Uncertainty, Disagreement and the Idiosyncratic Volatility Puzzle*

Anand Bharadwaj†

NOVEMBER 8, 2026

Abstract

Why do stocks with high idiosyncratic volatility earn low returns? Greater fundamental volatility raises both uncertainty and disagreement, with disagreement increasing proportionately more. Under short sale constraints, these forces have opposing effects on prices. We establish a sharp common prior bound relating uncertainty and disagreement and characterize an exact nonlinear rational expectations equilibrium with risk-averse investors, continuously varying positions, and learning from prices. Greater fundamental volatility can raise prices when ownership is sufficiently narrow. Consistent with the theory, we find empirically that the negative relationship between idiosyncratic volatility and subsequent returns strengthens as borrowing costs rise and ownership breadth declines, and that disagreement accounts for much of the anomaly around earnings announcements, when firm-specific public information is revealed.

JEL classification: G12, G14, D82, D83.

Keywords: Idiosyncratic volatility; disagreement; uncertainty; rational expectations equilibrium; short sale constraints; information aggregation.

*I am indebted to Alan Auerbach, Yuriy Gorodnichenko, Eben Lazarus, Martin Lettau, Emi Nakamura, Chris Shannon, Jón Steinsson, and Johan Walden for their guidance on this paper. I am also especially grateful to Michael Woodford, Nicholas Bloom, and Richard Stanton for many suggestions that have greatly improved the paper. For valuable comments and feedback, I also thank Hassan Afrouzi, Robert Barro, Onur Bayar, Jamie Cross, David Frankel, Jonathon Hazell, Barry Hughes, Michael Jansson, Palani-Rajan Kadapakkam, Roman Kapuscinski, Chen Lian, Qingmin Liu, Peter Maxted, 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

The idiosyncratic volatility puzzle, the finding that stocks with high idiosyncratic volatility (IVOL) earn low subsequent returns (Ang et al., 2006, 2009), is one of the most durable anomalies in empirical asset pricing. The pattern is neither confined to a single episode nor fading in recent data. Figure 1 shows that since 1962, a value-weighted portfolio that buys the highest IVOL decile and sells the lowest has earned a three-factor alpha of -1.46 percentage points per month ($t = -6.35$), with a negative alpha in 96 percent of rolling five-year windows. The puzzle arises because the standard logic of risk compensation points, if anything, in the opposite direction. If firm-specific risk is fully diversifiable, additional volatility need not command any premium, but if investors must bear that risk, greater uncertainty should require greater compensation, lowering the current price and raising the expected return. In this paper, we revisit the puzzle by asking how an increase in fundamental volatility affects an asset's equilibrium price and expected return. Answering this question requires understanding how greater volatility changes the beliefs on which investors trade, and how those beliefs enter market prices when some investors face constraints on the positions they can take.

Indeed, two properties of investors' beliefs turn out to be particularly important. An investor may remain uncertain about the eventual payoff of an asset after taking account of all the information available to her, while different investors may also disagree about the payoff they expect to receive. We distinguish uncertainty within investors' posterior distributions from disagreement across their posterior means, and show that these two forms of dispersion alter asset demand through different margins. Holding an investor's posterior mean fixed, greater posterior variance reduces desired exposure through the usual compensation for bearing risk, whereas greater dispersion in posterior means redistributes desired positions across investors. With unrestricted positions and a common posterior variance, a widening of posterior means induces relatively optimistic investors to buy more and relatively pessimistic investors to sell more, leaving aggregate demand, and hence the market-clearing price, unchanged. Once pessimistic positions are constrained, however, that cancellation fails: sufficiently low valuations can no longer be expressed through correspondingly negative positions, although high valuations can continue to support larger positive holdings. Greater disagreement can consequently raise the current price, as Miller (1977) first argued, even as uncertainty discourages investors from bearing risk. The effect of greater fundamental volatility on prices may therefore depend on how it changes uncertainty and disagreement, whose implications for asset demand can run in opposite directions.

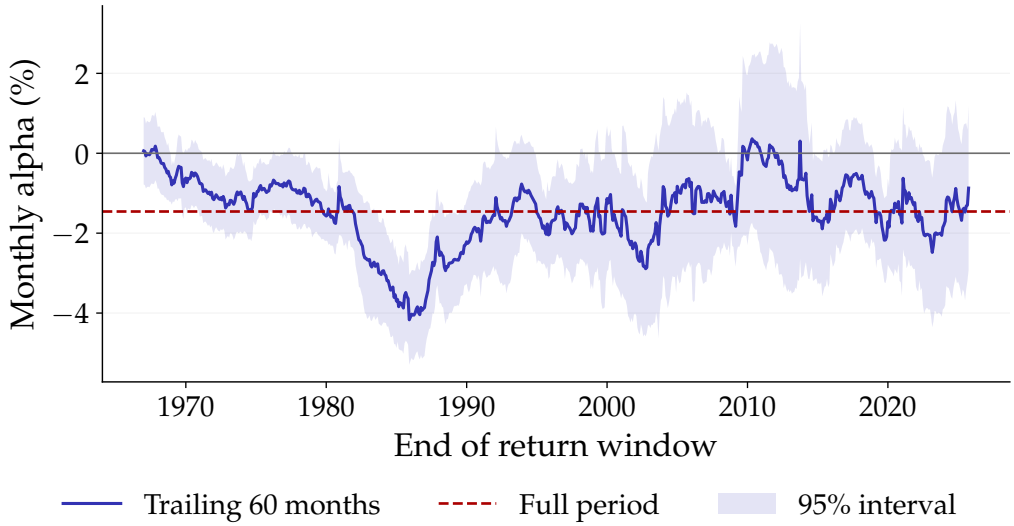


Figure 1: The IVOL–return relation since 1962. The figure plots the intercept from a regression of the value-weighted return spread between the highest and lowest deciles of one-year daily FF3 IVOL on the market, size, and value factors of Fama and French (1993), estimated over trailing 60-month windows. Portfolio returns are CRSP returns, including delisting returns, without winsorization. Returns run from January 1962 through September 2025, and the first window ends in December 1966. Shading gives pointwise 95 percent confidence intervals based on six lag Newey–West standard errors, and the dashed line marks the full period estimate. Values are in monthly percentage points, and negative values indicate the IVOL anomaly.

We begin, therefore, by analyzing uncertainty and disagreement in a Bayesian learning model, in which both arise from the information investors receive about a common underlying payoff. An important implication of this approach is that uncertainty and disagreement cannot vary independently: when investors share a common prior and form their beliefs rationally, the distribution of their posterior means is restricted by the uncertainty that remains within their individual posterior distributions. We establish a general bound showing that expected disagreement cannot exceed expected posterior uncertainty, together with a sharper bound relating both quantities to the prior variance of the payoff. The intuition is that investors who have learned very little cannot disagree greatly because their posterior means remain close to the common prior mean, while investors who have learned almost everything must largely agree on the realized payoff. Substantial disagreement can arise between these extremes, but its magnitude is constrained by how much uncertainty remains. In a Gaussian learning model, the sharper bound is attained, and the same structure reveals how uncertainty and disagreement respond differently to changes in fundamentals and information. Greater fundamental

variance raises both, but raises disagreement proportionately more, whereas additional common information draws posterior means together faster than it narrows individual posterior distributions. We examine these predictions in the Survey of Professional Forecasters, where respondents report both point forecasts and subjective probability distributions for the same eventual outcome. The survey data support the common prior restriction and the predicted response of uncertainty and disagreement as information accumulates.

The relationship between uncertainty and disagreement is particularly important in financial markets because investors learn not only from their own private information but also from the prices at which assets trade. A market price reflects the demands of other investors and thereby conveys information about the signals they have received, allowing investors to learn from the trading process itself. Even if private information initially generates substantial differences in investors' posterior means, the information revealed through market clearing may compress that disagreement, and in familiar rational expectations settings can eliminate its separate upward effect on valuation once investors have conditioned on the price. Whether disagreement survives this learning process, how much survives, and whether the remaining dispersion affects prices when pessimistic positions are constrained must therefore be determined jointly with the information revealed in equilibrium. The question cannot be settled by considering disagreement and short sale constraints separately from the information conveyed by market prices.

A central theoretical contribution of this paper is an exact characterization of a rational expectations equilibrium in which risk-averse investors receive heterogeneous private signals, learn from the market price, and face a finite short sale floor. Stochastic asset supply prevents the price from fully revealing the payoff, while the truncation of continuously varying demands generally makes market clearing nonlinear. Consequently, the equilibrium price need not be affine, posterior beliefs conditioned on that price need not remain Gaussian, and the familiar mean-variance demand formula is no longer available. With CARA utility and Gaussian private signals, however, weighting an investor's posterior by marginal utility is equivalent to translating the private signal entering the likelihood while leaving the common information inferred from price intact. This allows each investor's desired holding to be expressed through the difference between her private signal and a public statistic inferred from the price. Although the position floor truncates those holdings, their distribution across investors can still be integrated exactly, since the fundamental payoff cancels from the differences that determine their desired positions. Market clearing determines the noise contained in the statistic revealed by price, while Bayesian inference determines the valuation attached to that statistic. Short sale

constraints therefore affect the informational content of prices even when the precision of private information is held fixed. We obtain an exact nonlinear rational expectations equilibrium without imposing an affine price function or a Gaussian payoff prior, and recover the classical unrestricted information structure of Hellwig (1980) as the short sale floor recedes.

The equilibrium also allows us to examine how constraints affect the information contained in market prices. When an investor's desired position lies below the short sale floor, her actual holdings cannot respond further to changes in her private signal, so the information conveyed through her demand is correspondingly limited. A given disturbance to asset supply must therefore produce a larger movement in the public statistic revealed by price when more investors are constrained. Relaxing the short sale constraint allows a larger fraction of privately informed investors to adjust their positions, reducing the influence of supply disturbances on the information conveyed by price and lowering posterior uncertainty. When the price becomes sufficiently informative, disagreement contracts proportionately more than uncertainty, mirroring the response to additional common information in the analysis of beliefs. The model thus establishes how constraints affect both the disagreement that remains after investors learn from prices and the extent to which that disagreement influences valuation.

The exact solution also yields a characterization of how greater fundamental volatility affects equilibrium prices. A useful property of the model is that fundamental risk can vary without changing the distribution of equilibrium holdings, allowing the effect of volatility on valuation to be separated from changes in the allocation of the asset. The sign of this effect depends on the breadth of ownership, which we measure by the fraction of potential investors who hold a positive quantity of the asset. When relatively few investors hold the asset, market clearing draws disproportionately on the valuations of investors who have received favorable private information, so greater fundamental volatility can raise expected prices as the increase in disagreement outweighs the additional compensation required for bearing uncertainty. As the asset comes to be held by a broader set of investors, progressively less optimistic valuations enter market demand, until the conventional effect of risk on valuation dominates. In the Gaussian specialization, the boundary occurs when half of the potential investors hold the asset, and we establish corresponding ownership conditions in the more general equilibrium. The same reasoning clarifies the role of firm size. Holding the expected liquidating payoff per unit fixed, a larger outstanding supply requires the asset to be held by a broader set of investors, extending market clearing farther down the distribution of valuations. Market capitalization provides an empirical

counterpart to this supply variation, while the principal empirical tests examine ownership breadth and constraints on short selling directly.

These theoretical results suggest that the negative relation between IVOL and subsequent returns should be strongest where trading constraints allow disagreement to have the greatest effect on equilibrium prices, and that uncertainty and disagreement should have different return implications even though both increase with fundamental volatility. In U.S. stock return data, we examine these predictions using option-implied variance as a measure of uncertainty and dispersion in analysts' earnings forecasts as a measure of disagreement. IVOL is strongly associated with both measures, but their implications for subsequent returns differ: disagreement predicts lower returns, whereas uncertainty has a positive coefficient in regressions that also control for observed volatility, consistent with compensation for bearing risk. The conditional evidence follows the comparative statics of the model closely. As implied stock borrowing costs rise, the negative return relations associated with both disagreement and IVOL become more pronounced, consistent with the greater influence of optimistic valuations when pessimistic investors face tighter trading constraints. We also examine the role of ownership directly, using the fraction of reporting mutual funds that hold a stock, following Chen et al. (2002). As ownership broadens, the negative IVOL return slope approaches zero, and the disagreement coefficient becomes much less negative. Uncertainty and disagreement together account for most of the difference in the IVOL slope between narrowly and broadly held stocks. Finally, we examine ownership and borrowing costs jointly by sorting stocks along both dimensions. Among stocks that are held by relatively few investors and are expensive to short, both IVOL and disagreement predict significantly lower subsequent returns, while the difference in the IVOL relation across the two dimensions is also statistically significant. To this end, the negative relationships of IVOL and disagreement with subsequent returns are particularly pronounced among stocks with narrow ownership and high borrowing costs, consistent with the theory.

The information structure of the model also suggests when the negative return relation associated with disagreement should be most pronounced. Disagreement can sustain elevated valuations when pessimistic investors are constrained, but the arrival of common information draws posterior means together and can dissipate this premium. Earnings announcements provide a setting in which substantial information becomes public at identifiable dates, allowing us to examine whether the return implications of disagreement are concentrated around these events. Consistent with this mechanism, the negative return relation associated with prior disagreement is concentrated in the three trading

days surrounding earnings announcements and is essentially absent otherwise, while uncertainty displays no comparable pattern. The low returns of stocks with high IVOL are similarly concentrated around announcements, and controlling for disagreement accounts for much of the negative IVOL coefficient during these windows. By contrast, the measures of lottery demand examined here do little to account for this concentration of negative returns, even after controlling for recent return reversal. Firms with greater prior disagreement subsequently report more disappointing earnings, yet disagreement continues to predict lower announcement returns after conditioning on the realized earnings surprise.

The negative average IVOL coefficient can therefore coexist with conventional compensation for bearing risk. Greater fundamental volatility changes both uncertainty within investors' beliefs and disagreement across their valuations, but their effects on equilibrium prices differ, while learning from market prices and constraints on trading determine how strongly those effects are expressed. The breadth of ownership governs which valuations enter market demand, and the arrival of public information changes the disagreement that remains among investors. An unconditional IVOL coefficient averages across these differences and does not identify a single price of risk. The IVOL puzzle can thus be understood as a conditional phenomenon arising from rational learning and constrained trading.

1.1 Relation to the Literature

Because the IVOL puzzle has resisted a single explanation, the literature has developed accounts that differ substantially in what they take idiosyncratic volatility to represent. Fu (2009) emphasizes expected rather than lagged realized volatility, whereas Chen and Petkova (2012) and Herskovic et al. (2016) show that variation measured at the firm level can still contain exposure to aggregate priced risks. Other explanations emphasize reversal and illiquidity (Huang et al., 2010; Han and Lesmond, 2011), preferences for lottery-like payoffs (Boyer et al., 2010; Bali et al., 2011), or the asymmetry created when correcting overpricing requires a short position (Stambaugh et al., 2015), while Jiang et al. (2009) show that high IVOL predicts negative earnings surprises, which accounts for part of the puzzle. Hou and Loh (2016) document substantial heterogeneity in the explanatory power of these channels, while Goulding et al. (2026) develop a theoretically motivated composite measure of disagreement that combines IVOL with short interest, analyst dispersion, expected skewness, and related signals. The framework developed here organizes part of that heterogeneity, since greater fundamental volatility moves both uncertainty and

disagreement, whose effects on price depend on the positions investors are able to take. Several established explanations are mediated in part through these belief components or the trading conditions that govern their pricing, though others, notably lottery demand, retain some independent explanatory power. Around earnings announcements, however, the negative IVOL relation is largely unaffected by measures of lottery demand, whereas controlling for disagreement substantially reduces it.

Since the contribution of Miller (1977), the potential for disagreement to affect asset prices when pessimistic investors are constrained has occupied an important place in theories of heterogeneous beliefs. Harrison and Kreps (1978) and Scheinkman and Xiong (2003) develop related valuation effects in settings with disagreement and opportunities for subsequent trade, while Banerjee and Kremer (2010) study the implications of differing interpretations of public information for trading volume and return dynamics. Unlike these differences of opinion models, in which disagreement can persist even when investors have the same information, our framework derives disagreement from heterogeneous private signals under a common prior and accounts for the information revealed by equilibrium prices. Investor participation also affects valuation in Merton (1987), where investors know and hold only subsets of available securities, while Chen et al. (2002) connect the selection of optimistic owners to the breadth of the investor base. Empirically, Diether et al. (2002) document lower subsequent returns among stocks with greater dispersion in analyst forecasts, and Boehme et al. (2006) show that the valuation effects associated with differences of opinion are particularly pronounced when short selling is constrained. Berkman et al. (2009) further find that the low returns of high-dispersion stocks are realized disproportionately around earnings announcements. Johnson (2004) offers a different rational interpretation, treating forecast dispersion as a proxy for information risk rather than as differences of opinion.

Ang et al. (2006) also find that controlling for dispersion in analysts' earnings forecasts substantially attenuates the IVOL anomaly. We develop this connection into an explanation of the puzzle by deriving uncertainty and disagreement jointly from Bayesian learning and showing how their opposing effects on prices depend on short sale constraints and the breadth of ownership. The resulting framework explains both the variation in the anomaly across stocks and its concentration around earnings announcements. Establishing those implications under rational expectations requires a further step, because investors condition their valuations on a price that reflects information held by others, while their resulting demands determine that same price. Learning from price can compress the dispersion generated by private signals and, in the unrestricted Gaussian setting, eliminate its independent upward effect on valuation. We characterize the disagreement that

survives investors' learning from price and its implications for price when pessimistic positions are constrained.

We build on the noisy rational expectations tradition of Grossman (1976), Grossman and Stiglitz (1980), and Hellwig (1980), together with subsequent work on constrained trading and information aggregation (Diamond and Verrecchia, 1987; Bai et al., 2006; Yuan, 2006; Nezafat et al., 2017; Glebkin et al., 2021).

Related tractable models accommodate trading constraints by restricting the information available to different classes of investors or by introducing nontradable endowment risks correlated with an unlearnable component of the asset payoff (Nezafat et al., 2017; Glebkin et al., 2021). While these approaches have yielded useful and economically interpretable equilibria, we extend the noisy supply framework of Hellwig (1980) to short sale constraints without either restriction, allowing disagreement and the informational content of prices to be studied jointly as constraints vary. Nevertheless, we also investigate fixed supply and trading motivated by endowment risks in Appendix A.2.2, and show that the central predictions of the theory concerning ownership are robust to these alternative assumptions.

A close comparison is Albagli et al. (2024), who characterize information aggregation when investors face bounded positions, including a lower bound that rules out short sales. In their tractable constrained case, investors are risk-neutral and choose at the bounds of their feasible position set, so market clearing is governed by the valuation of the marginal investor. The present setting instead involves risk-averse investors who choose continuously varying positions above a finite short sale floor, requiring the entire distribution of unconstrained desired holdings to be aggregated rather than merely locating a cutoff investor. Breon-Drish (2015) obtains exact solutions beyond the jointly Gaussian setting by combining CARA preferences with conditional distributions of the payoff that belong to an exponential family. In his economy with a continuum of investors, desired positions are additively separable in private signals and a function of price. He uses market clearing to identify a linear combination of the payoff and supply that the equilibrium price must reveal. A floor on short positions breaks this separability for actual holdings, because the set of investors at the floor changes with price. The argument used to identify a linear statistic from market clearing therefore does not carry over directly. We retain the translation property of CARA utility and Gaussian private signals, but use the valuation of an investor who would optimally hold zero to construct the candidate price. Desired positions then depend on the difference between private signals and a public statistic inferred from price. The payoff cancels from these differences, allowing us to aggregate constrained demands exactly. The result is a rational expectations equilibrium in which

price may be nonlinear and posterior beliefs need not be Gaussian. The characterization allows a broad class of payoff priors and delivers comparative statics for the information contained in price, the breadth of ownership, and the response of expected price to fundamental risk, while recovering the unrestricted information structure of Hellwig (1980) as the position floor recedes.

Separating uncertainty within beliefs from disagreement across posterior means also connects the paper to a broad literature on expectations formation. Survey research has long distinguished the dispersion of point forecasts from uncertainty within an individual's probability distribution (Zarnowitz and Lambros, 1987; Patton and Timmermann, 2010; Lahiri and Sheng, 2010). Models of noisy information (Woodford, 2003), rational inattention (Sims, 2003), and sticky information (Mankiw and Reis, 2002) provide different accounts of how beliefs evolve as information arrives, while Coibion and Gorodnichenko (2012, 2015) use survey forecasts to test models of information rigidity and distinguish between competing explanations of expectations formation. The common prior bound developed here provides a new perspective on these models: regardless of the particular information structure, expected disagreement is restricted by the uncertainty remaining within individual posterior beliefs, with a sharper bound determined by prior payoff variance. The Gaussian benchmark then shows how fundamental risk and common information move the two moments in different proportions, providing additional predictions that can be confronted directly with survey data.

2 Uncertainty and Disagreement

We first consider an asset with terminal payoff X and a population of investors indexed by i. Each investor observes information $\mathcal{I}_i$ and forms a posterior distribution for the payoff. We write

$$m_i = \mathbb{E}[X \mid \mathcal{I}_i], \qquad v_i = \mathrm{Var}(X \mid \mathcal{I}_i)$$

for the posterior mean and variance. The variance v_i measures the uncertainty that remains in investor i's beliefs about the terminal payoff.

To measure disagreement, we compare posterior means across investors rather than the variation in those means across different realizations of the payoff. This distinction is important even when every investor receives a perfectly revealing signal. In that case all investors know the realized payoff, so neither uncertainty nor disagreement remains, although their common posterior mean $m_i = X$ is still random before the payoff is realized.

The unconditional variance of posterior means would consequently attribute disagreement to variation in the underlying state, even though investors agree in every state.

We therefore compare investors who face the same realized payoff, defining uncertainty and disagreement within each state by

$$U(x) = \mathbb{E}_i[v_i \mid X = x], \qquad D(x) = \mathrm{Var}_i(m_i \mid X = x). \tag{1}$$

This conditioning is a device for comparing investors, not an assumption that they knew the payoff when forming their beliefs. For simplicity, (1) displays only the conditioning on X, although other information common to investors is also held fixed, including the information conveyed by market price in the equilibrium developed below.

Figure 2 illustrates these definitions: wide posterior distributions can coexist with closely aligned posterior means, producing substantial uncertainty with little disagreement, whereas relatively precise posterior distributions may instead be centered far apart, producing substantial disagreement with little uncertainty.

