 31, 573–585.

Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.

Gu, S., B. T. Kelly, and D. Xiu (2020). Empirical asset pricing via machine learning. *Review of Financial Studies* 33(5), 2223–2273.

Harrison, J. M. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics* 92(2), 323–336.

Hellwig, M. (1980). On the aggregation of information in competitive markets. *Journal of Economic Theory* 22(3), 477–498.

Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* 54(6), 2143–2184.

Huang, W., Q. Liu, S. G. Rhee, and L. Zhang (2010). Return reversals, idiosyncratic risk, and expected returns. *Review of Financial Studies* 23(1), 147–168.

Jones, C. M. and O. A. Lamont (2002). Short-sale constraints and stock returns. *Journal of Financial Economics* 66(2–3), 207–239.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review* 101(6), 2590–2615.

Kaplan, S. N., T. J. Moskowitz, and B. A. Sensoy (2013). The effects of stock lending on security prices: An experiment. *Journal of Finance* 68(5), 1891–1936.

Lahiri, K. and X. Sheng (2010). Measuring forecast uncertainty by disagreement: The missing link. *Journal of Applied Econometrics* 25(4), 514–538.

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics* 117(4), 1295–1328.

Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics* 132(1), 367–433.

Martin, I. and C. Wagner (2019). What is the expected return on a stock? *Journal of Finance* 74(4), 1887–1929.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance* 32(4), 1151–1168.

Nagel, S. (2005). Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics* 78(2), 277–309.

Patton, A. and A. Timmermann (2010). Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics* 57(7), 803–820.

Pontiff, J. (2006). Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics* 42(1–2), 35–52.

Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of Political Economy* 111(6), 1183–1220.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.

Stambaugh, R. F., J. Yu, and Y. Yuan (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance* 70(5), 1903–1948.

Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. In P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (Eds.), *Knowledge, Information, and Expectations in Modern Macroeconomics*, pp. 25–58. Princeton University Press.

Zarnowitz, V. and L. A. Lambros (1987). Consensus and uncertainty in economic prediction. *Journal of Political Economy* 95(3), 591–621.