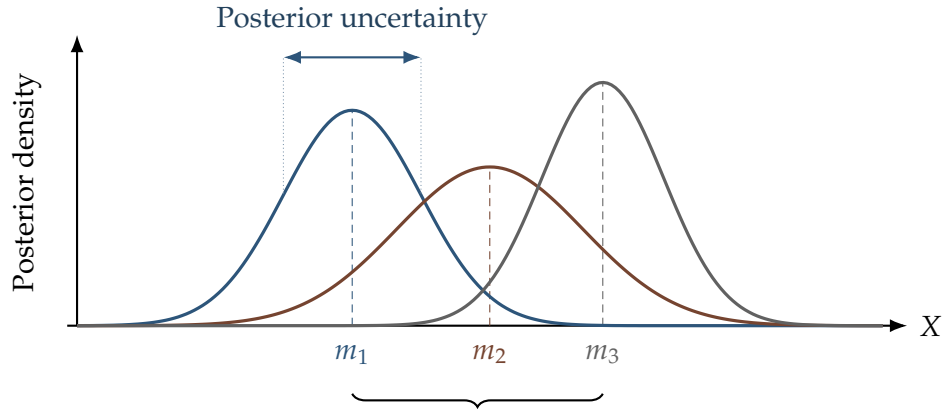


Figure 2: Posterior uncertainty and disagreement across posterior means. Each curve is an investor's posterior density for the same payoff. The width of a density captures that investor's uncertainty about the payoff, while differences in the centers of the densities capture disagreement across investors.

The restriction we derive concerns the corresponding ex ante quantities, obtained by averaging over the payoff and any other information common to investors:

$$U = \mathbb{E}_X[U(X)], \qquad D = \mathbb{E}_X[D(X)]. \tag{2}$$

Let $\bar{m}_c$ denote the average posterior mean across investors conditional on the realized common state. Under a common prior, writing $V = \mathrm{Var}(X)$ for prior payoff variance, the law of total variance gives the additive decomposition

$$V = \underbrace{U}_{\text{posterior uncertainty}} + \underbrace{D}_{\text{disagreement across beliefs}} + \underbrace{\mathrm{Var}(\bar{m}_c)}_{\text{common movement in posterior means}} . \quad (3)$$

To see the distinction in a setting without Gaussian assumptions, suppose that the payoff is the sum of two independent coin tosses, $X = c_1 + c_2$, where each c_k equals $+1$ or -1 with probability one half, so that $V = 2$. Each investor observes one of the two tosses, chosen independently and with equal probability, and learns nothing about the other. An investor who observes toss k therefore holds the posterior mean $m_i = c_k$ and the posterior variance $v_i = 1$, the variance of the unobserved toss. Uncertainty is thus the same in every state, $U(x) = 1$, but disagreement is not. When $X = \pm 2$, the two tosses agree, every investor holds the same posterior mean, and $D(x) = 0$. When $X = 0$, the tosses differ, so that half of the investors hold the posterior mean $+1$ and the other half -1, and $D(0) = 1$. Since $X = 0$ with probability one half, $D = 1/2$. The average posterior mean, $\bar{m}_c = X/2$, has variance $1/2$, so the decomposition in (3) reads $2 = 1 + \frac{1}{2} + \frac{1}{2}$. The unconditional variance of posterior means equals one, because it combines disagreement within each state with the common movement of beliefs across states, and conditioning on the realized state isolates disagreement from this common movement.

The fully revealing example places all prior payoff variance in the common movement of posterior means, with $U = D = 0$ even though those means vary with the state. Disagreement is nevertheless sometimes used empirically as a proxy for uncertainty, and Lahiri and Sheng (2010) examine when that practice is informative. The variance decomposition alone provides no reason for treating the two as equivalent. Understanding how they are related requires the additional restrictions imposed by rational learning from a common prior.

2.1 Survey Evidence

The Survey of Professional Forecasters (SPF) permits uncertainty and disagreement to be measured for the same outcome and horizon because respondents report both point forecasts and subjective probability distributions. The forecasts considered here concern a fixed annual outcome, such as inflation or unemployment over a given calendar year, which respondents predict in each of the four survey rounds. The point forecast is treated as the respondent's posterior mean, and posterior variance is estimated from the accompanying

probability histogram. Disagreement is the cross-sectional variance of point forecasts, while uncertainty is the average fitted posterior variance. Appendix B describes the respondent sample and measurement procedure, and Appendix Tables A1 and A2 report the moments by variable and by survey quarter.

Figure 3 documents two patterns that the theory helps explain. First, disagreement is generally much smaller than uncertainty: median D/U is 0.232 for core CPI, 0.218 for core PCE, and 0.061 for unemployment, and the ratio exceeds one in only four variable-quarter observations. Second, disagreement falls sharply relative to uncertainty as the fixed event approaches: pooling the three series, mean D/U declines from 0.412 in the first quarter to 0.069 in the fourth, even though forecasters continue to predict the same annual outcome.

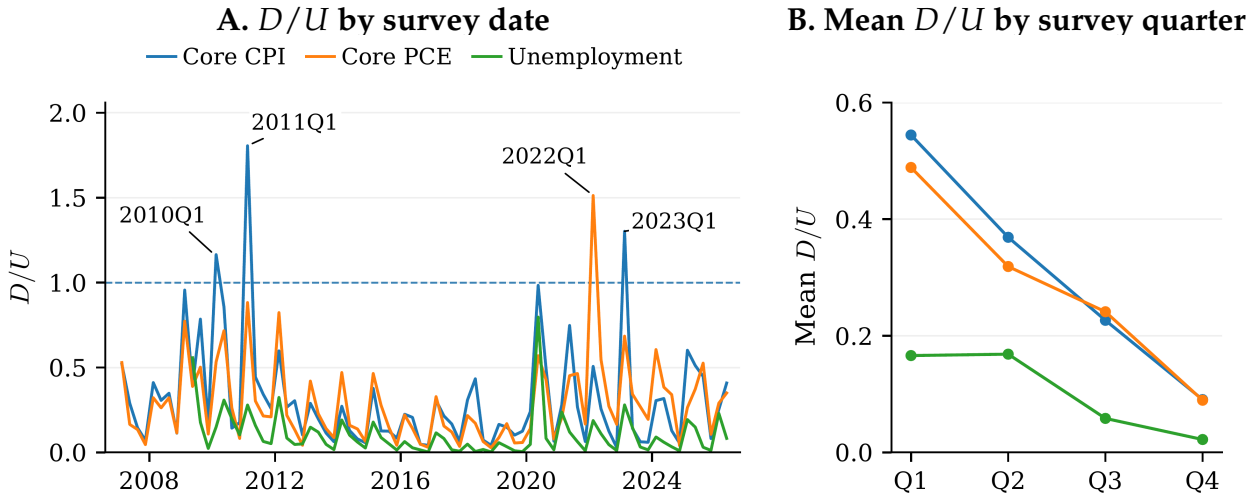


Figure 3: Disagreement relative to uncertainty in the Survey of Professional Forecasters. Panel A plots D/U by survey date for core CPI, core PCE, and unemployment. Within a calendar year, the first through fourth survey rounds all concern the same annual outcome. Panel B averages these observations by survey quarter separately for the three variables, showing how disagreement relative to uncertainty changes as information accumulates before that fixed outcome is realized.

These comparisons invite a more precise question: if forecasters begin from a common prior and update rationally, how much disagreement can remain for a given level of posterior uncertainty? We first derive a restriction that does not depend on how their private information is generated, before examining the different responses of uncertainty and disagreement as information accumulates.

2.2 A Bound on Disagreement

When investors share a common prior and update rationally, the dispersion of their posterior means cannot be chosen independently of the uncertainty remaining within their beliefs. The next result makes this restriction precise without specifying how their information is distributed.

Theorem 1 (Common prior projection bound)**.** *Let $X \in L^2$ with $V > 0$, and suppose investors share a common prior and form rational conditional expectations. For the ex ante quantities U and D in* (2)*, with statewise disagreement conditioned on the realized payoff and any other information common to investors,*

$$D \leq U\left(1 - \frac{U}{V}\right) \leq U. \tag{4}$$

Equivalently, whenever $U > 0$,

$$\frac{D}{U} \leq 1 - \frac{U}{V}. \tag{5}$$

Proof. See Appendix A.1.1. □

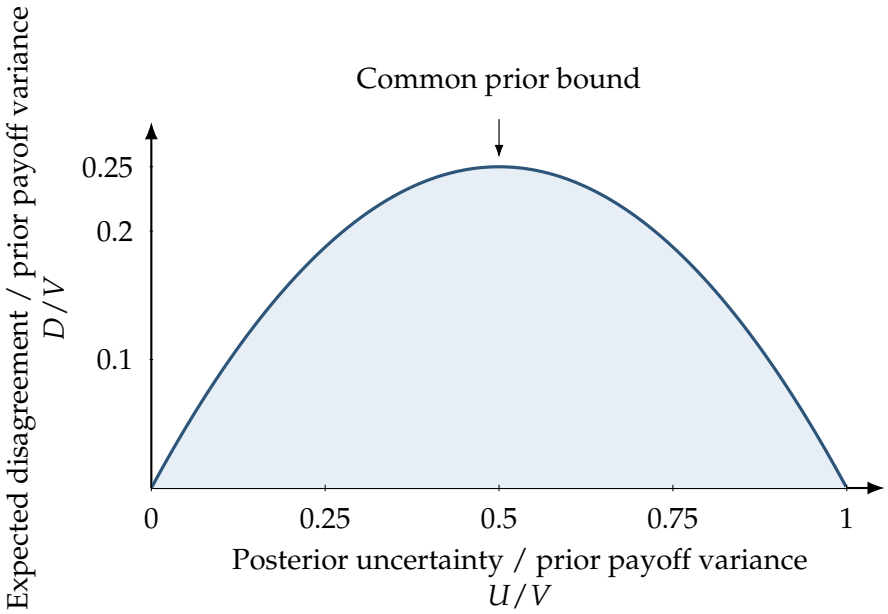


Figure 4: Ex ante posterior uncertainty and expected disagreement under a common prior. For each value of U/V, the curve gives the largest expected disagreement consistent with the projection bound, $D/V \leq (U/V)(1 - U/V)$. With a Gaussian payoff prior and private Gaussian signals of equal variance, changing signal precision traces the curve. Very precise signals leave little uncertainty or disagreement, while nearly uninformative signals leave high uncertainty but little disagreement because posterior means remain close to the common prior. Expected disagreement is largest at intermediate precision.

The argument follows from a familiar property of conditional expectations. For a randomly drawn investor with posterior mean m, rational learning implies $\text{Var}(m) = \text{Cov}(m, X) = V - U$: the variation in posterior means is tied to the amount of prior payoff variance resolved by the investor's information. Projecting m on the realized payoff removes the component of beliefs explained by the underlying state, while conditioning on all commonly observed information can remove still more. The variance left after the first projection is exactly $U(1 - U/V)$, giving the sharper inequality in (4). In the two-coin example, that inequality holds with equality, since $D = U(1 - U/V) = 1/2$.

Equation (5) also restricts the circumstances in which disagreement can serve as an empirical proxy for uncertainty. When private information is almost uninformative, individual posteriors remain diffuse but their means stay close to the common prior, leaving high uncertainty and little disagreement. At the opposite extreme, nearly revealing private information leaves little uncertainty and little disagreement because investors learn almost the same payoff. The greatest scope for disagreement lies between these extremes, when private information moves posterior means materially away from the prior but remains noisy enough for those means to differ across investors.

Equality in the first inequality obtains when the realized payoff is the only common conditioning variable and, after centering, the common component of posterior means is exhausted by the linear projection on X:

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

A Gaussian payoff prior with private Gaussian signals of equal variance satisfies this condition, and varying their precision while holding prior payoff variance fixed traces out the entire boundary in Figure 4, with expected disagreement largest at an intermediate signal precision.

The predominance of $D/U < 1$ in Figure 3 is consistent with the common prior restriction, although the theorem concerns expectations across states rather than every realized cross section. If investors differ in the precision or timing of their information, disagreement can exceed uncertainty in a particular state without violating the ex ante bound. We next examine how this distinction appears in noisy information and sticky information models, and how the accumulation of common information bears on the decline in D/U in the SPF.

2.3 Ex Ante and Statewise Disagreement

Consider first a Gaussian noisy information model in which investors observe imperfect signals about a common state (Sims, 2003; Woodford, 2003). Let

$$X \sim N(0, a), \qquad s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, b),$$

with signal errors independent across investors and of the payoff. Because all investors receive signals of the same precision, posterior uncertainty and the conditional dispersion of posterior means are constant across payoff states, so the statewise and ex ante quantities coincide. Bayesian updating gives

$$U = \frac{ab}{a+b}, \qquad D = \frac{a^2 b}{(a+b)^2}, \qquad \frac{D}{U} = \frac{a}{a+b}.$$

Because $V = a$, the common prior bound holds with equality in this economy. Holding signal noise b fixed, greater fundamental variance raises posterior uncertainty and disagreement and increases their ratio, D/U. This comparison changes the payoff variation investors are trying to learn about while keeping the precision of their private signals unchanged.

The same Gaussian model also allows us to examine what happens when investors receive more information in common, rather than facing a change in fundamental risk. Let τ_c denote the precision of the common posterior before private signals arrive, incorporating the prior and any information observed by all investors, and let τ_x denote the precision of the independent private signals. Then

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

Holding private precision fixed, an increase in common precision satisfies

$$\frac{\partial \log U}{\partial \tau_c} = -\frac{1}{\tau_c + \tau_x}, \qquad \frac{\partial \log D}{\partial \tau_c} = -\frac{2}{\tau_c + \tau_x}, \qquad \frac{\partial (D/U)}{\partial \tau_c} = -\frac{\tau_x}{(\tau_c + \tau_x)^2} < 0.$$

Additional common precision reduces uncertainty while drawing posterior means together at twice the proportional rate. This gives an interpretation of the fixed event pattern in Figure 3: the annual target remains unchanged, but the information available in common becomes more extensive as the year progresses, lowering D/U in the Gaussian model.

A different source of heterogeneity arises in models of sticky information, where agents differ in when they last updated rather than in the noise of a contemporaneous signal

(Mankiw and Reis, 2002). A stark one-period version makes the contrast between statewise and ex ante disagreement explicit. Normalizing the prior mean of X to zero, suppose that a fraction $1-\lambda$ of investors receives a perfectly informative signal about X, while a fraction λ receives no new information. Average posterior uncertainty is

$$U^S = \lambda V,$$

whereas disagreement conditional on the realized payoff $X = x$ is

$$D^S(x) = \lambda(1-\lambda)x^2.$$

Statewise disagreement therefore exceeds average posterior uncertainty whenever $x^2 > V/(1-\lambda)$. Ex ante, however,

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

The occasional observations with $D/U > 1$ in Figure 3 are therefore entirely compatible with common prior learning: large realized states can separate investors who have updated from those whose information is stale even though the ex ante restriction remains intact. Appendix A.2.1 develops dynamic versions of both the noisy information and sticky information models, allowing private information to accumulate over time and investors to update at different dates.

Disagreement can also reflect differences of opinion beyond those generated by heterogeneous information under a common prior. To represent such differences, add an idiosyncratic component η_i to investor i's valuation, with cross-sectional mean zero, variance σ_η^2, and zero covariance with the rational posterior mean m_i, conditional on the common state. If this component shifts the valuation center without changing posterior uncertainty, then

$$\widetilde{m}_i = m_i + \eta_i, \qquad \widetilde{U} = U, \qquad \widetilde{D} = D + \sigma_\eta^2.$$

The additional heterogeneity raises disagreement without changing uncertainty. For beliefs aligned on the same payoff, horizon, and probability measure, an ex ante violation of the common prior bound therefore rules out Bayesian learning from a common prior as a complete account of those beliefs.

2.4 Beliefs about Equity Returns

The SPF permits the most direct comparison because the same respondents report uncertainty and disagreement for the same outcome and horizon. Since these forecasts concern macroeconomic outcomes rather than tradable asset payoffs, we also examine forecasts of stock returns before turning to the equilibrium determination of prices.

Both the Livingston Survey, which collects forecasts from professional economists, and the Duke CFO Survey, which elicits expectations from corporate financial officers, contain forecasts of S&P 500 returns over the following year. The forecasts used here do not come with probability distributions comparable to those of the SPF, so we match disagreement in these return forecasts to option-implied variance, the natural market-based counterpart to uncertainty available at the same horizon. The comparison is not exact: survey forecasts represent respondents' subjective beliefs, whereas option-implied variance is defined under the risk-neutral distribution. Appendix B.2 gives the survey definitions, matching procedure, and construction of the option measures.

The resulting ratios are small in both surveys: the mean ratio of forecast disagreement to option-implied variance is 0.096 in the Livingston Survey and 0.057 in the Duke CFO Survey, and in all 74 matched observations, forecast disagreement lies below the corresponding variance inferred from options. This pattern is consistent with the effect of common information derived above. Investors forecasting returns on an actively traded stock index observe a market price that reflects information dispersed across investors, and the Gaussian model predicts that such information reduces disagreement relative to uncertainty. Because survey forecasts and option-implied variance refer to different probability measures, however, these comparisons are suggestive rather than a direct test of the bound.

In an asset market, the common information available to investors depends on the price generated by their own trading decisions. The extent of learning from that price, and the disagreement that remains when investors face trading constraints, must therefore be determined jointly with market clearing.

3 Information Aggregation and Learning from Prices

Having examined how uncertainty and disagreement arise under rational learning, we turn to the market in which investors act on those beliefs. We first consider a benchmark in which investors form beliefs from their private information but do not draw further inferences from price. This allows us to isolate how limits on short positions alter the

effects of uncertainty and disagreement on valuation. We then let investors learn from the very price that their demands determine, so that the remaining disagreement and its consequences for valuation are established in equilibrium rather than taken as given.

The focus on a single risky claim is also natural in a factor model of asset payoffs. Once exposures to common factors have been hedged, portfolio choice under CARA preferences separates across residual claims whose payoffs are independent conditional on investors' information. The model can therefore be interpreted as isolating how dispersed information and short-sale constraints affect the valuation of firm-specific risk.

A unit mass of investors trades a risky claim at date 0 that pays X at date 1, with the gross risk-free return normalized to one. Fix the realized payoff and any information common to investors, as in Section 2. Let m_i denote investor i's posterior mean. We write $\bar{m}$ and D for the mean and variance of these posterior means across investors, and U for the posterior variance common to all investors. Under conditionally Gaussian payoffs and CARA utility with coefficient $\gamma > 0$, the unconstrained desired holding is

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

We impose a common lower bound $q_i \geq -\ell$, where $\ell \geq 0$, so that actual holdings are given by

$$q_i = \max\left\{\frac{m_i - p}{\gamma U}, -\ell\right\}.$$

The unrestricted economy is recovered as $\ell \to \infty$, while $\ell = 0$ rules out short positions entirely.

3.1 Trading Constraints and Market Clearing

With unrestricted positions and fixed outstanding supply $S > 0$, market clearing gives

$$p = \bar{m} - \gamma U S.$$

Disagreement D affects the distribution of positions across investors but not aggregate demand, so a mean-preserving widening of posterior means changes the allocation without changing price. Greater posterior uncertainty, by contrast, reduces investors' willingness to absorb the fixed supply and lowers the clearing price through the risk compensation term $\gamma U S$.

With a lower bound on positions, optimistic and pessimistic investors no longer adjust their holdings symmetrically. In particular, if A_i denotes investor i's desired position at a given candidate price before the constraint is imposed, her actual holding is

$$A_i \longmapsto \max\{A_i, -\ell\}.$$

The lower bound preserves increases in the desired positions of optimists while limiting reductions among sufficiently pessimistic investors. The resulting convexity explains why more dispersed desired positions can increase aggregate demand even when their cross-sectional mean is unchanged.

Lemma 1 (Dispersion and constrained demand)**.** *At a fixed candidate price and common shorting limit, a mean-preserving spread in unconstrained desired positions weakly raises aggregate demand. The increase is strict whenever* $\int_{-\infty}^{-\ell}\{F_2(z) - F_1(z)\}\,dz > 0$*, where* F_1 *and* F_2 *are the distributions of desired positions before and after the spread.*

Proof. See Appendix A.1.3. □

The lemma requires neither Gaussian beliefs nor CARA preferences, since its conclusion follows from the convexity of actual holdings as a function of unconstrained desired positions. It also applies when investors learn from price, provided the comparison holds that inference rule fixed. A mean-preserving spread weakly raises aggregate demand at the initial price, with a strict increase when the spread increases expected holdings after truncation. If the aggregate demand schedule remains decreasing in price, such an increase requires a higher market-clearing price.

To obtain explicit comparisons for the clearing price, we now impose more structure on the distribution of posterior valuations. Throughout the remainder of this subsection, investors form their beliefs from private information and take those beliefs as given when the market clears, without yet drawing additional inferences about the payoff from the contemporaneous price. This distinction permits us to examine the effect of constrained demand before turning, in Section 3.2, to the rational expectations equilibrium in which price itself conveys information.

Set $\ell = 0$ and suppose that, within the realized state, posterior means are normally distributed across investors:

$$m_i \sim N(\bar{m}, D).$$

Define

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

and let

$$d = \frac{\bar{m} - p}{\sqrt{D}}.$$

The variable d is the standardized ownership cutoff: it measures the distance between the mean posterior valuation and the market price in units of cross-sectional dispersion. Since an investor holds a positive position precisely when $m_i > p$, the fraction of investors who hold the asset is

$$\Omega = \Phi(d).$$

Thus, $d < 0$ means that the price lies above the center of the valuation distribution and fewer than half of investors hold the asset, whereas $d > 0$ corresponds to ownership by more than half of investors.

Because $h(d) = \mathbb{E}[(d + Z)^+]$ for $Z \sim N(0, 1)$, market clearing becomes

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

Implicit differentiation gives

$$\left.\frac{\partial p}{\partial U}\right|_{D,\bar{m},S} = -\frac{\gamma S}{\Phi(d)} < 0, \qquad \left.\frac{\partial p}{\partial D}\right|_{U,\bar{m},S} = \frac{\phi(d)}{2\sqrt{D}\,\Phi(d)} > 0. \tag{6}$$

Greater uncertainty lowers the clearing price because investors require compensation for bearing more posterior risk, whereas greater disagreement raises it because the short sale constraint prevents negative positions from expanding symmetrically with positive ones.

The same cutoff also determines who holds the asset: among investors with positive positions,

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

When $\Omega < 1/2$, which we call *narrow ownership*, the price lies above the mean and median posterior valuation, and the asset is held by investors drawn disproportionately from the upper part of the valuation distribution. When $\Omega > 1/2$, the case of *broad ownership*, investors with progressively less optimistic valuations are represented among the holders of the asset.

The derivatives in (6) show what happens when uncertainty or disagreement changes on its own, but an increase in fundamental volatility will generally affect both. To examine

this joint response, suppose that investors share a Gaussian prior for the payoff and receive independent private signals, with

$$X \sim N(\mu_0, a), \qquad s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, b).$$

Bayesian updating gives

$$U = \frac{ab}{a+b}, \qquad D = \frac{a^2 b}{(a+b)^2}, \qquad \frac{\sqrt{D}}{U} = \frac{1}{\sqrt{b}}.$$

Holding signal noise b fixed, an increase in fundamental variance raises both posterior uncertainty and disagreement, with disagreement rising proportionately more, while the ratio $\sqrt{D}/U$ remains constant. Consequently, the standardized ownership cutoff d, defined above, is determined independently of a by

$$h(d) = \gamma S \sqrt{b}, \qquad \mathbb{E}[p] = \mu_0 - \frac{d}{\sqrt{b}} U. \tag{7}$$

Proposition 1 (Ownership threshold)**.** *Holding (b, S, γ, μ_0) fixed, a mean-preserving increase in fundamental variance raises expected price if and only if*

$$d < 0 \quad \Longleftrightarrow \quad \Omega < \frac{1}{2} \quad \Longleftrightarrow \quad S < \frac{\phi(0)}{\gamma \sqrt{b}}.$$

Expected price falls when the inequalities are reversed.

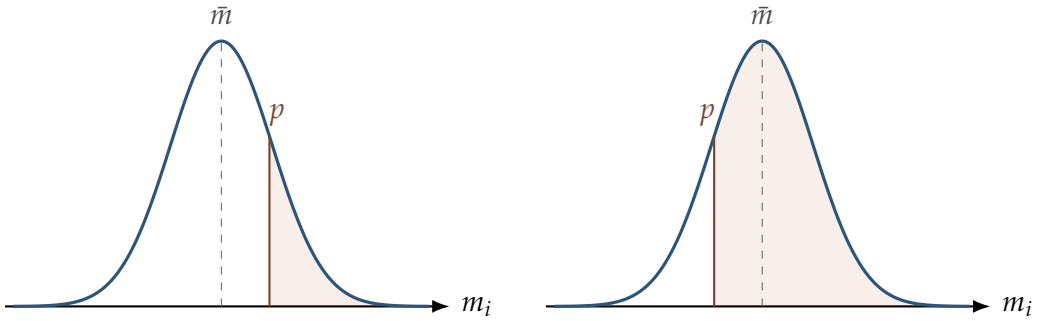
Proof. The result follows directly from (7), and Appendix A.1.3 establishes the corresponding result beyond the Gaussian case. □

The threshold at one half has a direct interpretation in terms of the investors who hold the asset. With narrow ownership, the clearing price lies above the center of the distribution of posterior valuations, so the effect of their greater dispersion outweighs the additional compensation required for bearing risk. When ownership is broad, the clearing price lies below that center and the conventional effect of uncertainty dominates. The equivalent condition on S also makes clear why supply is relevant: as more units must be absorbed, the asset comes to be held by a broader set of investors whose valuations extend progressively farther down the distribution.

The one half boundary does not depend on the particular shape of the Gaussian distribution. Suppose more generally that posterior valuations can be written as

$$m_i = \mu + cUZ_i,$$

where $c > 0$ and Z_i has a fixed continuous distribution with mean zero, unit variance, and unique median zero, so that $c = \sqrt{D}/U$. Market clearing again fixes the standardized cutoff independently of U, and an increase in U raises the price exactly when fewer than half of investors hold positive positions, as shown formally in Appendix A.1.3. Figure 5 illustrates the same comparison: under narrow ownership, the price cutoff lies above the center of the valuation distribution, while under broad ownership, it lies below it.



$\Omega < 1/2$
Narrow ownership

$\Omega > 1/2$
Broad ownership

Figure 5: Ownership and the price cutoff. The shaded portion of each valuation distribution consists of investors with positive holdings. Under narrow ownership, the price cutoff lies above the mean and median valuation, while under broad ownership, it lies below them.

The preceding results establish how constrained positions alter the clearing price for a given distribution of beliefs. In a rational expectations equilibrium, however, investors must also take account of what that price reveals about the information received by others, and we must determine whether the same effects persist once learning and market clearing are made mutually consistent.

3.2 Learning from Equilibrium Prices with Constrained Positions

Once investors learn from market prices, they base their decisions partly on information revealed by the demands of other investors. Their own demands, in turn, help determine the price from which everyone learns. With short sale constraints, this interaction is especially difficult because the identity of investors whose positions are truncated changes with price, so aggregate demand is generally nonlinear even when their private signals are Gaussian.

To study this interaction, we retain the heterogeneous private signals of the preceding model but replace the fixed outstanding supply S with an unobserved random supply T. Supply uncertainty prevents the market price from revealing the payoff perfectly, as in the

noisy rational expectations economy of Hellwig (1980). Unlike the unrestricted Gaussian benchmark, however, the constrained economy cannot be solved by assuming an affine price or Gaussian posterior beliefs, so investors' desired positions must be derived from the posterior distributions implied by the equilibrium price.

Before private signals are observed, the payoff has an atomless, nondegenerate prior with finite second moment. Supply is drawn independently of the payoff and of private signal errors according to

$$T \sim \text{Unif}[\underline{t}, \bar{t}], \qquad 0 < \underline{t} < \bar{t}.$$

Investor i then observes

$$s_i = X + \epsilon_i,$$

where the ϵ_i are independent Gaussian errors with variance b, together with the equilibrium market price, but investors do not observe T. Preferences remain CARA with coefficient γ, and positions satisfy $q_i \geq -\ell$. The distribution of T is held fixed when the shorting limit or the payoff distribution is later varied.

The lower bound on supply is necessary for an equilibrium with constrained positions. Since every investor must hold at least $-\ell$, aggregate demand cannot fall below $-\ell$ at any price, and realizations of supply below that level cannot be cleared. This rules out a Gaussian supply distribution under any finite position floor. We use a uniform distribution for the main construction, while Appendix A.1.4 extends the result to a broad class of distributions bounded below by $-\ell$, without requiring an upper bound.

Our approach is to determine the information that market clearing can reveal before calculating the equilibrium price associated with that information. We begin with a candidate public statistic of the form

$$Y_\ell = X + w_\ell(T),$$

where market clearing will determine the function w_ℓ. If equilibrium price is strictly increasing in Y_ℓ, observing the price conveys precisely the same information as observing this statistic. We therefore solve the portfolio problem conditional on $Y_\ell = y$, use market clearing to determine w_ℓ, and verify afterward that the resulting Bayesian price is indeed strictly increasing.

Conditional on $Y_\ell = y$, let ν_y denote the public posterior and write $M_y(s) = \mathbb{E}[X \mid s_i = s, Y_\ell = y]$ for the posterior mean after observing private signal s. Independence of supply and private signal errors implies that the investor's full posterior is proportional to

$$\exp\left[-\frac{(s-x)^2}{2b}\right]\nu_y(dx).$$

For an unconstrained position q at price p, the CARA first order condition is

$$p = \frac{\mathbb{E}[Xe^{-\gamma qX} \mid s, y]}{\mathbb{E}[e^{-\gamma qX} \mid s, y]}.$$

The right-hand side is the payoff mean under the investor's posterior after states have been weighted by marginal utility. CARA utility and the Gaussian private-signal likelihood give the identity

$$e^{-\gamma qx}e^{-(s-x)^2/(2b)} \propto \exp\left[-\frac{(s-\gamma bq-x)^2}{2b}\right].$$

After normalization, marginal-utility weighting is equivalent to evaluating the posterior mean at the shifted private signal $s - \gamma bq$, so the first-order condition becomes

$$p = M_y(s - \gamma bq).$$

The shift changes the weights entering marginal valuation, not the investor's actual information or beliefs, and requires no Gaussian assumption on the payoff prior. Moreover, the Gaussian likelihood implies

$$\frac{\partial M_y(s)}{\partial s} = \frac{\mathrm{Var}(X \mid s_i = s, Y_\ell = y)}{b} > 0,$$

where strictness follows from the nondegenerate posterior. Thus, a higher private signal raises the posterior mean even when that mean is not an affine function of the signal.

We choose the candidate price to equal the valuation of the investor whose private signal coincides with the public statistic:

$$p_\ell(y) = M_y(y) = \mathbb{E}[X \mid s_i = y, Y_\ell = y].$$

At this price, that investor's unconstrained optimal holding is zero. For any other investor, strict monotonicity of M_y implies that marginal valuation equals the same price precisely when

$$s_i - \gamma bq = y.$$

Hence, the exact unconstrained desired position is

$$q_i^* = \frac{s_i - y}{\gamma b}.$$

Strict concavity makes this position unique, and the short sale constraint simply truncates it at $-\ell$, with the full optimality argument given in Appendix A.1.4.

The investor whose desired position is zero provides a useful reference because she bears no payoff risk, and her marginal valuation therefore coincides with her ordinary posterior mean. By choosing the price to equal that valuation when her private signal matches the public statistic, we identify the price at which she is content not to trade. Other investors choose their positions by balancing the information in their private signals against the risk of their holdings: with CARA utility, each additional unit held reduces marginal valuation in exactly the same way as lowering the private signal by γb. Their desired positions consequently depend on the difference between their private signals and the public statistic inferred from price, rather than on the level of the fundamental payoff. Investors condition on both private information and price throughout this calculation. The cancellation of the payoff allows us to aggregate these positions exactly after imposing the short sale floor.

Market clearing delivers the second part of the characterization. Since

$$s_i = X + \epsilon_i \qquad \text{and} \qquad Y_\ell = X + w_\ell(T),$$

we have

$$s_i - Y_\ell = \epsilon_i - w_\ell(T). \tag{8}$$

The fundamental therefore drops out of the cross section of desired positions conditional on supply. Because the remaining heterogeneity is Gaussian, aggregate demand can be integrated exactly even after the floor is imposed. Using $h(d) = \mathbb{E}[(d + Z)^+]$, market clearing at $T = t$ becomes

$$t + \ell = \frac{1}{\gamma\sqrt{b}} h\left(\frac{-w_\ell(t) + \gamma b \ell}{\sqrt{b}}\right).$$

Because h is strictly increasing with range $(0, \infty)$, each feasible value of supply t determines a unique $w_\ell(t)$, the difference between the public statistic and the payoff. Define

$$d_\ell(t) = h^{-1}\left(\gamma\sqrt{b}(t + \ell)\right), \qquad w_\ell(t) = \gamma b \ell - \sqrt{b}\, d_\ell(t). \tag{9}$$

Proposition 2 (Equilibrium with constrained positions)**.** *An equilibrium is*

$$Y_\ell = X + w_\ell(T), \qquad P_\ell = p_\ell(Y_\ell), \qquad p_\ell(y) = \mathbb{E}[X \mid Y_\ell = y, s_i = y],$$
$$q_i = \max\left\{\frac{s_i - Y_\ell}{\gamma b}, -\ell\right\}. \tag{10}$$

The price function is continuous and strictly increasing on its feasible signal set, and the positions in (10) *clear supply in every state.*

Proof. See Appendix A.1.4. □

The clearing calculation determines Y_ℓ without first solving for the price as a function of that statistic. We must still verify, however, that the price investors actually observe conveys the information on which their portfolio decisions were based. Appendix A.1.4 establishes that p_ℓ is continuous and strictly increasing, using log concavity of the effective likelihood for the reference investor. The price and the public statistic therefore reveal precisely the same information, completing the verification that the positions in Proposition 2 are optimal and clear the market.

Equation (8) has a further implication for comparative statics: holding the supply law, private signal noise, risk aversion, and shorting limit fixed leaves the distribution of positions independent of the payoff distribution. Fundamental risk can therefore vary without altering the allocation or the fraction of investors holding the asset, although prices may change as investors infer and value the payoff.

The shorting limit also governs the informativeness of equilibrium prices, reflecting the interaction between short sale constraints and information aggregation studied by Bai et al. (2006). Let

$$\Omega_\ell^{\text{int}}(t) = \Phi(d_\ell(t))$$

denote the fraction of investors whose positions lie strictly above the floor, that is, the investors whose holdings can continue to respond to their private signals. When $\ell = 0$, this fraction coincides with the fraction of investors holding the asset, whereas when $\ell > 0$, it also includes investors with interior short positions. Differentiating the supply map gives

$$-w_\ell'(t) = \frac{\gamma b}{\Omega_\ell^{\text{int}}(t)}. \tag{11}$$

Equation (11) relates the supply response of the public statistic to the fraction of investors free to adjust their positions. When that fraction is small, a given supply disturbance requires a larger movement in the statistic, so supply noise is more prominent relative

to the fundamental. Relaxing the constraint expands the set of responsive positions and reduces this sensitivity to supply.

With uniform supply, this comparison can be stated directly in terms of posterior uncertainty. If $\ell_2 > \ell_1$, then

$$\mathbb{E}[\mathrm{Var}(X \mid P_{\ell_2})] \leq \mathbb{E}[\mathrm{Var}(X \mid P_{\ell_1})], \qquad \mathbb{E}[\mathrm{Var}(X \mid s_i, P_{\ell_2})] \leq \mathbb{E}[\mathrm{Var}(X \mid s_i, P_{\ell_1})].$$

The first inequality concerns information recovered from equilibrium price alone, while the second includes the investor's private signal, and Appendix A.1.5 proves both.

For any fixed pair of shorting limits, sufficiently small uniform supply disturbances also imply lower D and a lower ratio D/U under the looser constraint, provided the payoff density is positive and continuous almost everywhere. Disagreement then contracts proportionately more than uncertainty, as in the common-information comparison in Section 2, but now because constrained market clearing changes what investors learn from price. Appendix A.1.6 derives the approximation and states its regularity conditions.

The solution also permits a precise comparison with the unrestricted rational expectations economy of Hellwig (1980). First, if the supply distribution is held fixed while the short sale floor recedes, the nonlinear adjustment converges to the linear expression $w_\ell(t) \to -\gamma b t$, so

$$Y_\ell = X + w_\ell(T) \longrightarrow Y_\infty = X - \gamma b T.$$

In this limit, private information remains dispersed, but optimistic and pessimistic positions can adjust without truncation, and the public statistic has the familiar linear form without any need to assume Gaussian supply.

Recovering the Gaussian benchmark itself requires a further comparison. A Gaussian supply law assigns positive probability to values that no economy with a finite position floor can clear, so it cannot simply be imposed on each constrained economy. Appendix A.1.7 instead constructs feasible supply distributions that approach a Gaussian law as the floor recedes, while the allocations and the information revealed by price converge to those of the unrestricted economy. When the payoff prior is also Gaussian, Bayesian valuation yields the affine price of Hellwig (1980), in which greater fundamental variance lowers expected price for positive mean supply. Thus, the two limits establish both the convergence of the information structure with a fixed supply law and the recovery of the classical Gaussian equilibrium from a sequence of feasible constrained economies.

3.3 Fundamental Risk and Ownership

Section 3.1 establishes the ownership threshold when investors do not learn from price. We now examine whether it survives when they condition on the information conveyed by the equilibrium price. Equation (8) allows the payoff distribution to vary while the supply law, private signal noise, preferences, and shorting limit remain fixed, so changes in expected price reflect valuation and inference rather than changes in equilibrium holdings.

Return first to the economy with a fixed uniform supply law and finite shorting capacity. At supply t, the fraction of investors with positive positions is

$$\Omega_\ell^+(t) = \Phi\left(-\frac{w_\ell(t)}{\sqrt{b}}\right).$$

This is the ownership measure relevant for the price comparison, rather than the interior fraction $\Omega_\ell^{\text{int}}(t)$ introduced above, because with a finite shorting allowance, some investors may hold negative positions without being at the floor. We continue to call ownership *narrow* when $\Omega_\ell^+(t) < 1/2$ and *broad* when $\Omega_\ell^+(t) > 1/2$, and the fraction of investors holding the asset crosses one half when supply equals

$$t^*(\ell) = \frac{h(-\gamma\sqrt{b}\,\ell)}{\gamma\sqrt{b}}.$$

The breadth of ownership is determined in equilibrium rather than imposed as an assumption. Because the floor lifts each position from $q_i^* = (s_i - Y_\ell)/(\gamma b)$ to $q_i^* + (-\ell - q_i^*)^+$, and unconstrained desired positions average to $(X - Y_\ell)/(\gamma b) = -w_\ell(T)/(\gamma b)$, supply is absorbed only if

$$w_\ell(t) = \gamma b\Big(\mathbb{E}_i\big[(-\ell - q_i^*)^+\big] - t\Big). \tag{12}$$

The public statistic therefore lies above the payoff when the total quantity of short positions excluded by the floor exceeds supply, and below it when that quantity is smaller than supply. Ownership is narrow exactly in the first case, in which the asset trades as though it were in negative effective net supply. Solving the clearing condition for the public statistic, the fraction of investors holding the asset becomes

$$\Omega_\ell^+(t) = \Phi\big(d_\ell(t) - \gamma\sqrt{b}\,\ell\big),$$

so ownership depends on the supply realization, the shorting limit, and the product $\gamma\sqrt{b}$, but not on the payoff distribution. Since $h' = \Phi$, differentiating $d_\ell(t) - \gamma\sqrt{b}\,\ell$ with respect to t, ℓ, and $\gamma\sqrt{b}$, respectively, gives

$$\frac{\gamma\sqrt{b}}{\Phi(d_\ell(t))}, \qquad \frac{\gamma\sqrt{b}\,[1-\Phi(d_\ell(t))]}{\Phi(d_\ell(t))}, \qquad \frac{t+\ell\,[1-\Phi(d_\ell(t))]}{\Phi(d_\ell(t))}$$

and each derivative is positive when supply is positive. These derivatives establish that ownership broadens with outstanding supply, since additional units must be absorbed by investors farther down the distribution of valuations. Ownership also broadens when the shorting limit is relaxed, because the larger short positions of pessimistic investors must be matched by holdings on the long side, and when risk aversion or private signal noise rises, because each investor is willing to hold less. Correspondingly, the threshold $t^*(\ell)$ falls from $\phi(0)/(\gamma\sqrt{b})$ at $\ell = 0$, the threshold in the economy with fixed beliefs, toward zero as the floor recedes.

To isolate the valuation effect of additional fundamental risk, start from a payoff X_1 and add an independent mean-zero risk Z, producing $X_2 = X_1 + Z$. The information technology, private signal variance, preferences, shorting limit, and supply distribution are held fixed throughout, and we compare expected equilibrium prices over payoff and supply realizations.

Proposition 3 (Additional fundamental risk)**.** *Suppose X_1 has a log concave density and finite second moment. Let $X_2 = X_1 + Z$, where Z is independent of X_1, supply, and private signal errors, has mean zero, and has finite variance. Holding the supply distribution, private signal variance, preferences, and shorting limit fixed,*

$$\begin{aligned} \Omega_\ell^+(t) \le \tfrac{1}{2} \text{ for every } t \quad &\Longrightarrow \quad \mathbb{E}[P_2] \ge \mathbb{E}[P_1], \\ \Omega_\ell^+(t) \ge \tfrac{1}{2} \text{ for every } t \quad &\Longrightarrow \quad \mathbb{E}[P_2] \le \mathbb{E}[P_1]. \end{aligned}$$

Proof. See Appendix A.1.9. □

Proposition 3 shows that the ownership result survives endogenous learning from equilibrium prices. If ownership remains narrow for every supply realization, adding independent fundamental risk weakly raises expected price, whereas if it remains broad for every supply realization, the price weakly falls. In terms of the support of the supply distribution, the first condition is $\bar{t} \le t^*(\ell)$, and the second is $\underline{t} \ge t^*(\ell)$.

Relaxing the shorting limit moves the economy toward the conventional price effect of fundamental risk through several related changes. A looser limit permits more of the

short demand of pessimistic investors to be expressed, raising effective supply in (12) and lowering the valuation required to clear the market for given beliefs. It also broadens ownership at every supply realization and lowers the threshold $t^*(\ell)$, narrowing the range of supply over which greater fundamental volatility raises expected price. At the same time, more investors can adjust their positions in response to private information, making price more informative and reducing posterior uncertainty and, for small supply disturbances, disagreement proportionately more. As the floor recedes, $t^*(\ell)$ approaches zero, and with positive supply every realization eventually lies on the broad side of the threshold, so that the conventional decline in expected price reappears, as in the unrestricted economy of Hellwig (1980).

In a Gaussian specialization that retains learning from price, we can replace the requirement that ownership remain on the same side of one half in every state with an exact condition on average ownership. Set $\ell = 0$ and choose a fixed distribution for positive supply T such that the public statistic has Gaussian noise:

$$W = w_0(T) \sim N(m_W, v).$$

Appendix A.1.8 constructs the corresponding supply distribution explicitly, and the payoff is now also taken to be Gaussian,

$$X \sim N(\mu_0, a).$$

Posterior uncertainty, disagreement, and expected price are then

$$U = (a^{-1} + b^{-1} + v^{-1})^{-1}, \qquad D = \frac{U^2}{b}, \qquad \mathbb{E}[P] = \mu_0 + \frac{m_W}{b}U.$$

The supply law is held fixed as a changes, so varying a again isolates an increase in fundamental risk. Both uncertainty and disagreement rise, while the allocation continues to be governed by the same supply distribution. In this Gaussian economy, average ownership is

$$\mathbb{E}[\Omega^+] = \Phi\left(-\frac{m_W}{\sqrt{b+v}}\right),$$

and the price comparative static satisfies

$$\frac{d\mathbb{E}[P]}{da} > 0 \quad \Longleftrightarrow \quad \mathbb{E}[\Omega^+] < \frac{1}{2}.$$

The derivative is negative above one half and zero at the boundary, so the threshold from Section 3.1 survives learning from price even when ownership crosses one half across supply realizations. Average ownership alone determines the sign in this Gaussian economy, as Appendix A.1.8 establishes from the posterior moments and the supply distribution.

The clearing condition in (12) gives another way to understand this result. If $\bar{\xi}$ denotes the expected short demand excluded by the floor, then $m_W = \gamma b(\bar{\xi} - \mathbb{E}[T])$, and

$$\mathbb{E}[P] = \mu_0 - \gamma U\big(\mathbb{E}[T] - \bar{\xi}\big).$$

This is the price of an asset whose effective supply is $\mathbb{E}[T] - \bar{\xi}$, so compensation for risk and the premium sustained by constrained pessimism are not separate forces but two parts of a single quantity, scaled by posterior uncertainty. Greater fundamental variance raises U and so amplifies whichever part is larger. In the unrestricted economy, no short demand is excluded and the expression reduces to the familiar $\mu_0 - \gamma U\mathbb{E}[T]$. Narrow ownership instead corresponds to negative expected effective supply, and the same increase in uncertainty then raises the expected price.

Under the same Gaussian assumptions, the effect of learning from price on the two belief moments can be compared with the survey evidence in Section 2. Relative to an investor who observes only private information, define

$$\rho_p = \frac{a^{-1} + b^{-1}}{a^{-1} + b^{-1} + v^{-1}} \in (0, 1).$$

Learning from equilibrium price then changes the two belief moments according to

$$U^{\text{price}} = \rho_p U^{\text{priv}}, \qquad D^{\text{price}} = \rho_p^2 D^{\text{priv}}, \qquad \frac{D^{\text{price}}}{U^{\text{price}}} = \rho_p \frac{D^{\text{priv}}}{U^{\text{priv}}}.$$

Common price information reduces disagreement proportionately more than uncertainty, reproducing the Gaussian comparison used to interpret the fixed-event forecasts in Section 2. Nevertheless, the disagreement remaining after price is observed can still raise valuations when sufficiently pessimistic positions are constrained.

Finally, stochastic outstanding supply is not essential to the ownership result. Appendix A.2.2 considers fixed outstanding supply when the same investors who receive private signals also carry privately observed nontraded exposures to the payoff. Those exposures generate hedging demand, so price remains informative without fully revealing the fundamental. Investors infer the public statistic from price and choose constrained

positions using both their private signal and their exposure, while market clearing determines how aggregate exposure enters that statistic. In the Gaussian specialization, the effect of fundamental risk on expected price is again governed by the breadth of positive ownership, with the one half threshold preserved under fixed float.

Both the economy with fixed beliefs and the rational expectations equilibrium therefore connect the response of prices to fundamental volatility with the fraction of investors who hold the asset. Learning from market prices changes the uncertainty and disagreement entering their decisions, while short sale constraints affect not only their positions but also the information they obtain from price. In the empirical analysis, we examine the corresponding return patterns across stocks with different ownership breadth and borrowing costs, treating firm size as an additional implication of the supply mechanism.

4 From Theory to Measurement

The theory studies investors who receive different information about the terminal payoff of an asset and learn from the price at which they trade. Its principal prediction concerns the effect of greater fundamental volatility on equilibrium prices: a change in risk affects both the uncertainty investors perceive about the payoff and the disagreement among their valuations, while the extent to which those valuations enter demand depends on the constraints investors face and on how widely the asset is held. The empirical analysis examines these implications in stock return data, where beliefs, trading conditions, and ownership can be measured separately.

The one period formulation of the theory identifies the object about which investors form beliefs as the eventual liquidating payoff, rather than a future market price that must itself be determined in equilibrium. In practice, stocks trade repeatedly and information arrives over time, as emphasized in models of dynamic learning and asset prices (Wang, 1993; Banerjee, 2011; Daniel et al., 2023). The question pursued here is nevertheless closely related to the one studied in the model: whether volatility is associated with different return outcomes when investors disagree about future fundamentals and pessimistic positions are costly to establish. We examine that question using measures of expected variance, forecasts of earnings, the cost of borrowing shares, and the fraction of reporting mutual funds that hold each stock.

The sample consists of the ordinary common shares of U.S.-incorporated corporations listed on the NYSE, NYSE American, or Nasdaq, using the stock universe and firm characteristics assembled by Jensen et al. (2023). Monthly stock returns come from CRSP, and market capitalization is taken from the panel of Jensen et al. (2023). Option prices

and implied borrowing rates come from OptionMetrics, analysts' forecasts from I/B/E/S, and mutual fund holdings from S12 filings. Returns are observed from February 1996 through September 2025, with each predictor measured before the return it is used to explain. Some tests require both option prices and analyst forecasts. We distinguish the full universe of stocks from the *covered sample*, which also satisfies the common requirements for the availability of volatility and uncertainty measures described in Appendix C. This distinction is consequential because option and analyst coverage are concentrated among larger stocks, whereas the unconditional IVOL anomaly is strongest among smaller firms. The precise inclusion rules, return adjustments, and variable definitions are given in Appendix C.

We measure idiosyncratic volatility using the residuals from regressions of daily stock returns on the three factors of Fama and French (1993), estimated over the preceding year. Throughout the paper, *IVOL* denotes the standard deviation of these residuals. Ang et al. (2006) use daily residual volatility over the preceding month in their principal specification and also examine a longer formation period. We prefer the year of daily observations because the object of interest is the more persistent volatility associated with uncertainty about an eventual payoff, rather than fluctuations in realized volatility over a few trading days. Vuolteenaho (2002) finds that cash flow news accounts for much of the variation in firm-level returns, providing a reason to connect idiosyncratic return volatility with fundamental payoff risk. Averaging information from daily returns over a longer period also reduces the influence of short-lived movements in volatility, although it need not be preferable when underlying volatility changes rapidly. Because the IVOL anomaly is sensitive to the construction of the volatility measure (Bali and Cakici, 2008), the principal tests are repeated using weekly residual volatility over the preceding three years and daily residual volatility over the preceding month. We also control for the standard deviation of daily excess returns during the formation month. This measure captures total rather than idiosyncratic volatility and uses one month of observations rather than the year used for baseline IVOL.

The theoretical uncertainty measure is the variance of the payoff under an investor's posterior beliefs. Its empirical counterpart is constructed from the prices of options on individual stocks, which provide information about the distribution of future returns before those returns are realized. Following Martin (2017); Martin and Wagner (2019), we use firm SVIX squared, expressed as a variance, at a constant horizon of 365 days. The annual horizon corresponds approximately to the horizon of the earnings forecasts used to measure disagreement. The principal measure includes the total option-implied variance of the stock, while a second measure removes a component associated with market variance

to focus more closely on idiosyncratic risk. Option prices reflect risk-neutral valuations, rather than investors' subjective probability distributions, so these quantities are proxies for uncertainty rather than direct estimates of the posterior variance in the model. We use a 30-day horizon only when reproducing the cardinal specification of Martin and Wagner (2019).

To measure disagreement, we use the variance across analysts' forecasts of annual earnings per share, scaled by the square of the stock price. The forecasts concern a common fiscal outcome approximately one year ahead, making their dispersion a natural measure of differences in assessments of future firm fundamentals. Scaling by price places forecast dispersion in units comparable to squared stock returns. Here disagreement concerns earnings forecasts, whereas uncertainty is inferred from options on stock returns. The survey evidence reported earlier allows a more direct examination of the theoretical restrictions because forecasters report both subjective probability distributions and point forecasts for the same outcome.

Two further variables permit us to examine the trading mechanism. The stock-specific implied borrowing rate from OptionMetrics measures the cost associated with a short position, providing an observable counterpart to the constraints on pessimistic positions in the model. Although an option-implied borrowing rate is not a direct observation of the supply of lendable shares, higher values indicate that taking a negative position is more costly. To measure how broadly an asset is held, we follow Chen et al. (2002) and use the fraction of reporting mutual funds that hold the stock, with holdings lagged by one quarter. This measure of ownership breadth is not the fraction of all potential investors who hold positive positions, as in the theoretical economy, but it captures observable variation in the extent to which investors participate in holding a given stock. The model gives these variables separate roles: the cost of shorting governs the influence of pessimistic demand, while the breadth of ownership indicates how far down the distribution of valuations investors must be drawn to absorb the asset. We therefore examine them separately before asking how they operate together.

Firm size offers a further, less direct comparison with the model. A larger outstanding supply, for a given expected payoff per unit, must be absorbed by more investors and consequently brings less optimistic valuations into market clearing. Market capitalization is an imperfect proxy for this supply margin because it also incorporates differences in expected payoffs and prices, so size is treated principally as an additional implication and a control rather than as a substitute for the more direct evidence on borrowing costs and ownership. Measures of liquidity likewise enter the supporting analysis, since the ease of

trading can be correlated with borrowing conditions without representing the constraint on pessimistic positions itself.

The empirical tests proceed from the distinction between beliefs to the conditions governing how they enter prices. We first examine whether IVOL is associated with both uncertainty and disagreement, and whether the two measures have the opposing return implications suggested by the theory. We then ask whether the return coefficients become more negative as borrowing costs rise and less negative as ownership broadens, before considering the two dimensions jointly. The model also implies that public information can reduce disagreement by drawing investors' posterior expectations together. Earnings announcements therefore permit us to examine whether the negative return relation associated with disagreement, and the IVOL anomaly itself, are concentrated around the arrival of common information. Finally, we compare these results with alternative explanations of the IVOL puzzle, particularly those associated with lottery demand, return reversal, and limits to arbitrage.

For the principal return tests, we estimate regressions in each monthly cross-section and report the average coefficients, with t statistics based on Newey–West standard errors using six lags. The dependent variable is the subsequent monthly return in percentage points, while the explanatory variables are measured at the end of the preceding month. The regressions introduce uncertainty and disagreement together, then add the conventional firm characteristics, and finally control for IVOL and realized volatility. The characteristics are market equity, the book-to-market ratio, momentum, one-month reversal, operating profitability, asset growth, and analyst coverage where available. The financial characteristics follow Jensen et al. (2023), while analyst coverage is measured by the I/B/E/S analyst count. Missing values of the five financial characteristics other than market equity are replaced with monthly medians, as described in Appendix C; the other regressors are not imputed. The table notes specify whether IVOL and the belief measures enter as standardized logarithms or ranks, and Appendix C gives the transformations of the controls. These transformations make the coefficients comparable across stocks. Appendix C describes the treatment of extreme observations and the specifications for which the estimation procedure differs. Throughout, *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels.

5 Empirical Evidence

5.1 Idiosyncratic Volatility, Uncertainty, and Disagreement

Before examining the belief measures, it is useful to establish the negative relation between IVOL and subsequent returns in the stock universe used for the analysis. Table 1 reports monthly cross-sectional regressions of next-month returns on standardized log IVOL, both with and without the firm characteristics described above. Among all stocks, an increase of one standard deviation in log IVOL is associated with a reduction in the subsequent monthly return of 0.566 percentage points ($t = -3.05$) without controls and 0.519 points ($t = -2.98$) after the characteristics are included. These are economically substantial differences: the latter coefficient corresponds to more than six percentage points at an annual rate when the monthly estimates are added. The negative relation is especially pronounced among smaller stocks, where the coefficient in the smallest size quintile is -1.151 ($t = -6.20$), compared with -0.144 ($t = -0.82$) in the largest quintile.

Table 1: IVOL and subsequent returns across samples

	All stocks		Covered sample		Smallest size		Largest size	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
IVOL	-0.566^{***}	-0.519^{***}	-0.218	-0.200	-1.170^{***}	-1.151^{***}	-0.106	-0.144
	(-3.05)	(-2.98)	(-0.98)	(-0.99)	(-5.29)	(-6.20)	(-0.50)	(-0.82)
Characteristics	✗	✓	✗	✓	✗	✓	✗	✓
Months	356	356	356	356	356	356	356	356
Stocks	15,736	15,736	7,614	7,614	7,798	7,798	3,880	3,880
Mean R^2	0.032	0.062	0.043	0.084	0.025	0.054	0.047	0.111

Monthly Fama–MacBeth regressions of next-month excess returns (percentage points) on standardized log one-year daily FF3 IVOL. Columns (1) and (2) use all stocks; columns (3) and (4) use the covered sample of stocks with both option and analyst forecast data; columns (5)–(8) use the smallest and largest quintiles of market equity among all stocks. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4 other than analyst coverage, which is not available for all stocks. t statistics in parentheses use Newey–West standard errors with six lags.

Requiring both option prices and analyst earnings forecasts changes the composition of the sample. Many of the smaller firms with the most negative IVOL coefficients are not covered by both sources, and the coefficient in the covered sample is correspondingly smaller, at -0.218 ($t = -0.98$) without controls and -0.200 ($t = -0.99$) with them. Within this more selective sample, the subsequent tests examine how the IVOL return relation varies with firm characteristics and trading conditions.

The historical portfolio evidence complements these regressions. From February 1996 through September 2025, the value-weighted portfolio that buys stocks in the highest IVOL decile and sells those in the lowest has a three-factor alpha of -1.70 percentage points per month ($t = -4.51$), while its raw return spread is -1.10 points ($t = -1.85$). Accounting for exposure to the market, size, and value factors therefore does not remove the negative relation. Figure 1 places the result in the longer history beginning in 1962 and shows that the anomaly has persisted through different market conditions, although its magnitude has varied appreciably over time. Neither the historical result nor the size variation by itself identifies the mechanism, but together they establish the empirical phenomenon that the more detailed tests seek to explain.

We next examine whether the two belief measures are associated with observed IVOL. Table 2 reports regressions of standardized log IVOL on the standardized ranks of uncertainty and disagreement within each month. With total option-implied uncertainty, the loading on uncertainty is 0.831 ($t = 104.91$), while disagreement has a smaller but highly significant loading of 0.067 ($t = 9.24$). The average cross-sectional R^2 is 0.767. The coefficient on uncertainty is considerably larger, as IVOL and option variance both measure return volatility, while the positive coefficient on disagreement indicates that the information in analysts' forecasts is not exhausted by the option measure.

Table 2: Uncertainty and disagreement in idiosyncratic volatility

	Covered sample		S&P 500	
	(1)	**(2)**	**(3)**	**(4)**
U:	Total	Market-adj.	Total	Market-adj.
U	0.831***	0.677***	0.789***	0.575***
	(104.91)	(49.07)	(96.20)	(28.73)
D	0.067***	0.206***	0.060***	0.284***
	(9.24)	(14.53)	(9.29)	(15.38)
Months	356	356	347	347
Stocks	7,594	7,594	1,090	1,090
Mean R^2	0.767	0.650	0.673	0.515

Monthly Fama–MacBeth regressions of standardized log one-year daily FF3 IVOL on standardized ranks of uncertainty U and disagreement D, without further controls. U is 365-day total uncertainty in columns (1) and (3) and 365-day market-adjusted uncertainty in columns (2) and (4). Within each month, uncertainty, disagreement, and log IVOL are ranked or standardized within the reported sample. t statistics in parentheses use Newey–West standard errors with six lags.

The comparison is sensitive to whether the option measure includes aggregate market risk, as one would expect when relating firm residual volatility to option variance. After

removing a component associated with market variance, the uncertainty loading falls to 0.677, while the disagreement loading rises to 0.206 ($t = 14.53$). A similar change appears among historical S&P 500 constituents, where disagreement has loadings of 0.060 with total uncertainty and 0.284 with the market-adjusted measure, both estimated very precisely. Thus IVOL is closely associated with uncertainty on average, but disagreement is also significantly related to its cross-sectional variation, including among large firms.

The availability of options and analyst forecasts may also attenuate the observed role of disagreement by excluding smaller and less liquid stocks. Appendix Table A3 examines this possibility by reweighting the covered observations so that, within each month, their distribution across size and liquidity groups more closely resembles that of the full stock universe. The disagreement loading in the IVOL regression rises from 0.079 to 0.107 under the 2×2 reweighting scheme, while its return coefficient changes from -0.186 ($t = -2.01$) to -0.328 ($t = -2.95$). Giving more weight to the firms that are relatively scarce in the covered sample therefore strengthens rather than weakens the negative return relation associated with disagreement.

5.2 Pricing Uncertainty and Disagreement

The central prediction is not merely that uncertainty and disagreement are both associated with IVOL, but that their implications for prices and expected returns differ. Higher uncertainty ordinarily discourages investors from bearing risk, whereas disagreement can raise prices when pessimistic positions are constrained. We begin by comparing the option measure with the expected return specification of Martin and Wagner (2019), which relates returns to option-implied variance in its original units, and then examine the two belief measures together in monthly cross-sectional regressions. The latter specifications allow the relation associated with disagreement to be distinguished from that of uncertainty and from the variation already captured by past return volatility.

Table 3 uses 30-day option variance, matching the horizon of Martin and Wagner (2019), and reports results for historical S&P 500 constituents both during and after the period in their study, as well as for the broader covered sample. The estimated uncertainty coefficients among the S&P 500 stocks are positive in each specification, ranging from 0.174 to 0.224 in the earlier period and from 0.306 to 0.456 subsequently. Their signs are consistent with a positive relation between uncertainty and expected returns, although the coefficients are not statistically distinguishable from zero at the monthly horizon. In the wider covered sample, the corresponding estimates are close to zero. We therefore regard this comparison as suggestive of a more positive uncertainty relation among larger,

more widely held firms, rather than as strong independent evidence of a premium in the cardinal option variance specification.

Table 4 turns to the principal measures, using option variance over one year and the separately observed dispersion in analysts' forecasts. Disagreement predicts lower subsequent returns across specifications: its coefficient is -0.186 percentage points per month ($t = -2.01$) without controls, -0.205 ($t = -2.99$) with the firm characteristics, and -0.177 ($t = -2.68$) after IVOL and realized volatility are also included. The last estimate implies that a one standard deviation increase in log disagreement is associated with about 0.18 percentage points lower returns in the following month, an economically meaningful magnitude that remains statistically significant after controlling for the familiar return predictors.

Uncertainty does not exhibit the same negative relation. Its coefficient is -0.079 ($t = -0.45$) alongside disagreement without controls and -0.004 ($t = -0.03$) after the characteristics are added, neither of which is distinguishable from zero. Once IVOL and realized volatility are included, however, the uncertainty coefficient is positive, at 0.253 percentage points per month ($t = 2.25$). A similar estimate of 0.174 ($t = 2.63$) obtains with the market-adjusted option measure in Panel B. After controlling for volatility already observed in returns, the option measure has the positive coefficient suggested by compensation for bearing risk, while disagreement continues to predict lower returns. The positive uncertainty relation is also visible among more broadly held stocks, where it reaches 0.275 ($t = 2.17$) in Table 7 after the volatility controls are included.

We also allow the return relation associated with disagreement to depend on the level of uncertainty. The coefficient on the interaction $U \times D$ is consistently negative, at -0.186 ($t = -3.96$) without controls, -0.138 ($t = -3.38$) with the characteristics, and -0.138 ($t = -3.62$) with the volatility controls. In the specification with characteristics, a one standard deviation increase in log disagreement is associated with approximately 0.31 percentage points lower monthly returns when ranked uncertainty is one standard deviation above its mean, compared with only 0.03 points when it is one standard deviation below. The economic implication is that disagreement predicts especially low returns among stocks for which the underlying payoff is also perceived to be uncertain. This accords with the demand mechanism developed in the theory: when valuation differences are large in an economy where pessimistic investors cannot take correspondingly negative positions, those differences can have a greater effect on the market-clearing price.

Two comparisons in the appendix help interpret the positive coefficient on option-implied uncertainty. The option measure predicts realized idiosyncratic volatility over the following year even after past IVOL and realized volatility are controlled for. In

Table 3: 30-day option-implied variance and subsequent returns

	(1)	(2)	(3)	(4)
Panel A: S&P 500, February 1996–October 2014				
U	0.199	0.186	0.174	0.224
	(0.86)	(0.79)	(0.73)	(0.90)
D		0.061	−0.057	−0.011
		(0.48)	(−0.51)	(−0.11)
IVOL				0.026
				(0.16)
Realized vol.				−0.195
				(−1.47)
Observations	105,234	105,234	105,234	105,234
Panel B: S&P 500, November 2014–December 2024				
U	0.306	0.336	0.386	0.456
	(1.03)	(1.13)	(1.29)	(1.47)
D		−0.105	−0.055	0.090
		(−0.51)	(−0.31)	(0.54)
IVOL				−0.407**
				(−2.10)
Realized vol.				−0.036
				(−0.24)
Observations	53,653	53,653	53,653	53,653
Panel C: Covered sample, full period				
U	−0.015	−0.007	−0.001	0.004
	(−0.34)	(−0.18)	(−0.02)	(0.09)
D		−0.082	−0.094	−0.063
		(−0.81)	(−1.14)	(−0.74)
IVOL				0.011
				(0.08)
Realized vol.				−0.116
				(−1.22)
Observations	695,671	695,671	695,671	695,671
Characteristics	✗	✗	✓	✓
IVOL and realized volatility	✗	✗	✗	✓

Pooled regressions in the specification of Martin and Wagner (2019). All specifications include an intercept and no stock or month fixed effects. Standard errors are clustered by stock and month. The dependent variable is the next-month return net of the value-weighted sample return, divided by the gross risk-free return. U is firm SVIX variance over the 30-day horizon, net of the value-weighted average, in variance units. The coefficient on U is reported in these units, and the remaining coefficients are in percentage points. D is standardized log EPS disagreement; IVOL is standardized log one-year daily FF3 IVOL. Panel A ends in October 2014, the last month of the sample in Martin and Wagner (2019). A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics are reported in parentheses.

Table 4: Return slopes for uncertainty and disagreement

	Additive			Interaction		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Panel A: 365-day total uncertainty						
U	-0.079	-0.004	0.253^{**}	-0.089	-0.020	0.244^{**}
	(-0.45)	(-0.03)	(2.25)	(-0.51)	(-0.12)	(2.13)
D	-0.186^{**}	-0.205^{***}	-0.177^{***}	-0.130	-0.168^{**}	-0.144^{**}
	(-2.01)	(-2.99)	(-2.68)	(-1.44)	(-2.51)	(-2.20)
$U \times D$				-0.186^{***}	-0.138^{***}	-0.138^{***}
				(-3.96)	(-3.38)	(-3.62)
IVOL			-0.132			-0.138
			(-1.40)			(-1.49)
Realized volatility			-0.225^{***}			-0.226^{***}
			(-3.77)			(-3.82)
Mean R^2	0.059	0.096	0.102	0.064	0.100	0.105
Panel B: 365-day market-adjusted uncertainty						
U	-0.085	-0.006	0.174^{***}	-0.069	0.005	0.201^{***}
	(-0.59)	(-0.05)	(2.63)	(-0.48)	(0.04)	(3.05)
D	-0.177^{*}	-0.201^{**}	-0.163^{**}	-0.125	-0.169^{**}	-0.124^{*}
	(-1.70)	(-2.48)	(-2.41)	(-1.23)	(-2.09)	(-1.83)
$U \times D$				-0.195^{***}	-0.145^{***}	-0.159^{***}
				(-4.33)	(-3.71)	(-4.31)
IVOL			-0.094			-0.118
			(-0.77)			(-0.97)
Realized volatility			-0.206^{***}			-0.210^{***}
			(-3.02)			(-3.12)
Mean R^2	0.050	0.091	0.101	0.055	0.094	0.104
Characteristics	✗	✓	✓	✗	✓	✓
IVOL and realized volatility	✗	✗	✓	✗	✗	✓
Months	356	356	356	356	356	356
Stocks	7,593	7,593	7,593	7,593	7,593	7,593

Monthly Fama–MacBeth regressions of next-month returns (percentage points) on standardized ranked uncertainty U, standardized log disagreement D, and their product. IVOL is one-year daily FF3 IVOL and realized volatility is total return volatility, each logged, winsorized, and standardized within month. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Appendix Table A5, its coefficient is 0.364 ($t = 19.68$), and remains 0.282 ($t = 17.81$) when disagreement and the characteristics are included. The option measure therefore contains information about future volatility beyond that in the historical return series. Appendix Table A6 also considers the spread between option-implied and realized volatility and a proxy for the variance risk premium. Adding either measure separately leaves the uncertainty coefficient positive, at 0.318 ($t = 2.09$) or 0.336 ($t = 2.93$), respectively. The two proxies are highly correlated, however, and including both makes the uncertainty estimate imprecise.

5.3 Constraints and the IVOL Relation

We now turn from the average return relations to the trading conditions under which the model predicts that disagreement will have the greatest influence on prices. Short sale constraints make it costly for pessimistic investors to offset the demand of optimistic investors, while the breadth of ownership indicates how widely the stock is held across investors. These are related but distinct aspects of the mechanism.

Table 5 considers borrowing costs, ownership breadth, and, as an additional comparison, firm size. Its first three columns examine how uncertainty and disagreement are associated with IVOL at different levels of each characteristic. The next three columns estimate how their return coefficients vary with the same characteristics, while the final three examine IVOL directly. Each characteristic is standardized within month before entering the interactions, and the standard firm characteristics are included throughout. This arrangement allows us to distinguish two related observations: the association of measured IVOL with the belief measures may vary across trading environments, and the return implications of disagreement and uncertainty may vary even after their measured levels are held fixed.

The results for borrowing costs are particularly clear. Among stocks that are more expensive to short, IVOL is more closely associated with disagreement: the interaction in the IVOL regression is 0.042 ($t = 9.62$), while its interaction with uncertainty is -0.024 ($t = -8.34$). More importantly for the price mechanism, the return coefficient on disagreement becomes more negative as borrowing costs rise, with an interaction of -0.106 ($t = -3.42$). The corresponding interaction for IVOL is -0.139 ($t = -5.18$). These estimates are statistically precise, and their economic interpretation is consistent across the two tests: disagreement and IVOL are associated with especially low subsequent returns where it is relatively costly to establish the pessimistic positions that would otherwise offset optimistic demand.

Table 5: Constraints, the composition of IVOL, and return slopes

	Dependent variable: IVOL			Dependent variable: next-month return					
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
M:	Borrow.	Breadth	Size	Borrow.	Breadth	Size	Borrow.	Breadth	Size
U	0.744^{***}	0.724^{***}	0.731^{***}	-0.006	-0.001	-0.024			
	(95.97)	(80.37)	(88.40)	(-0.03)	(-0.01)	(-0.14)			
D	0.095^{***}	0.103^{***}	0.104^{***}	-0.141^{**}	-0.141^{**}	-0.147^{**}			
	(11.86)	(12.29)	(12.22)	(-2.26)	(-2.26)	(-2.38)			
$U \times M$	-0.024^{***}	0.016^{***}	0.017^{***}	-0.060^{***}	0.092^{***}	0.053			
	(-8.34)	(4.11)	(3.18)	(-2.88)	(2.91)	(1.55)			
$D \times M$	0.042^{***}	-0.034^{***}	-0.033^{***}	-0.106^{***}	0.099^{***}	0.099^{***}			
	(9.62)	(-6.95)	(-6.81)	(-3.42)	(2.98)	(2.74)			
IVOL							-0.181	-0.185	-0.218
							(-1.17)	(-1.20)	(-1.39)
IVOL$\times M$							-0.139^{***}	0.159^{***}	0.135^{***}
							(-5.18)	(4.94)	(3.73)
Characteristics	✓	✓	✓	✓	✓	✓	✓	✓	✓
Months	356	356	356	356	356	356	356	356	356
Stocks	7,575	7,592	7,594	7,573	7,591	7,593	7,573	7,591	7,593
Mean R^2	0.803	0.796	0.797	0.100	0.101	0.102	0.093	0.094	0.094

Monthly Fama–MacBeth regressions. The dependent variable is standardized log one-year daily FF3 IVOL in columns (1)–(3) and the next-month return (percentage points) in columns (4)–(9). U and D enter as standardized within-month ranks, and each column interacts them, or IVOL, with the standardized rank of the characteristic M named in the column: implied borrowing cost, mutual fund ownership breadth, or market equity. Implied borrowing cost measures the cost of short positions, ownership breadth the fraction of mutual funds holding the stock, and market equity firm size. The level of M and the standard characteristics are included but not shown. U is 365-day total uncertainty, and Table A7 reports the market-adjusted version. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Ownership breadth produces the opposite ordering, as the theory would suggest. When a larger fraction of reporting mutual funds holds the stock, IVOL is less closely associated with disagreement, with an interaction of -0.034 ($t = -6.95$), while the interaction with uncertainty is 0.016 ($t = 4.11$). The return coefficients also become less negative as ownership broadens: the interaction is 0.099 ($t = 2.98$) for disagreement and 0.159 ($t = 4.94$) for IVOL. These results provide evidence for the ownership implication independently of the borrowing cost comparison, while the corresponding size interactions have the same general ordering.

Ownership breadth and firm size are closely related, so we also include their interactions with IVOL jointly. The breadth interaction remains positive at 0.235 ($t = 3.48$), while the size interaction is -0.070 ($t = -0.93$) (Appendix Table A9). The ownership gradient therefore persists after allowing the IVOL return slope to vary with firm size.

The findings do not depend on using total option-implied uncertainty or on the particular year of daily returns used to construct IVOL. With the market-adjusted option measure, the disagreement interaction with borrowing cost is 0.032 ($t = 7.75$) in the IVOL regression and -0.105 ($t = -3.46$) in the return regression (Appendix Table A7). The negative IVOL return interaction with borrowing cost also remains when IVOL is measured from weekly returns over three years or daily returns over one month, at -0.128 ($t = -5.28$) and -0.121 ($t = -4.91$), respectively, while the corresponding ownership interactions remain positive (Appendix Table A8). The magnitudes vary, as one would expect from volatility measures constructed over different periods, but the conditional return patterns are similar.

Having considered the two dimensions separately, we next examine them together. Table 6 divides the stocks in the lowest and highest terciles of implied borrowing cost into groups with broad and narrow mutual fund ownership, and estimates the return regressions separately within each group. This comparison asks whether the IVOL relation is particularly negative when both circumstances identified by the model are present: relatively few investors hold the stock, and pessimistic positions are costly to establish. It also permits the difference between the borrowing cost groups to be compared across the two ownership groups, rather than relying only on the level of the coefficient in the most constrained group.

The IVOL coefficient is most negative among stocks that are narrowly held and costly to short, at -0.499 percentage points per month ($t = -2.61$). In each of the other three groups, the coefficient is smaller in magnitude and statistically insignificant. The coefficient is nearly half a percentage point per month in magnitude, an economically substantial return difference for the standardized IVOL measure. Within narrowly held stocks, the difference in the IVOL coefficient between high and low borrowing costs is -0.434 points ($t = -4.30$). Among stocks with high borrowing costs, the difference between narrow and broad ownership is -0.352 ($t = -3.48$). The double difference is also significant, at -0.385 ($t = -3.14$), indicating that the effect of borrowing costs on the IVOL slope is more negative when ownership is narrow.

Disagreement behaves similarly in the groups that the theory identifies as most exposed to constrained pessimism. Its return coefficient is -0.379 ($t = -3.61$) among narrowly held stocks that are expensive to short, and it remains -0.290 ($t = -2.78$) after IVOL and realized volatility are included. The corresponding coefficients in the other groups are considerably smaller, and the within-group return relation in the most constrained group is both economically and statistically substantial. The double difference for disagreement has the predicted negative sign, although it is not estimated precisely enough to establish

Table 6: Return slopes by ownership breadth and borrowing cost

	Low borrowing cost		High borrowing cost		Contrasts		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
	Broad	Narrow	Broad	Narrow	(4)−(2)	(4)−(3)	Double difference
Panel A: IVOL							
IVOL	−0.098	−0.065	−0.147	−0.499***	−0.434***	−0.352***	−0.385***
	(−0.66)	(−0.44)	(−0.88)	(−2.61)	(−4.30)	(−3.48)	(−3.14)
Panel B: Uncertainty and disagreement							
U	0.045	−0.011	0.013	−0.097	−0.085	−0.110	−0.054
	(0.20)	(−0.07)	(0.07)	(−0.52)	(−0.63)	(−0.83)	(−0.36)
D	−0.045	−0.054	−0.150*	−0.379***	−0.325***	−0.229**	−0.221
	(−0.40)	(−0.58)	(−1.74)	(−3.61)	(−2.82)	(−2.15)	(−1.55)
$U \times D$	−0.030	0.038	−0.068	−0.064	−0.102	0.004	−0.065
	(−0.37)	(0.47)	(−1.04)	(−1.01)	(−1.03)	(0.05)	(−0.52)
Panel C: Uncertainty and disagreement with volatility controls							
U	0.423	0.047	0.213	0.278	0.231	0.065	0.441*
	(1.63)	(0.30)	(1.46)	(1.50)	(1.17)	(0.37)	(1.68)
D	−0.016	−0.053	−0.129	−0.290***	−0.237**	−0.160	−0.124
	(−0.15)	(−0.58)	(−1.54)	(−2.78)	(−2.04)	(−1.41)	(−0.84)
$U \times D$	−0.002	0.021	−0.081	−0.104	−0.125	−0.022	−0.045
	(−0.02)	(0.26)	(−1.25)	(−1.57)	(−1.26)	(−0.30)	(−0.37)
IVOL	−0.172	0.075	−0.118	−0.183	−0.259*	−0.066	−0.313
	(−1.41)	(0.60)	(−0.83)	(−1.19)	(−1.69)	(−0.48)	(−1.49)
Realized vol.	−0.185**	−0.149**	−0.136	−0.428***	−0.279***	−0.292***	−0.327**
	(−2.20)	(−1.99)	(−1.52)	(−4.49)	(−2.59)	(−2.71)	(−2.42)
Characteristics	✓	✓	✓	✓	—	—	—
Months	356	356	356	356	356	356	356
Stocks	2,110	6,277	4,901	5,899	7,051	7,436	7,544

Separate monthly Fama–MacBeth regressions of next-month returns within each group. Groups are halves of mutual fund ownership breadth within the lowest and highest terciles of implied borrowing cost. Column (7) equals (4)−(2)−[(3)−(1)]. Differences are formed month by month before inference. U is 365-day total uncertainty. U and D enter as standardized logarithms. Stocks in the contrast columns count the union of the groups involved. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4 other than analyst coverage. t statistics in parentheses use Newey–West standard errors with six lags.

an interaction of the same strength as that for IVOL. Once the belief measures and realized volatility are included, IVOL itself is no longer significant in the narrowly held, expensive-to-short group, with a coefficient of -0.183 ($t = -1.19$). An analogous division by firm size and borrowing cost yields the same broad ordering, though the joint effect is less pronounced (Appendix Table A10).

5.4 Ownership and the IVOL Relation

The preceding comparisons indicate that the negative return coefficients on IVOL and disagreement become less pronounced as mutual fund ownership broadens. We examine this relation more closely in Table 7, which divides stocks each month into five groups according to the lagged fraction of reporting mutual funds that hold them and repeats the return regressions within each group. Market capitalization is included among the standard characteristics, so the comparison is not simply another way of dividing firms by size. It allows us to examine the model's prediction that disagreement should have less influence on prices when the asset is held by a broader set of investors, while also asking whether the uncertainty coefficient behaves differently across these stocks.

The decline in the anomaly across ownership groups is economically large and statistically clear. In the narrowest ownership quintile, the IVOL coefficient is -0.434 percentage points per month ($t = -2.79$), whereas in the broadest it is -0.095 ($t = -0.97$), giving a difference of 0.339 points ($t = 3.06$). The corresponding relation for disagreement is particularly striking: its coefficient changes from -0.390 ($t = -4.23$) among narrowly held stocks to -0.035 ($t = -0.53$) among broadly held stocks, a difference of 0.355 points ($t = 3.98$). When IVOL and realized volatility are also included, the disagreement coefficient remains -0.352 ($t = -3.91$) in the narrowest group and is close to zero in the broadest, with a difference of 0.324 ($t = 3.69$). The fact that this gradient persists after controlling for observed volatility is important, since it indicates that the ownership result is not merely a comparison between stocks with different levels of IVOL.

The uncertainty coefficients are less stable across specifications, although the broadest ownership group has a positive coefficient of 0.275 ($t = 2.17$) when IVOL and realized volatility are included. This is consistent with the conventional compensation for risk becoming more visible where the disagreement mechanism is weaker. At the same time, the difference between the uncertainty coefficients in the extreme ownership groups is not estimated precisely, so we do not infer a sharply identified ownership gradient for uncertainty from this table. A related finding is that, after the two belief measures and realized volatility are held fixed, the difference in the IVOL coefficients between the

Table 7: IVOL, uncertainty, and disagreement across ownership breadth

	Ownership breadth quintile					
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
	1 (low)	2	3	4	5 (high)	5−1
Panel A: IVOL						
IVOL	−0.434***	−0.125	−0.112	−0.077	−0.095	0.339***
	(−2.79)	(−0.81)	(−0.81)	(−0.65)	(−0.97)	(3.06)
Panel B: Uncertainty and disagreement						
U	−0.168	−0.002	−0.058	0.005	0.025	0.193
	(−1.05)	(−0.02)	(−0.43)	(0.04)	(0.20)	(1.50)
D	−0.390***	−0.134**	−0.059	−0.037	−0.035	0.355***
	(−4.23)	(−1.97)	(−0.89)	(−0.54)	(−0.53)	(3.98)
Panel C: Uncertainty and disagreement with volatility controls						
U	0.125	0.170	0.065	0.168	0.275**	0.149
	(0.99)	(1.44)	(0.68)	(1.39)	(2.17)	(0.96)
D	−0.352***	−0.122*	−0.051	−0.025	−0.028	0.324***
	(−3.91)	(−1.81)	(−0.76)	(−0.37)	(−0.43)	(3.69)
IVOL	−0.169	−0.047	−0.012	−0.058	−0.154**	0.015
	(−1.47)	(−0.42)	(−0.11)	(−0.77)	(−2.19)	(0.12)
Realized vol.	−0.349***	−0.240***	−0.167**	−0.174**	−0.171**	0.178*
	(−4.38)	(−3.02)	(−2.32)	(−2.05)	(−2.57)	(1.86)
Characteristics	✓	✓	✓	✓	✓	—
Months	356	356	356	356	356	356
Stocks	5,267	4,768	3,729	2,610	1,192	6,105

Separate monthly Fama–MacBeth regressions of next-month returns within quintiles of lagged mutual fund ownership breadth. Column (6) is the month-by-month difference between quintiles 5 and 1. U is 365-day total uncertainty. Within each quintile, IVOL, U, and D enter as standardized ranks and realized volatility as a standardized logarithm. Stocks in column (6) count both quintiles. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

broadest and narrowest groups becomes small and insignificant. The evidence for a changing return relation is consequently most direct for disagreement itself, while the additional controls account for much of the ownership variation in the IVOL coefficient.

To examine how much of the ownership pattern in the IVOL anomaly is associated with the two belief measures, Table 8 decomposes the coefficient from a regression of returns on standardized log IVOL alone. The exercise retains the monthly transformations calculated before the ownership split, whereas Table 7 uses ranks calculated within each quintile and includes the firm characteristics. Within each month and group, the univariate

coefficient equals the covariance of IVOL with returns divided by the variance of IVOL. We allocate the covariance explained jointly by uncertainty and disagreement using their Shapley contributions, which average the contributions obtained when each measure enters first and after the other. This treats the correlated measures symmetrically. The two contributions and the residual sum to the original IVOL coefficient; they are not return coefficients from a joint regression.

Table 8: Decomposition of the IVOL slope across ownership breadth

	Ownership breadth quintile					
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
	1 (low)	2	3	4	5 (high)	5−1
IVOL slope	−0.599***	−0.137	−0.130	−0.052	−0.048	0.551***
	(−2.76)	(−0.74)	(−0.69)	(−0.29)	(−0.28)	(5.05)
U contribution	−0.272*	−0.042	−0.083	−0.001	0.033	0.305***
	(−1.92)	(−0.32)	(−0.60)	(−0.01)	(0.21)	(3.76)
D contribution	−0.149***	−0.019	−0.012	0.005	0.002	0.152***
	(−3.39)	(−0.52)	(−0.38)	(0.17)	(0.07)	(4.50)
Residual	−0.177***	−0.076	−0.035	−0.055	−0.083**	0.095
	(−2.81)	(−1.64)	(−0.68)	(−1.39)	(−2.46)	(1.38)
Months	356	356	356	356	356	356
Stocks	5,267	4,768	3,729	2,610	1,192	—

The first row reports monthly Fama–MacBeth coefficients from univariate regressions of next-month returns on standardized log IVOL within ownership quintiles. The other rows allocate the coefficient to *U*, *D*, and a residual using the Shapley method in Appendix C; they are not return slopes. *U* is 365-day total uncertainty and enters as a standardized rank; *D* enters as a standardized logarithm. Monthly transformations precede the ownership split, after which variables are demeaned without rescaling. Column (6) is the monthly difference between quintiles 5 and 1. *t* statistics use Newey–West standard errors with six lags.

The decomposition shows that the measured belief variables account for most of the difference in the univariate IVOL slope between narrowly and broadly held stocks. The coefficient rises from -0.599 points ($t = -2.76$) in the lowest ownership quintile to -0.048 in the highest, a difference of 0.551 points ($t = 5.05$). The contributions associated with uncertainty and disagreement change by 0.305 ($t = 3.76$) and 0.152 ($t = 4.50$), respectively, while the change in the residual contribution, 0.095, is not significant. Together the measured belief variables therefore account for approximately 83 percent of the difference. This allocation does not imply that uncertainty has a negative price of risk in the narrowest ownership group. Rather, it describes how IVOL and returns covary with the two measures in that group. Read alongside Table 7, it indicates that a substantial part of the weakening of the IVOL anomaly as ownership broadens is associated with uncertainty and disagreement, and that disagreement is the measure whose conditional return coefficient changes most clearly across the ownership distribution.

5.5 Temporal Structure of Disagreement and the IVOL Anomaly

The preceding results examine which stocks exhibit the most negative return relations associated with disagreement and IVOL. The information structure of the model also provides a reason to examine when those returns are realized. Investors may hold different posterior expectations after receiving private information, and constraints on pessimistic positions can allow this disagreement to support elevated prices. When information subsequently becomes available to investors in common, their posterior means may move closer together and the effect of the constraint on valuation may diminish. Earnings announcements offer an opportunity to examine this implication, because they reveal information about firm fundamentals at identifiable dates around which investors have previously formed different expectations.

Berkman et al. (2009) document that stocks with greater differences of opinion earn lower returns around earnings announcements. We extend that evidence by examining whether disagreement helps explain the concentration of the IVOL anomaly around earnings announcements. We distinguish uncertainty from disagreement as channels through which fundamental volatility affects prices. Disagreement accounts for a substantial portion of the negative relation around announcements and continues to predict lower returns after conditioning on realized earnings surprises. Using the announcement dates recorded by Compustat and the corresponding CRSP trading days, we define the announcement period as the day before the announcement, the announcement day, and the following trading day. We then compare the average return

per trading day in this three-day period with the average return on the other trading days of the month, using identical stock-month observations within each specification. Table 9 reports the monthly average regression coefficients for these two sets of days and for their difference.

The distinction is pronounced for disagreement. With uncertainty and the standard firm characteristics included, a one standard deviation increase in prior disagreement is associated with an average daily return 0.0816 percentage points lower during the announcement period ($t = -3.55$), whereas the coefficient on other trading days is 0.0004 ($t = 0.10$). The difference of -0.0821 points per day is itself statistically significant ($t = -3.53$). Over the three trading days surrounding an announcement, the coefficient corresponds to roughly one quarter of a percentage point in cumulative returns, although the regressions are estimated using average daily returns. The uncertainty coefficient, by comparison, is small and statistically insignificant both around announcements and on other days. The interaction between uncertainty and disagreement is also more negative during the announcement period, with a coefficient of -0.0587 ($t = -4.11$) compared with approximately zero on other days. The finding that disagreement predicts low returns when information becomes public is therefore not simply a reflection of a general negative relation between uncertainty and returns.

Other return anomalies have been found to be concentrated around earnings announcements (Engelberg et al., 2018), so the timing alone cannot establish that the mechanism is specific to disagreement. It is nevertheless informative to ask whether the IVOL anomaly is realized at the same time and whether conditioning on disagreement changes its magnitude. In the full stock sample, Appendix Table A11 shows that the IVOL coefficient per trading day is -0.0483 percentage points ($t = -2.44$) around earnings announcements and 0.0165 ($t = 1.62$) on other days, a significant difference of -0.0648 ($t = -3.65$). Although the announcement periods represent fewer than five percent of trading days, they make a disproportionately negative contribution to the monthly IVOL coefficient. In the decomposition of compounded monthly excess returns described in Appendix C, the contribution from these days is -0.0979 ($t = -5.59$), compared with a full monthly coefficient of -0.2470, and it is the only part of that monthly relation estimated with comparable precision.

Within the sample for which uncertainty and disagreement can both be observed, Table 9 compares specifications estimated on the same observations for each stock and month. After controlling for the standard firm characteristics, including recent return reversal, the IVOL coefficient around earnings announcements is -0.0611 percentage points per day ($t = -3.04$). Adding disagreement reduces its magnitude to -0.0322

Table 9: IVOL, uncertainty, and disagreement around earnings announcements

	(1)	(2)	(3)	(4)
Panel A: Earnings-window days				
IVOL	-0.0611^{***}	-0.0322		
	(-3.04)	(-1.46)		
U			0.0048	0.0051
			(0.20)	(0.21)
D		-0.0729^{***}	-0.0816^{***}	-0.0755^{***}
		(-3.33)	(-3.55)	(-3.31)
$U \times D$				-0.0587^{***}
				(-4.11)
Panel B: Other trading days				
IVOL	0.0025	0.0020		
	(0.29)	(0.24)		
U			0.0071	0.0061
			(0.85)	(0.73)
D		0.0016	0.0004	0.0005
		(0.35)	(0.10)	(0.12)
$U \times D$				-0.0004
				(-0.10)
Panel C: Difference				
IVOL	-0.0636^{***}	-0.0343^{*}		
	(-3.31)	(-1.68)		
U			-0.0023	-0.0010
			(-0.11)	(-0.05)
D		-0.0745^{***}	-0.0821^{***}	-0.0761^{***}
		(-3.39)	(-3.53)	(-3.27)
$U \times D$				-0.0583^{***}
				(-3.81)
Characteristics	✓	✓	✓	✓
Months	355	355	355	355
Stocks	7,453	7,453	7,453	7,453
Mean R^2	0.042	0.047	0.046	0.050

Monthly Fama–MacBeth regressions of the average daily return (percentage points) on earnings-window days $[-1, +1]$, on other trading days, and their difference, on predictors measured before the announcement. U is 365-day total uncertainty. U and D enter as standardized within-month ranks and IVOL as a standardized logarithm. All columns use identical stock-months. Mean R^2 refers to Panel C. Table A12 adds lottery characteristics to the IVOL regressions. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

($t = -1.46$), a decline of almost one half, while disagreement itself has a coefficient of -0.0729 ($t = -3.33$). The IVOL coefficient is no longer statistically distinguishable from zero in the latter specification. The economic interpretation is that stocks with greater prior disagreement account for a considerable part of the low returns associated with IVOL during the brief periods in which earnings information arrives.

We also consider whether familiar alternative explanations account for these announcement returns. The lottery characteristics examined in Appendix Table A12 are expected idiosyncratic skewness and the presence of unusually large positive daily returns, two measures often used to capture a preference for lottery-like stock payoffs. On the common observations, the coefficient on IVOL is -0.0661 ($t = -3.14$) before adding either these characteristics or disagreement. Including both lottery characteristics leaves the coefficient almost unchanged, at -0.0654 ($t = -2.91$), whereas including disagreement reduces it to -0.0389 ($t = -1.70$). The standard characteristics already account for recent reversal. Thus the competing variables considered here do not explain the concentration of the IVOL anomaly around earnings announcements nearly as well as disagreement does, even though they may have explanatory power for other aspects of the anomaly. The conclusion is also similar when IVOL is measured from weekly returns over the preceding three years: adding disagreement reduces the announcement coefficient from -0.0403 ($t = -1.96$) to -0.0082.

The earnings news itself provides a further way to examine the interpretation of these returns. Using the difference between reported earnings and the latest analyst consensus before the announcement, scaled by the stock price, Appendix Table A13 shows that greater prior disagreement predicts more disappointing earnings, with a coefficient of -0.038 ($t = -3.60$). Uncertainty also predicts more negative surprises in this specification, so the sign of the surprise coefficient alone does not distinguish disagreement from uncertainty. More revealing is the relation between prior disagreement and the stock's return after taking the realized earnings surprise into account. In regressions of the three-day market-adjusted announcement return, the disagreement coefficient is -0.201 ($t = -3.93$) before controlling for the surprise and remains -0.165 ($t = -3.13$) after it is included. Allowing the relation between the surprise and announcement returns to vary across surprise deciles also leaves a negative disagreement coefficient of -0.213 ($t = -3.44$). The uncertainty coefficient is less stable across these specifications, whereas the disagreement coefficient remains negative and statistically significant.

The announcement results complement the cross-sectional evidence in two respects. They show that the negative return relation associated with disagreement is concentrated around the arrival of public information, as the learning argument would suggest, and they

establish that disagreement accounts for a substantial part of the IVOL anomaly during precisely those periods. These comparisons show that lottery characteristics explain little of the IVOL relation around announcements, while disagreement continues to predict returns after conditioning on realized earnings surprises.

5.6 Alternative Explanations of the IVOL Puzzle

The preceding evidence connects the negative IVOL return relation to disagreement, and shows that the relation varies with ownership, borrowing costs, and the arrival of public information in ways suggested by the model. The IVOL puzzle has also motivated a substantial literature emphasizing other sources of predictable returns, including preferences for lottery-like payoffs, reversal and illiquidity, and differences in the ease with which investors can act on overpricing and underpricing. These accounts need not be mutually exclusive with the mechanism developed here. We next ask whether these alternative characteristics account for the conditional patterns in returns associated with IVOL and disagreement.

One explanation is that investors are willing to pay more for stocks that offer a small probability of an unusually large gain, so that such stocks earn low subsequent returns. Boyer et al. (2010) examine expected idiosyncratic skewness, while Bali et al. (2011) use the largest recent daily stock return, commonly called MAX, to capture this preference. Since these characteristics can be correlated with volatility and with dispersion in analysts' forecasts, it is useful to ask whether they account for the negative return relation associated with disagreement. Appendix Table A14 makes that comparison on identical observations. The disagreement coefficient is -0.158 ($t = -2.55$) with uncertainty and the standard firm characteristics, and remains between -0.153 and -0.161 as the lottery characteristics are introduced, with a coefficient of -0.155 ($t = -2.50$) when both are included. The coefficient on MAX remains negative and statistically significant in the latter regression, at -0.158 ($t = -3.65$), so the lottery measures contain return information of their own without absorbing the relation associated with disagreement. Around earnings announcements, however, they leave the negative IVOL coefficient almost unchanged, whereas including disagreement substantially reduces it, as shown in Appendix Table A12. Lottery preferences may therefore help explain some of the unconditional IVOL anomaly, but the measures examined here do little to explain the concentration of the anomaly around earnings announcements.

Short-term reversal and illiquidity provide other explanations for why stocks with volatile recent returns may subsequently perform poorly (Huang et al., 2010; Han and

Lesmond, 2011). The principal regressions already control for one-month reversal, and Appendix Table A15 examines the further role of the bid–ask spread, Amihud illiquidity, and the frequency of days without trading. These controls have little effect on the interactions involving implied borrowing cost. Across these specifications, the disagreement interaction remains between -0.102 and -0.106, while the IVOL interaction remains between -0.137 and -0.145. When borrowing cost is first adjusted for its cross-sectional association with size and a composite of the liquidity measures, its interaction with IVOL remains -0.115 ($t = -4.88$). The borrowing cost result is therefore not simply a reflection of the observed variation in size and liquidity, although borrowing conditions and the ease of trading are naturally related in the data.

The asymmetry between correcting overpricing and underpricing is especially relevant because it also depends on the possibility of short selling. Stambaugh et al. (2015) classify stocks using a collection of return anomalies and find that the negative IVOL relation is concentrated among stocks identified as relatively overpriced. Their result is present in our data over the period used in their study. In Appendix Table A17, the value-weighted return spread between high and low IVOL stocks is 0.546 percentage points ($t = 1.85$) among the most underpriced firms and -1.026 ($t = -2.86$) among the most overpriced. We then examine whether uncertainty and IVOL have different return implications within these mispricing groups.

The two return relations remain distinct even among stocks classified as most underpriced. There, the coefficient on option-implied uncertainty is positive, at 0.426 ($t = 2.42$), even when IVOL is included, while the IVOL coefficient in the same regression is -0.340 ($t = -2.51$). The uncertainty coefficient remains 0.451 ($t = 2.68$) after realized volatility is added (Appendix Table A16). These estimates show that a positive return relation associated with uncertainty and a negative relation associated with IVOL can coexist even within a group formed using an independent measure of apparent mispricing. The Stambaugh–Yu–Yuan classification therefore identifies an important asymmetry, but does not eliminate the distinction between uncertainty and the return variation associated with observed IVOL. The construction of the mispricing groups is described in Appendix D.5.

A related approach is to combine several characteristics associated with differences of opinion into a single measure. Goulding et al. (2026) construct such a measure, denoted DIS, using analyst dispersion, IVOL, short interest, expected skewness, and other variables. Because IVOL is one of its inputs, we compare DIS with uncertainty and disagreement in return regressions rather than using it to explain IVOL itself. Table 10 shows that DIS predicts lower subsequent returns, with a coefficient of -0.204 ($t = -1.95$) when

it enters alone and -0.157 ($t = -2.81$) after uncertainty and disagreement are added. The interaction between uncertainty and disagreement remains negative in the latter specification, at -0.090 ($t = -2.38$). Moreover, when we allow the return coefficients on the belief measures to vary with borrowing cost and size, the interaction of disagreement with borrowing cost remains -0.068 ($t = -2.44$) after DIS is included.

Table 10: Uncertainty, disagreement, and the composite DIS

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
U		-0.096	0.010	-0.082	0.018
		(-0.56)	(0.06)	(-0.48)	(0.11)
D		-0.110	-0.083	-0.110	-0.084
		(-1.58)	(-1.14)	(-1.62)	(-1.19)
$U \times D$		-0.084^{**}	-0.090^{**}	-0.040	-0.047
		(-2.20)	(-2.38)	(-0.84)	(-0.99)
DIS	-0.204^{*}		-0.157^{***}		-0.151^{***}
	(-1.95)		(-2.81)		(-2.72)
$U \times B$				-0.036^{*}	-0.037^{*}
				(-1.71)	(-1.73)
$D \times B$				-0.071^{**}	-0.068^{**}
				(-2.52)	(-2.44)
$U \times S$				-0.004	0.000
				(-0.10)	(0.01)
$D \times S$				0.001	-0.002
				(0.04)	(-0.05)
Characteristics	✓	✓	✓	✓	✓
B and S main effects	✗	✗	✗	✓	✓
Months	324	324	324	324	324
Stocks	6,663	6,663	6,663	6,651	6,651
Mean R^2	0.087	0.103	0.105	0.111	0.113

Monthly Fama–MacBeth regressions of next-month returns (percentage points). U, D, DIS, implied borrowing cost B, and size S are standardized within-month ranks. DIS is the released composite measure of Goulding et al. (2026), whose construction includes IVOL. U is 365-day total uncertainty. All regressions include the standard characteristics. Table A18 adds interactions of DIS with B and S. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Appendix Table A18 allows the return coefficient on DIS itself to vary with borrowing cost, so that the overlap between the two measures can be examined more directly. The interaction between DIS and borrowing cost is -0.079 ($t = -3.23$) before the interactions involving the belief measures are added, but falls to -0.033 ($t = -0.97$) afterward. The interaction between disagreement and borrowing cost remains negative at -0.063 ($t =$

−2.13). Thus much of the variation in the composite measure's return relation across borrowing conditions is shared with the separately observed disagreement measure, although DIS continues to contain information about returns that is not captured by that interaction.

Finally, the positive conditional coefficient on option-implied uncertainty is relevant to explanations that emphasize compensation for expected risk rather than realized volatility. Fu (2009) studies expected idiosyncratic volatility, while Chen and Petkova (2012) and Herskovic et al. (2016) connect volatility measured at the firm level with priced aggregate risks. The option measure used here is forward-looking, and its return coefficient can be positive even where the coefficient on realized IVOL is negative. This distinction is especially visible once disagreement and the trading conditions governing its effect on prices are taken into account.

The comparisons therefore suggest overlap among several explanations without reducing them to the same mechanism. The lottery characteristics retain a negative return relation of their own, while failing to explain the pronounced IVOL effect around earnings announcements. The borrowing cost interactions survive controls for reversal and liquidity, and the mispricing classification leaves room for uncertainty and IVOL to have opposite return relations within the same group. Finally, the conditional evidence for disagreement persists alongside the composite measure of differences of opinion. Taken together with the ownership and announcement results, these comparisons support the interpretation that uncertainty and disagreement, and the trading conditions under which investors act on them, explain an important part of the IVOL puzzle that is not captured by the alternatives considered here.

6 Conclusion

The negative relation between idiosyncratic volatility and subsequent stock returns presents a longstanding challenge to the conventional understanding of risk compensation. We have approached this puzzle by examining how greater fundamental volatility affects equilibrium asset prices when investors possess different information about a common payoff and face constraints on trading. Greater fundamental volatility increases not only the uncertainty investors perceive about the eventual payoff, but also the disagreement across their valuations. While uncertainty discourages investors from bearing risk, disagreement can raise prices when pessimistic positions are constrained. The effect of greater volatility on expected returns therefore depends on how these two consequences of fundamental risk interact in equilibrium.

We begin by establishing restrictions on the relationship between uncertainty and disagreement that follow from rational Bayesian learning. Under a common prior, expected disagreement is bounded by expected posterior uncertainty, with a sharper bound that also depends on the prior variance of the payoff. In a Gaussian learning model, greater fundamental volatility increases both quantities but raises disagreement proportionately more, while additional common information reduces disagreement more rapidly than uncertainty. These results provide a systematic account of how rational beliefs respond to changes in fundamental risk and information, and their predictions find support in the Survey of Professional Forecasters.

We then examine how uncertainty and disagreement affect prices when investors also learn from the information those prices convey. This introduces a substantial theoretical difficulty, since investors infer information from a price that must simultaneously clear the market when their desired positions are constrained. We characterize the resulting nonlinear rational expectations equilibrium exactly, allowing for risk-averse investors with heterogeneous private information and continuously varying demands subject to short sale constraints. The solution shows how these constraints affect the information conveyed by prices, how much disagreement survives investors' learning from prices, and how the resulting beliefs enter equilibrium valuations. In particular, the breadth of ownership determines whether greater fundamental volatility raises or lowers expected prices. When relatively few investors hold the asset, greater disagreement can raise valuations by more than the additional uncertainty reduces them, even though investors continue to require compensation for bearing risk.

Indeed, we find that the comparative statics derived from the model are closely reflected in U.S. stock return data. Uncertainty and disagreement have opposing implications for subsequent returns, while the negative relations associated with disagreement and IVOL become more pronounced when short selling is costly and less pronounced as ownership broadens. When we examine these conditions jointly, both IVOL and disagreement predict substantially lower returns among stocks that are narrowly held and expensive to short, with the interaction between ownership and borrowing costs also statistically significant for IVOL. The timing of these returns provides further evidence of the mechanism. The negative relation associated with disagreement is concentrated around earnings announcements, when substantial common information becomes available, and disagreement accounts for much of the IVOL anomaly realized during these periods. The measures of lottery demand examined here do considerably less to account for this concentration of negative returns, after controlling for recent return reversal.

These findings suggest that the IVOL puzzle is best understood not as a failure of conventional risk compensation, but as a consequence of how fundamental risk changes investors' beliefs and how those beliefs are reflected in market prices. Greater volatility can simultaneously increase the compensation investors require for bearing uncertainty and raise valuations through disagreement, with the balance between these effects governed by information, ownership, and trading constraints. The negative average IVOL coefficient conceals this variation across trading environments and the arrival of information. Once these distinctions are taken into account, the anomaly emerges as a conditional phenomenon whose principal features can be understood through rational learning and constrained trading.

A Theory

A.1 Proofs and Equilibrium Derivations

A.1.1 Common Prior Projection Bound

For the formal argument, draw an investor index independently from the cross-sectional population measure and let $\mathcal{H}$ be the sigma field generated by that investor's identity and information, so that $m = \mathbb{E}[X \mid \mathcal{H}]$. Let $\mathcal{G}$ contain the realized payoff and any other common conditioning variables, but not the information specific to an individual investor. Then

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

Proof of Theorem 1. Center X without loss of generality, and let $P_{\mathcal{G}} m = \mathbb{E}[m \mid \mathcal{G}]$ denote the L^2 projection of m onto $L^2(\mathcal{G})$. The law of total variance and iterated expectations give

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

Moreover,

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

Because $\sigma(X) \subseteq \mathcal{G}$, the one dimensional span of X is contained in $L^2(\mathcal{G})$. The residual from projecting m onto $L^2(\mathcal{G})$ therefore cannot exceed the residual from projecting it onto X. The latter projection is

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

so

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

Since $0 \leq U \leq V$, the second inequality in Theorem 1 follows. For the nested subspaces $\mathrm{span}\{X\} \subseteq L^2(\mathcal{G})$, the Pythagorean identity gives equality if and only if $P_{\mathcal{G}} m = P_X m$. When $\mathcal{G} = \sigma(X)$, this is equivalent, after centering, to

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

Joint Gaussianity of (m, X) is sufficient for this condition to hold. □

A.1.2 Conditioning on Common Information

Let $\mathcal{F}$ be public information contained in both $\mathcal{H}$ and $\mathcal{G}$, and set $\mu_F = \mathbb{E}[X \mid \mathcal{F}]$, $A = \mathbb{E}[m \mid \mathcal{G}]$, and $C = \mathbb{E}\mathrm{Var}(X \mid \mathcal{F})$. If $C > 0$, residualizing the projection argument gives

$$D = U\left(1 - \frac{U}{C}\right) - \mathcal{R}, \qquad \mathcal{R} = \mathbb{E}\left[\left\{A - \mu_F - \left(1 - \frac{U}{C}\right)(X - \mu_F)\right\}^2\right] \geq 0. \tag{13}$$

For $Z = X - \mu_F$ and $M = m - \mu_F$, projection gives $\mathbb{E}Z^2 = C$, $\mathbb{E}M^2 = \mathbb{E}MZ = C - U$, and $D = C - U - \mathbb{E}(A - \mu_F)^2$. Projecting $A - \mu_F$ onto Z then yields (13), whereas when $C = 0$, all three moments vanish. For any $0 < U < C$, a Gaussian public posterior with constant variance C and a private Gaussian error of variance $UC/(C - U)$ attains equality. The gap is also the additional consensus forecast error:

$$\mathbb{E}(X - A)^2 = U - D = U^2/C + \mathcal{R}.$$

Thus, equality requires a linear aggregate forecast, which may arise without Gaussian beliefs.

For example, allow investors to have an arbitrary fixed distribution of private Gaussian precisions τ_i, independently of the payoff, while the public posterior is Gaussian with constant variance C. Define $\alpha_i = C\tau_i/(1 + C\tau_i)$ and $\bar{\alpha} = \mathbb{E}_i\alpha_i$. Averaging posterior variances gives $U = C(1 - \bar{\alpha})$, while the consensus is $\mu_F + \bar{\alpha}(X - \mu_F)$, so

$$D = C\bar{\alpha}(1 - \bar{\alpha}), \qquad D/U = \bar{\alpha}.$$

As C rises, both U and $\bar{\alpha}$ increase, so $D = U\bar{\alpha}$ increases as well. The three ex ante comparisons continue to hold under arbitrary heterogeneity in private precision when the public posterior is Gaussian with constant variance. Statewise disagreement can nevertheless vary with the realized payoff because investors with different precisions respond differently to the same innovation.

A.1.3 Convex Order and Constrained Demand

Write $A_1 \leq_{cx} A_2$ if expectations of convex functions are weakly larger under A_2, including equality of means. For equal mean integrable desired positions with distribution functions F_1, F_2, the demand difference at a common floor is

$$\mathbb{E}\max\{A_2, -\ell\} - \mathbb{E}\max\{A_1, -\ell\} = \int_{-\infty}^{-\ell} \{F_2(z) - F_1(z)\}\, dz.$$

The identity follows by writing $\max\{A, -\ell\} = A + (-\ell - A)^+$ and using the equality of means. Under convex order, the integral is nonnegative because $(-\ell - A)^+$ is a convex function of A. It is positive under the strictness condition of the lemma, which proves Lemma 1.

The location–scale ownership result also extends to finite shorting limits. Suppose $M = \bar{m} + cUZ$, with $c > 0$ and a continuous distribution of Z with mean zero, unit variance, and unique median zero, so that $c = \sqrt{D}/U$. Define $d = (\bar{m} - P + \gamma U\ell)/(cU)$. Aggregate demand above the floor is proportional to $\mathbb{E}(d + Z)^+$, so clearing fixes d through $c\mathbb{E}(d + Z)^+ = \gamma(S + \ell)$, giving

$$P = \bar{m} + U(\gamma\ell - cd), \qquad P_U = (P - \bar{m})/U.$$

Positive ownership, rather than the fraction above the shorting floor, is $\mathbb{P}(Z > \gamma\ell/c - d)$, so $P_U > 0$ precisely when that fraction is below one half.

For Gaussian Z, write $h(d) = \mathbb{E}(d + Z)^+$ as in the main text and define the expected excess payoff relative to price by $r = \bar{m} - P = U(cd - \gamma\ell)$. Market clearing gives $ch(d) = \gamma(S + \ell)$, so d depends on c but not on U. Differentiating the clearing condition yields

$$\frac{\partial^2 r}{\partial U\, \partial c} = d + c\frac{\partial d}{\partial c} = -\frac{\phi(d)}{\Phi(d)} < 0.$$

Thus, greater uncertainty magnifies the negative return effect associated with a given degree of relative disagreement $c = \sqrt{D}/U$. The magnitude is larger when the ownership cutoff is more selective, since the inverse Mills ratio rises as d falls.

A.1.4 The Equilibrium with Stochastic Supply

Proof of Proposition 2. The proof proceeds in three steps: we construct the statistic Y_ℓ implied by market clearing, show that the proposed price is an invertible function of that statistic, and verify that each investor's position is optimal given the information conveyed by price. Since $h' = \Phi > 0$ and h maps the real line onto $(0, \infty)$, the clearing equation has a unique solution for $d_\ell(t)$ at every feasible supply realization, and (9) then determines $w_\ell(t)$. The map w_ℓ is strictly decreasing, so the support of the induced noise is

$$I_\ell = [L_\ell, H_\ell] = [w_\ell(\bar{t}), w_\ell(\underline{t})].$$

Changing variables from uniform supply T to $W_\ell = w_\ell(T)$ gives the density

$$f_\ell(w) = \frac{\Phi((\gamma b \ell - w)/\sqrt{b})}{\gamma b(\overline{t} - \underline{t})} \mathbf{1}_{I_\ell}(w).$$

The normal distribution function is log concave, so f_ℓ is log concave on its interval support.

To verify that price reveals the proposed signal, define

$$k(w) = e^{-w^2/(2b)}, \qquad K_\ell = \int f_\ell(w)k(w)\,dw, \qquad g_\ell(w) = f_\ell(w)k(w)/K_\ell.$$

The posterior of the investor with private signal $s_i = y$, conditional on $Y_\ell = y$, is proportional to $g_\ell(y - x)\,dF_X(x)$, and its mean is the proposed price $p_\ell(y)$. Because multiplication by k makes g_ℓ strictly log concave, the likelihood ratio of two such posteriors increases strictly with x as y rises, wherever their supports overlap. Outside that overlap, payoff values leave the posterior support at its lower end and enter at its upper end. The conditional payoff distribution is atomless and nondegenerate at every feasible signal with a positive normalizing integral, so its mean increases strictly with y.

Continuity follows by dominated convergence on each bounded neighborhood of y. The support restriction places the conditional payoff in $[y - H_\ell, y - L_\ell]$, and atomlessness assigns zero probability to its moving endpoints. Thus, p_ℓ is continuous and strictly increasing on the feasible signal set, with a measurable inverse, so observing price conveys exactly the same information as observing Y_ℓ.

To verify demand, note that given (s_i, y), investor i has a posterior proportional to

$$f_\ell(y - x) \exp\left[-\frac{(s_i - x)^2}{2b}\right] dF_X(x).$$

For CARA utility, the first order condition for an unrestricted position a_i at price p is

$$\frac{\mathbb{E}[Xe^{-\gamma a_i X} \mid s_i, y]}{\mathbb{E}[e^{-\gamma a_i X} \mid s_i, y]} = p.$$

With $a_i = (s_i - y)/(\gamma b)$, the factor by which marginal utility reweights the posterior satisfies

$$e^{-\gamma a_i x} \exp\left[-\frac{(s_i - x)^2}{2b}\right] = \exp\left[\frac{y^2 - s_i^2}{2b}\right] \exp\left[-\frac{(y - x)^2}{2b}\right].$$

The first term on the right is independent of x and cancels from the ratio in the first order condition. The remaining distribution is precisely the posterior of the investor with $s_i = y$, whose mean is $p_\ell(y)$, so the proposed position satisfies the condition.

The conditional payoff is bounded, so the required exponential moments exist, and nondegeneracy makes expected utility strictly concave in the position. The stationary point is therefore the unique unrestricted optimum, and imposing the lower bound gives $q_i = \max\{a_i, -\ell\}$.

Finally, average these optimal positions over private signals at the realized payoff. Since $s_i = X + \epsilon_i$ and the cross section of errors is Gaussian, aggregate demand is

$$\mathbb{E}_i q_i = \frac{h((X - y + \gamma b \ell)/\sqrt{b})}{\gamma\sqrt{b}} - \ell = T.$$

The last equality is the clearing equation used to construct $y = X + w_\ell(T)$. The optimal positions therefore clear the market, and the posterior used to choose those positions is consistent with the clearing price. Moreover, $s_i - Y_\ell = \epsilon_i - w_\ell(T)$, so the distribution of holdings is independent of the payoff law. The bound $|p_\ell(Y_\ell) - Y_\ell| \leq \max\{|L_\ell|, |H_\ell|\}$ and the finite second moment of X imply a finite second moment for price, which completes the verification of the proposed equilibrium. □

With uniform supply, conditioning on the price index restricts feasible payoff values to a compact interval around that index. The proof therefore accommodates any atomless, nondegenerate payoff prior with finite second moment, including priors with disconnected or unbounded support and the Gaussian and log concave cases used below.

For a supply law other than the uniform distribution, take primitive $T > -\ell$ independent of the payoff and private errors, form $W = w_\ell(T)$, and let f denote its density on an interval. It is sufficient that $g(w) \propto f(w)e^{-w^2/(2b)}$ be strictly log concave, that the relevant posteriors be nondegenerate, and that the conditional exponential moments required for optimality exist. Log-concavity of f is a stronger sufficient condition, and for compact induced noise, the required exponential moments are automatic. Gaussian induced noise also satisfies the conditions in the Gaussian payoff specialization below.

One sufficient condition on primitive supply is that its density $\rho(t)$ be positive, nonincreasing, and log concave on the feasible interval. Inverting the supply map gives

$$d(w) = \frac{\gamma b \ell - w}{\sqrt{b}}, \qquad t(w) = \frac{h(d(w))}{\gamma\sqrt{b}} - \ell, \qquad f(w) = \rho(t(w))\frac{\Phi(d(w))}{\gamma b}.$$

Here, $t'(w) < 0$ and $t''(w) > 0$, so a concave nonincreasing $\log\rho$ composed with the convex $t(w)$ remains concave, and adding the concave $\log\Phi(d(w))$ and the strictly concave $-w^2/(2b)$ proves the condition on g. Uniform supply satisfies this sufficient condition, but

the argument only requires concavity of the combined log density and does not impose a linear price rule.

At the optimum of every investor above the floor, marginal utility weights the posterior by $e^{-\gamma a_i X}$. The signal translation identity in the proof makes all these weighted distributions identical to that of the investor with $s_i = y$, so they share the common variance $U_0(p_\ell(y))$, even though individual posterior variances may differ. Holding the public posterior ν fixed,

$$\frac{\partial}{\partial y}\frac{\int x e^{-(y-x)^2/(2b)}\nu(dx)}{\int e^{-(y-x)^2/(2b)}\nu(dx)} = \frac{U_0}{b}. \tag{14}$$

The pricing derivatives therefore depend on the variance under the marginal investor's weighted posterior.

A.1.5 Short Selling and Information in Price

Differentiating the supply map gives

$$w_t = -\frac{\gamma b}{\Phi(d_\ell(t))}, \qquad \partial_\ell w_\ell(t) = -\frac{\gamma b[1-\Phi(d_\ell(t))]}{\Phi(d_\ell(t))}, \qquad w_{t\ell} = \frac{\gamma^2 b^{3/2}\phi(d_\ell(t))}{\Phi(d_\ell(t))^3} > 0.$$

Relaxing the shorting limit therefore increases the fraction of investors above the floor and contracts the noise interval between any two fixed supply quantiles. Letting $\ell \to \infty$ recovers the unrestricted signal, as derived in Appendix A.1.7.

To prove the uncertainty comparisons, let A denote the higher capacity experiment and B the lower capacity experiment. Write F_A, F_B for their noise distribution functions. At marginal signal quantiles y_A, y_B of the same rank u, contraction of all noise quantile intervals implies that

$$\Xi(x) = F_A(y_A - x) - F_B(y_B - x)$$

crosses zero at most once, from positive to negative. Since both signals have marginal rank u, $\mathbb{E}\Xi(X) = 0$, and therefore $\mathbb{E}[X\Xi(X)] \le 0$. Log-concavity of the noise densities makes each public posterior mean nondecreasing in its signal. If Q_A, Q_B denote the quantile functions of those posterior means, their integrated lower quantiles satisfy

$$\int_0^u Q_A(v)\,dv = \mathbb{E}[XF_A(y_A - X)] \le \mathbb{E}[XF_B(y_B - X)] = \int_0^u Q_B(v)\,dv.$$

Because the two posterior means have the same expectation, the displayed inequalities place the posterior mean under A above that under B in convex order: more information produces greater variation in the conditional expectation of the payoff. In particular, the

variance of the posterior mean is weakly larger under A, and writing $C_j = \mathbb{E}[\mathrm{Var}(X \mid Y_j)]$ for $j = A, B$, the law of total variance gives $C_A \leq C_B$. The same proof applies under the conditional payoff law given almost every private signal, because the noise remains independent of (X, s_i), and averaging over the private signal yields $U_A \leq U_B$. The argument compares location experiments in the sense of Lehmann (1988) and establishes the ordering of posterior uncertainty under squared error loss.

A.1.6 Uncertainty and Disagreement after Learning from Price

Let $m(s, y) = \mathbb{E}[X \mid s, y]$ and $v(s, y) = \mathrm{Var}(X \mid s, y)$. Throughout this subsection,

$$U = \mathbb{E}[v(s_i, Y)], \qquad D = \mathbb{E}[\mathrm{Var}_i(m(s_i, Y) \mid X, Y)]. \tag{15}$$

By conditioning on the payoff and public signal before averaging across states, D measures variation across investors rather than variation in the public posterior mean. We also write $C = \mathbb{E}[\mathrm{Var}(X \mid Y)]$ for uncertainty after public information alone. Differentiating each investor's forecast with respect to the private signal then relates this disagreement to posterior uncertainty without requiring Gaussian beliefs.

Lemma 2 (Posterior variance bounds for disagreement)**.** *Let public information Y be independent of the Gaussian private error conditional on X. Let $m(s, y) = \mathbb{E}[X \mid s, y]$, $v(s, y) = \mathrm{Var}(X \mid s, y)$, and define U, D by* (15)*. Whenever the moments and derivatives below are integrable,*

$$\frac{U^2}{b} \;\leq\; D \;\leq\; \frac{\mathbb{E}[v(s, Y)^2]}{b}.$$

Proof. Differentiating the Gaussian private signal likelihood gives the exact identity

$$\partial_s m(s, y) = \frac{v(s, y)}{b}.$$

At any fixed common state $(X, Y) = (x, y)$, Gaussian integration by parts and the Cauchy–Schwarz inequality together imply

$$\mathrm{Cov}_\epsilon\{m(x + \epsilon, y), \epsilon\} = \mathbb{E}_\epsilon v(x + \epsilon, y), \qquad \mathrm{Var}_\epsilon m(x + \epsilon, y) \geq \frac{\{\mathbb{E}_\epsilon v(x + \epsilon, y)\}^2}{b}.$$

The conditional variance on the left is disagreement at the common state (x, y). Averaging this inequality over (X, Y) and applying Jensen's inequality to the squared conditional

mean gives $D \geq U^2/b$. For the upper bound, the Gaussian Poincaré inequality bounds the variance of a function of the private error by the expected square of its derivative:

$$\mathrm{Var}_\epsilon m(x+\epsilon, y) \leq b\mathbb{E}_\epsilon\{\partial_s m(x+\epsilon, y)\}^2 = \frac{\mathbb{E}_\epsilon v(x+\epsilon, y)^2}{b}.$$

Averaging over the common state gives $D \leq \mathbb{E}[v(s, Y)^2]/b$, as required. □

The bounds coincide when posterior variance is constant across investors and common states, and they become asymptotically tight when residual posterior variance is nearly constant, a property used below as prices become increasingly informative.

Small Supply Disturbances. Around a fixed supply level $t > 0$, consider the independent disturbance

$$T_\delta = t + \delta Z_T, \qquad Z_T \sim \mathrm{Unif}[-1, 1], \qquad 0 < \delta < t.$$

At each fixed δ, both shorting limits are compared under the same supply law, without imposing Gaussian noise on the signal revealed by price. Let

$$\Omega_\ell^{\mathrm{int}}(t) = \Phi\Big(h^{-1}(\gamma\sqrt{b}(t+\ell))\Big), \qquad K_\ell(t) = \frac{\gamma b}{\Omega_\ell^{\mathrm{int}}(t)} = -w_\ell'(t).$$

Proposition 4 (Belief moments with small supply disturbances)**.** *Suppose X has a finite second moment and a density that is positive and continuous at almost every realization of X. For any fixed finite capacity ℓ,*

$$C_{\ell,\delta} = \frac{K_\ell(t)^2\delta^2}{3} + o(\delta^2), \qquad U_{\ell,\delta} = \frac{K_\ell(t)^2\delta^2}{3} + o(\delta^2), \tag{16}$$

$$D_{\ell,\delta} = \frac{K_\ell(t)^4\delta^4}{9b} + o(\delta^4), \qquad \frac{D_{\ell,\delta}}{U_{\ell,\delta}} = \frac{K_\ell(t)^2\delta^2}{3b} + o(\delta^2). \tag{17}$$

Consequently, for any fixed $0 \leq \ell_1 < \ell_2 < \infty$, all four quantities $C, U, D, D/U$ are strictly lower at ℓ_2 for every sufficiently small positive δ.

Proof. Couple economies by the same (X, Z_T, ϵ) and write $Y = X + w_\ell(t + \delta Z_T)$, $s = X + \epsilon$. Conditional on (Y, s), the posterior density of a candidate $z \in [-1, 1]$ is proportional to

$$\pi\{Y - w_\ell(t+\delta z)\} \exp\left[-\frac{\{s - Y + w_\ell(t+\delta z)\}^2}{2b}\right].$$

For almost every true realization, both factors converge uniformly in z to positive constants as $\delta \downarrow 0$, so the posterior of z converges to the uniform distribution, and omitting the private likelihood proves the same assertion for the public posterior.

The uniform Taylor expansion

$$w_\ell(t+\delta z) = w_\ell(t) + \delta w'_\ell(t) z + O(\delta^2)$$

then gives $v(s,Y)/\delta^2 \to K_\ell(t)^2/3$ almost surely, and the same limit for public conditional variance. Because the conditional support of X has length $w_\ell(t-\delta) - w_\ell(t+\delta) = O(\delta)$, both conditional variances are bounded by a constant times δ^2, uniformly over realizations, and dominated convergence gives (16) and

$$\frac{\mathbb{E}v(s,Y)^2}{\delta^4} \longrightarrow \frac{K_\ell(t)^4}{9}.$$

The inequalities in Lemma 2 squeeze D/δ^4 to $K_\ell(t)^4/(9b)$, proving (17). Finally, $\Omega^{\text{int}}_\ell(t)$ increases strictly with ℓ, so $K_{\ell_2}(t) < K_{\ell_1}(t)$ and the leading coefficients are strictly ordered, which means that each asserted finite comparison holds for sufficiently small, nonzero δ. □

Even with an asymmetric or multimodal prior, a sufficiently informative price leaves investors with nearly identical small residual uncertainty. The comparison therefore does not rely on Gaussian posteriors, and for two fixed shorting limits it gives

$$\frac{U_{\ell_2,\delta}}{U_{\ell_1,\delta}} \longrightarrow \left(\frac{\Omega^{\text{int}}_{\ell_1}(t)}{\Omega^{\text{int}}_{\ell_2}(t)}\right)^2, \qquad \frac{D_{\ell_2,\delta}}{D_{\ell_1,\delta}} \longrightarrow \left(\frac{\Omega^{\text{int}}_{\ell_1}(t)}{\Omega^{\text{int}}_{\ell_2}(t)}\right)^4.$$

With fewer investors able to respond, disagreement is amplified proportionately more than uncertainty. For every fixed pair of shorting limits, the strict comparisons therefore hold once supply uncertainty is sufficiently small.

Gaussian Fundamentals and Nonlinear Price Learning.

Proposition 5 (Nonlinear price learning under a Gaussian payoff prior)**.** *Let $X \sim N(\mu, a)$ and retain the exact uniform supply equilibrium at any finite capacity. Put $u_* = ab/(a+b)$, the posterior variance from a private signal alone. Then*

$$\frac{U^2}{b} \le D \le \frac{u_*}{b} U \le \frac{u_*^2}{b}, \qquad \frac{D}{U} \le \frac{a}{a+b}. \tag{18}$$

With private information alone, uncertainty is u_*, disagreement is u_*^2/b, and their ratio is $u_*/b = a/(a+b)$. With nondegenerate finite uniform supply, the upper bounds are strict.

Proof. The effective noise density is log concave, and the posterior of an individual investor is proportional to

$$\exp\left[-\frac{(x-\mu)^2}{2a} - \frac{(s-x)^2}{2b}\right] f_\ell(y-x).$$

Its negative log density has curvature at least $a^{-1}+b^{-1}$ on its interval support, so $v(s,y) \leq u_*$. One direct proof of this scalar variance bound is as follows: if a density $r \propto e^{-V}$ on $[L,H]$ has $V'' \geq \kappa$, an independent copy argument gives $\mathrm{Cov}(Z, V'(Z)) \geq \kappa \mathrm{Var} Z$. Integration by parts gives

$$\mathrm{Cov}(Z, V'(Z)) = 1 - \{(H - \mathbb{E}Z)r(H) - (L - \mathbb{E}Z)r(L)\} \leq 1.$$

Thus, $\mathrm{Var} Z \leq 1/\kappa$, and in the present compact support posterior, the boundary term is strictly positive. Applying Lemma 2 now gives $D \leq \mathbb{E}v^2/b \leq u_* U/b$, which, together with $U \leq u_*$, proves (18). □

The comparison therefore establishes that conditioning on the nonlinear equilibrium price reduces disagreement relative to the private information benchmark.

A.1.7 Limiting Result from Hellwig (1980)

This appendix establishes the convergence stated in Section 3.2 and derives the unrestricted Gaussian economy of Hellwig (1980). The information statistic converges with the original supply law unchanged, whereas the Gaussian posterior and pricing formulas require an additional distributional specialization. We also construct feasible constrained supply distributions approaching Gaussian supply, keeping this sequence distinct from the comparison that holds the supply law fixed.

Convergence under a fixed supply law. For each fixed supply realization t, equation (9) gives $d_\ell(t) \to \infty$ as $\ell \to \infty$. The Gaussian tail identity

$$h(d) - d = \phi(d) - d\Phi(-d) \longrightarrow 0$$

and the definition of $d_\ell(t)$ imply

$$w_\ell(t) = -\gamma b t + \sqrt{b}\{h(d_\ell(t)) - d_\ell(t)\} \longrightarrow -\gamma b t.$$

The convergence is uniform on compact sets of supply realizations because $h(d) - d$ decreases to zero. Consequently, the public statistic and individual positions converge to

$$Y_\infty = X - \gamma b T, \qquad q_{i,\infty} = \frac{s_i - Y_\infty}{\gamma b} = T + \frac{\epsilon_i}{\gamma b}. \tag{19}$$

Private signal errors average to zero across investors, so these positions absorb the realized supply. This limit retains the uniform supply law of Section 3.2, and does not require Gaussian supply.

Valuation in the unrestricted Gaussian economy. To recover the Gaussian specification of Hellwig (1980), take $X \sim N(\mu_0, a)$ and unrestricted net supply $T \sim N(\bar{T}, \Delta^2)$, with $\Delta^2 > 0$, independently of each other and of private signal errors. Adding the known constant $\gamma b \bar{T}$ to Y_∞ gives the unbiased common signal $\widehat{Y} = Y_\infty + \gamma b \bar{T}$, whose precision is

$$\tau_p^H = (\gamma^2 b^2 \Delta^2)^{-1}.$$

The posterior mean and desired position are

$$m_i = U_H(a^{-1}\mu_0 + b^{-1}s_i + \tau_p^H \widehat{Y}), \qquad q_i = \frac{m_i - P_\infty}{\gamma U_H}, \tag{20}$$

where

$$U_H = (a^{-1} + b^{-1} + \tau_p^H)^{-1}.$$

Averaging posterior means and clearing supply gives

$$P_\infty = \bar{m} - \gamma U_H T = \pi_0 + \pi_X X - \pi_T T,$$

with

$$\pi_0 = U_H(a^{-1}\mu_0 + \tau_p^H \gamma b \bar{T}), \qquad \pi_X = U_H(b^{-1} + \tau_p^H), \qquad \pi_T = \gamma b \pi_X.$$

Since $\pi_X > 0$ and $\pi_T / \pi_X = \gamma b$, observing price is equivalent to observing Y_∞. Substitution into (20) recovers the positions in (19), verifying both market clearing and the information used in investors' decisions.

Conditional on payoff and supply, posterior means differ only through private signal errors. Disagreement and expected price therefore satisfy

$$D_H = \frac{U_H^2}{b}, \qquad \frac{D_H}{U_H} = \frac{U_H}{b}, \qquad \mathbb{E}[P_\infty] = \mu_0 - \gamma U_H \bar{T}.$$

Relative to private information alone, common price information reduces uncertainty and disagreement, with disagreement falling proportionately more. Holding the supply law fixed, greater fundamental variance raises the two belief moments, but when $\bar{T} > 0$ its effect on expected price is

$$\frac{d\mathbb{E}[P_\infty]}{da} = -\gamma \bar{T} \frac{U_H^2}{a^2} < 0.$$

Without the floor, optimistic and pessimistic positions can offset in aggregate, so disagreement continues to affect holdings without independently raising price.

Convergence from feasible constrained economies. A Gaussian supply law is not feasible under a finite common floor because nondegenerate clearing requires $T > -\ell$. To approach the Gaussian economy while respecting this requirement, let $T \sim N(\bar{T}, \Delta^2)$ denote the target unrestricted supply and define

$$T_\ell = \frac{h(\gamma\sqrt{b}\,(T+\ell))}{\gamma\sqrt{b}} - \ell, \qquad w_\ell(T_\ell) = -\gamma b T.$$

For every finite ℓ, $T_\ell > -\ell$, so each approximating economy is feasible and generates the same Gaussian information statistic as the unrestricted economy. Individual holdings are

$$q_{i,\ell} = \max\left\{T + \frac{\epsilon_i}{\gamma b}, -\ell\right\}, \qquad \mathbb{E}_i q_{i,\ell} = T_\ell.$$

Moreover,

$$0 \le T_\ell - T = \frac{h(-\gamma\sqrt{b}\,(T+\ell))}{\gamma\sqrt{b}} \longrightarrow 0,$$

and the difference is dominated by $(-T)^+ + 1/(\gamma\sqrt{2\pi b})$. Dominated convergence gives $T_\ell \to T$ in L^2, while the truncation error in individual holdings is dominated by the negative part of the unrestricted position and therefore also vanishes in L^2. The Gaussian information statistic, and hence its Bayesian pricing function, is unchanged along this sequence, while allocations converge to (19). Each economy in the sequence satisfies the conditions of Appendix A.1.4, because the induced noise $-\gamma b T$ is Gaussian and the effective likelihood is therefore strictly log concave. The argument does not rely on the Gaussian payoff prior. For any atomless, nondegenerate payoff prior with finite second moment, the pricing function along the sequence is that of the unrestricted economy with Gaussian supply, which falls within the class analyzed by Breon-Drish (2015), and the price is affine in the statistic only when the prior is Gaussian.

Other approximating sequences. The limit does not depend on the construction above. Suppose instead that supply in the economy with floor ℓ is $T \sim N(\bar{T}, \Delta^2)$ conditioned on $T > -\ell/2$, so that its density ρ_ℓ is proportional to the Gaussian density on $(-\ell/2, \infty)$. With $d(w)$ and $t(w)$ defined as in Appendix A.1.4, the induced noise has density $f_\ell(w) = \rho_\ell(t(w))\Phi(d(w))/(\gamma b)$ below $w_\ell(-\ell/2)$. Since $t'(w) = -\Phi(d)/(\gamma b)$ and $t''(w) = \phi(d)/(\gamma b^{3/2})$, direct differentiation gives

$$\frac{d^2}{dw^2}\log\left(f_\ell(w)e^{-w^2/(2b)}\right) = -\frac{\Phi(d)^2}{\gamma^2 b^2 \Delta^2} + \frac{(\bar{T} - t)\,\phi(d)}{\gamma b^{3/2}\Delta^2} - \frac{\phi(d)\,h(d)}{b\,\Phi(d)^2} - \frac{1}{b},$$

evaluated at $d = d(w)$ and $t = t(w)$. Only the second term can be positive. On the support, $t > -\ell/2$ and $d > \underline{d}_\ell = h^{-1}(\gamma\sqrt{b}\,\ell/2)$, where $\underline{d}_\ell \geq \gamma\sqrt{b}\,\ell/2 - \phi(0)$ once $\gamma\sqrt{b}\,\ell/2 \geq \phi(0)$, so that term is at most $(\bar{T} + \ell/2)^+\phi(\underline{d}_\ell)/(\gamma b^{3/2}\Delta^2)$, which vanishes as $\ell \to \infty$. For all sufficiently large ℓ, the effective likelihood is therefore strictly log concave, and the construction of Proposition 2 yields an equilibrium under this supply law.

As the floor recedes, $t(w) \to -w/(\gamma b)$ and $\Phi(d(w)) \to 1$ uniformly on compact sets, while ρ_ℓ converges uniformly to the Gaussian density. Because f_ℓ is uniformly bounded and the factor $e^{-w^2/(2b)}$ controls the tails, $f_\ell(w)e^{-w^2/(2b)}$ converges uniformly on the real line to its counterpart under Gaussian induced noise. The payoff prior has a finite first moment, so the reference valuations converge uniformly on compact sets of y to the pricing function of the unrestricted economy with Gaussian supply. The conditioned supply law converges in total variation to $N(\bar{T}, \Delta^2)$, and $w_\ell(t) \to -\gamma b t$ uniformly on compact sets, so the public statistic and individual holdings converge in distribution as before. The truncation point must recede with the floor for this argument, because the second term grows with ℓ when supply can approach $-\ell$.

A.1.8 A Gaussian Specialization

To retain Gaussian inference while prohibiting short sales, set $\ell = 0$ and take $W \sim N(m_W, v)$ independently of X. Define supply by

$$T = \frac{h(-W/\sqrt{b})}{\gamma\sqrt{b}}.$$

This law has strictly positive supply and induces exactly the Gaussian noise W through (9). In the notation below, $a = \tau_0^{-1}$, $b = \tau_x^{-1}$, and $v = \kappa^2\sigma_z^2$, with $W = \kappa z - \sqrt{b}\,d_0$. The public index in the general construction differs from $X + \kappa z$ only by the known constant $-\sqrt{b}\,d_0$, so the two indices convey identical information. Because the supply law does not depend

on fundamental variance a, that variance can change without altering the distribution of asset supply.

Let

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

and let $z \sim N(0, \sigma_z^2)$ be independent. Investors observe (s_i, p), but neither z nor realized supply separately, and for a fixed $\kappa > 0$, price will reveal the noisy public signal

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

Fix $d_0 \in \mathbb{R}$, which indexes the standardized ownership cutoff when $z = 0$. Asset supply is specified by

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

This schedule permits price to be strictly monotone in Y, so observing price is equivalent to observing that scalar signal, while constrained demand clears exactly in every state.

Proposition 6 (Informative price equilibrium)**.** *There is a rational expectations equilibrium with*

$$U = \mathrm{Var}(X \mid s_i, p) = \frac{1}{\tau}, \qquad D = \mathrm{Var}_i(\mathbb{E}[X \mid s_i, p] \mid X, z) = \frac{\tau_x}{\tau^2}, \tag{21}$$
$$\beta_p = \frac{\tau_x + \tau_p}{\tau}, \qquad \alpha = \frac{\tau_0 \mu_0}{\tau} - d_0 \sqrt{D},$$

and price

$$p(X, z) = \alpha + \beta_p (X + \kappa z).$$

Investor i has posterior mean and demand

$$m_i = U(\tau_0 \mu_0 + \tau_x s_i + \tau_p Y), \qquad q_i = \frac{(m_i - p)^+}{\gamma U}. \tag{22}$$

The ownership cutoff in each state is

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

so aggregate demand equals $\mathcal{S}(z)$ in every state.

Proof. Because price is strictly increasing in Y, observing p is equivalent to observing the Gaussian public signal Y, and Bayesian updating gives (21) and (22). 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. The price coefficients imply

$$\bar{m} - p = \sqrt{D}\,[d_0 - \kappa\sqrt{\tau_x}z].$$

Aggregating constrained demand gives

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

The equilibrium makes the information effect of price explicit: without information from price,

$$U^{\text{priv}} = \frac{1}{\tau_0 + \tau_x}, \qquad D^{\text{priv}} = \frac{\tau_x}{(\tau_0 + \tau_x)^2}.$$

With price information, let

$$\rho_p = \frac{\tau_0 + \tau_x}{\tau_0 + \tau_x + \tau_p}.$$

Then

$$U^{\text{price}} = \rho_p U^{\text{priv}}, \qquad D^{\text{price}} = \rho_p^2 D^{\text{priv}}, \qquad \frac{D^{\text{price}}}{U^{\text{price}}} = \rho_p \frac{D^{\text{priv}}}{U^{\text{priv}}}.$$

Equivalently,

$$\frac{D}{U} = \frac{\tau_x}{\tau_0 + \tau_x + \tau_p}. \tag{23}$$

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

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

Proof. Divide the ratio in (23) by its private-information counterpart $\tau_x/(\tau_0 + \tau_x)$. The displayed expression follows immediately, and it is strictly below one whenever $\tau_p > 0$. $\square$

Information contained in price reduces disagreement relative to uncertainty in the posterior beliefs entering demand, after which the position constraint determines how the remaining disagreement is expressed in positions. The general common prior bound

permits D/U to approach one, whereas an informative common price yields the lower ratio in (23).

Price learning leaves the benchmark ownership threshold intact in this specification. Define

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

Expected positive ownership and mean supply are

$$\Omega := \mathbb{E}_z[\Phi(d(z))] = \Phi\left(\frac{d_0}{\sqrt{1+\chi^2}}\right), \tag{24}$$

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

Expected price is

$$\bar{p}(a) = \mu_0 - d_0\sqrt{D} = \mu_0 - d_0\sqrt{\tau_x}\, U, \tag{25}$$

where $a = \tau_0^{-1}$ is prior fundamental variance.

Proposition 7 (Ownership threshold with price learning)**.** *Holding $(\tau_x, \kappa, \sigma_z, d_0, \gamma)$ fixed,*

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

and

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

Thus, greater fundamental variance raises expected price if and only if

$$d_0 < 0 \quad \Longleftrightarrow \quad \Omega < \frac{1}{2} \quad \Longleftrightarrow \quad \bar{T}(d_0) < \bar{T}^*,$$

where

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

Expected price falls when the inequalities are reversed.

Proof. Use $U = (a^{-1} + \tau_x + \tau_p)^{-1}$, $D = \tau_x U^2$, and (25). Equation (24) gives $\Omega < 1/2$ exactly when $d_0 < 0$, and strict monotonicity of h gives the equivalent mean supply threshold. □

A.1.9 Additional Fundamental Risk

Fix all trading primitives and suppress the capacity subscript. Let f be the actual compact noise density, $k(w) = e^{-w^2/(2b)}$, $K = \int fk$, and $g = fk/K$. For the proof, introduce a probability measure $\mathbb{Q}$ under which W has density g and remains independent of the payoff, whose law is unchanged. Define

$$L(w) = \frac{f(w)}{g(w)} = Ke^{w^2/(2b)}, \qquad m(y) = \mathbb{E}_{\mathbb{Q}}[W \mid Y = y], \qquad r(y) = \mathbb{E}_{\mathbb{Q}}[L(W) \mid Y = y].$$

The equilibrium price is $p(y) = y - m(y)$, and the auxiliary measure expresses the marginal investor's valuation as a conditional expectation while keeping the actual supply distribution in the likelihood weight L.

Proof of Proposition 3. Use the same realizations of X_1, W, and Z to define $Y_1 = X_1 + W$ and $Y_2 = Y_1 + Z$ under $\mathbb{Q}$, a construction that preserves the supply noise and changes only the fundamental. For $j = 1, 2$, let m_j and r_j be the conditional expectations defined above under payoff X_j.

The difference between price and realized payoff is $P_j - X_j = W - m_j(Y_j)$. To take its expectation under the actual distribution, multiply by $L(W)$ under $\mathbb{Q}$. Conditioning the second term on Y_j gives

$$\mathbb{E}P_j - \mathbb{E}X_j = \mathbb{E}_{\mathbb{Q}}[WL] - \mathbb{E}_{\mathbb{Q}}[m_j(Y_j)r_j(Y_j)].$$

The first term is identical in the two economies because their supply laws coincide, and because the added risk has mean zero, the unconditional mean payoff is also unchanged.

Since Z is independent of (X_1, W), Y_2 is obtained by adding independent noise to Y_1, so conditional on Y_1, observing Y_2 provides no further information about W. Iterated expectations therefore imply

$$m_2(Y_2) = \mathbb{E}_{\mathbb{Q}}[m_1(Y_1) \mid Y_2], \qquad r_2(Y_2) = \mathbb{E}_{\mathbb{Q}}[r_1(Y_1) \mid Y_2].$$

Subtracting the two price expressions and applying the conditional covariance identity yields

$$\mathbb{E}P_2 - \mathbb{E}P_1 = \mathbb{E}_{\mathbb{Q}}\mathrm{Cov}_{\mathbb{Q}}(m_1(Y_1), r_1(Y_1) \mid Y_2). \tag{26}$$

The price comparison has thus been reduced to the sign of a covariance between two functions of the original signal.

Log concavity of the density of X_1 implies that the posterior distribution of W shifts upward in likelihood ratio order as Y_1 increases, so its conditional mean m_1 is nondecreasing. If the noise support is nonnegative, $L(w) = K\exp(w^2/(2b))$ is also nondecreasing, so its conditional expectation r_1 is nondecreasing. Two nondecreasing functions of the same random variable have nonnegative covariance under every conditional distribution. For completeness, if A' is an independent copy of A under that distribution, their covariance is

$$\frac{1}{2}\mathbb{E}\big[(m_1(A) - m_1(A'))(r_1(A) - r_1(A'))\big] \geq 0.$$

Equation (26) then gives $\mathbb{E}P_2 \geq \mathbb{E}P_1$.

On a nonpositive noise support, L is nonincreasing, and the same posterior ordering makes r_1 nonincreasing while m_1 remains nondecreasing, so the covariance has the opposite sign and $\mathbb{E}P_2 \leq \mathbb{E}P_1$. Finally, $\Omega_\ell^+(t) = \Phi(-w_\ell(t)/\sqrt{b})$ is at most one half exactly when $w_\ell(t) \geq 0$, and at least one half exactly when $w_\ell(t) \leq 0$, which are the two ownership conditions in the proposition. The argument uses independence and the zero mean of Z, together with its stated moment condition, and allows an arbitrary distribution for the added risk. □

For Gaussian fundamentals, center Y by μ_0 and put $f_a = f * \phi_a$ and $g_a = g * \phi_a$, where ϕ_a is the centered Gaussian density with variance a. Let m_f, m_g denote the conditional means of W under the two noise laws and V_g its conditional variance under g. For $j = f, g$, differentiating the Gaussian likelihood with respect to the observed signal gives

$$m_j(y) = y + a\partial_y \log j_a(y), \qquad \partial_y m_g(y) = V_g(y)/a.$$

Each convolution satisfies the heat equation $\partial_a j_a = \frac{1}{2}\partial_{yy} j_a$. To differentiate the posterior mean, put $s_g = \partial_y \log g_a$. Then $\partial_a s_g = \frac{1}{2}\partial_{yy} s_g + s_g \partial_y s_g$ and $m_g = y + a s_g$, so

$$\partial_a m_g = \tfrac{1}{2}\partial_{yy} m_g + s_g \partial_y m_g.$$

Both the conditional mean and the actual density of the observed signal vary with a, and differentiating their integral and integrating by parts gives

$$\begin{aligned}\frac{d}{da}\int m_g f_a \, dy &= \int \left(\tfrac{1}{2}\partial_{yy} m_g + s_g \partial_y m_g\right) f_a \, dy + \tfrac{1}{2}\int m_g \partial_{yy} f_a \, dy \\ &= \int \left(\partial_{yy} m_g + s_g \partial_y m_g\right) f_a \, dy\end{aligned}$$

$$= \int (s_g - \partial_y \log f_a)(\partial_y m_g) f_a \, dy = -\frac{1}{a^2}\mathbb{E}[(m_f - m_g)V_g].$$

In the last equality, $s_g - \partial_y \log f_a = (m_g - m_f)/a$ and $\partial_y m_g = V_g/a$. The boundary terms vanish because the noise has compact support and the convolutions have Gaussian tails. The expected price premium is $\mathbb{E}_f W - \mathbb{E}m_g$, the valuation gap is $m_f - m_g$, and $V_g = U_{0,a}$. Hence

$$\frac{d\mathbb{E}[P_a]}{da} = \frac{1}{a^2}\mathbb{E}[U_{0,a}(P_a)\,\{P_a - \mathbb{E}[X_a \mid P_a]\}]\,.$$

Compact noise support bounds posterior moments and provides Gaussian tail bounds for the differentiations. The same identity holds in the Gaussian noise specialization, where all integrals can be evaluated directly.

The same calculation clarifies the effect of a more precise signal observed by every investor. Let $H = X + \zeta$, where $\zeta \sim N(0, \tau^{-1})$ is independent of supply and private errors, and retain $X \sim N(\mu_0, a)$. Conditional on H, fundamental variance is $a_\tau = (a^{-1} + \tau)^{-1}$. Translation equivariance gives $\mathbb{E}[P] = \mu_0 + B(a_\tau)$, where $B(v)$ is the expected premium in the economy with fundamental variance v. At differentiable points,

$$\frac{\partial \mathbb{E}[P]}{\partial \tau} = -a_\tau^2 B\prime(a_\tau) = -a^2 \frac{\partial \mathbb{E}[P]}{\partial a} = -\mathbb{E}[U_0\{P - \mathbb{E}[X \mid P, H]\}]\,.$$

Thus, holding the supply law, private signal variance, preferences, and shorting limit fixed, more precise public information reverses the expected price effect of greater fundamental variance, even when equilibrium price is nonlinear.

A.2 Theoretical Extensions

A.2.1 Dynamic Noisy and Sticky Information

With persistent fundamentals, investors carry earlier private information into their current forecasts, so disagreement can reflect both recent signals and differences inherited from previous dates. To distinguish these sources, let

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

Noisy Information. In a Gaussian noisy information benchmark in the spirit 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

$$A_t = \rho^2 U_{t-1} + \sigma_\epsilon^2, \qquad K_t = \frac{A_t}{A_t + \sigma_\nu^2}, \qquad U_t = (1 - K_t)A_t = \sigma_\nu^2 K_t. \tag{27}$$

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

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

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

$$D_t = \rho^2(1 - K_t)^2 D_{t-1} + K_t^2 \sigma_\nu^2. \tag{28}$$

The inherited term in (28) carries dispersion from earlier private signals and gives the dynamic recursion its persistence.

In the stationary equilibrium, write K for the steady state gain. Equations (27)–(28) give

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

The statewise unit ceiling remains intact:

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

The Kalman gain also governs the comparative statics, since combining $A = \rho^2 U + \sigma_\epsilon^2$, $U = \sigma_\nu^2 K$, and $A = \sigma_\nu^2 K/(1 - K)$ yields

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

The right hand side is strictly increasing in K, and the stationary ratio satisfies

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

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, while 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, sticky information is the limiting case of a noisy information process with heterogeneous precision: at each date, the fraction $1-\lambda$ that updates receives an infinitely precise incremental signal, while the fraction λ that does not update receives no signal. An age-s vintage records s consecutive dates without an update 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

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

where $P_0 = 0$. Disagreement conditional on the realized state is therefore

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

Summing over vintages gives average posterior uncertainty and expected disagreement,

$$U^S = \sum_{s=0}^{\infty} w_s P_s = \frac{\lambda \sigma_\epsilon^2}{1-\lambda\rho^2},$$
$$\bar{D}^S = \mathbb{E}[D_t^S] = \frac{\lambda(1-\lambda)\sigma_\epsilon^2}{(1-\lambda\rho^2)(1-\lambda^2\rho^2)}, \qquad \frac{\bar{D}^S}{U^S} = \frac{1-\lambda}{1-\lambda^2\rho^2} < 1.$$

The sticky information benchmark therefore preserves the ex ante common prior restriction while allowing realized disagreement to vary with information histories. In particular, the ex ante ratio remains below one even though realized disagreement can exceed average posterior uncertainty along some histories, providing a common prior benchmark for the occasional statewise exceedances in the survey data.

A.2.2 Fixed Float with Endogenous Price Inference

Stochastic asset supply is not required for the ownership result. Consider instead an asset with fixed outstanding supply $\bar{S} > 0$, held by investors who receive private information about the payoff and also carry privately observed nontraded exposures to the same risk. Hedging demand then prevents the price from revealing the payoff completely, while market clearing determines how strongly aggregate exposure enters the information conveyed by price.

Investor i has CARA utility with risk aversion $\gamma > 0$ and a position satisfying

$$q_i \geq -\ell.$$

The asset pays X, with prior distribution F_X, and investor i observes

$$s_i = X + \epsilon_i, \qquad \epsilon_i \sim N(0, b),$$

and a nontraded exposure

$$e_i = H + \nu_i, \qquad \nu_i \sim N(0, c), \qquad H \sim N(0, v_H),$$

where H is independent of X. Conditional on (X, H), the cross-sectional distribution of (ϵ_i, ν_i) is the product of the corresponding Gaussian distributions. Apart from terms irrelevant for portfolio choice, terminal wealth is

$$W_i = (q_i + e_i)X - q_i P.$$

A larger positive exposure therefore gives an investor a stronger reason to sell the traded claim, and market clearing requires $\int q_i \, di = \bar{S}$.

Suppose price reveals the statistic

$$Y = X - kH, \qquad P = p(Y),$$

where $k > 0$ is determined in equilibrium. Conditional on $Y = y$, the posterior density of X given (s_i, e_i) is proportional to

$$\exp\left\{ -\frac{(x-y)^2}{2k^2 v_H} - \frac{(s_i - x)^2}{2b} - \frac{(e_i - (x-y)/k)^2}{2c} \right\} dF_X(x).$$

Thus, both private observations contain information about X once price is observed.

For an unconstrained position q_i, CARA optimality implies

$$P = \frac{\mathbb{E}\left[Xe^{-\gamma(q_i+e_i)X} \mid s_i, e_i, Y\right]}{\mathbb{E}\left[e^{-\gamma(q_i+e_i)X} \mid s_i, e_i, Y\right]}. \tag{29}$$

Consider desired demand

$$q_i^u = A + \frac{s_i - y}{\gamma b} + \left(\frac{1}{\gamma kc} - 1\right) e_i. \tag{30}$$

Substitution into (29) removes the investor-specific terms from the exponential tilt. Define

$$\frac{1}{r} = \frac{1}{b} + \frac{1}{k^2 c} + \frac{1}{k^2 v_H}.$$

The resulting price function is

$$p(y) = \frac{\int x \exp\left\{-\dfrac{(x-y)^2}{2r} - \gamma A x\right\} dF_X(x)}{\int \exp\left\{-\dfrac{(x-y)^2}{2r} - \gamma A x\right\} dF_X(x)}. \tag{31}$$

Conditional on (X, H), substituting $s_i = X + \epsilon_i$, $e_i = H + \nu_i$, and $y = X - kH$ into (30) gives

$$q_i^u = A + H\left[\frac{k}{\gamma b} + \frac{1}{\gamma kc} - 1\right] + \frac{\epsilon_i}{\gamma b} + \left(\frac{1}{\gamma kc} - 1\right)\nu_i.$$

Because actual demand is $q_i = \max\{q_i^u, -\ell\}$, a common shift in desired positions changes aggregate demand. Fixed supply can therefore clear for every realization of H only if

$$\frac{k}{b} + \frac{1}{kc} = \gamma. \tag{32}$$

Equivalently,

$$\frac{k^2}{b} - \gamma k + \frac{1}{c} = 0.$$

A positive solution exists when $\gamma^2 bc \geq 4$, and under strict inequality, the two positive roots are

$$k_\pm = \frac{\gamma b \pm \sqrt{\gamma^2 b^2 - 4b/c}}{2}.$$

The construction that follows applies to either of these roots.

Using (32), desired demand simplifies to

$$q_i^u = A + \frac{\epsilon_i - k\nu_i}{\gamma b}. \tag{33}$$

Its cross-sectional variance is

$$\sigma_q^2 = \frac{b + k^2 c}{\gamma^2 b^2}.$$

The common location A is the unique solution of

$$\mathbb{E}\big[\max\{A + \sigma_q Z, -\ell\}\big] = \bar{S}, \qquad Z \sim N(0,1). \tag{34}$$

Finally, differentiation of (31) gives

$$p'(y) = \frac{\mathrm{Var}_{R_{A,y}}(X)}{r} > 0, \tag{35}$$

where $R_{A,y}$ denotes the exponentially tilted distribution in (31). Price is therefore strictly increasing in Y, so observing P is equivalent to observing the statistic used in investors' posterior beliefs.

Proposition 8 (Constrained rational expectations equilibrium with fixed float)**.** *Suppose $\gamma^2 bc \geq 4$ and the payoff distribution is nondegenerate with the exponential moments required by* (31)*. For any positive solution k of* (32) *and the unique A satisfying* (34)*, equations* (31) *and* (33)*, together with $q_i = \max\{q_i^u, -\ell\}$, constitute a rational expectations equilibrium with fixed supply $\bar{S}$.*

Proof. Equation (29) and the exponential tilt leading to (31) establish individual optimality before the position floor is imposed. Since CARA expected utility is strictly concave in the traded position, truncating the unconstrained optimum at $-\ell$ gives the constrained optimum. Equation (32) removes the aggregate exposure term from the location of desired demand, and (34) then clears the fixed supply in every aggregate state. The strict monotonicity in (35) verifies that the observed price reveals exactly $Y = X - kH$, completing the rational expectations argument. □

Fundamental risk and ownership. With $X \sim N(\mu_0, a)$ and $U = \left(\frac{1}{a} + \frac{1}{r}\right)^{-1}$, the equilibrium price is

$$P = \mu_0 + \frac{U}{r}(Y - \mu_0) - \gamma U A, \qquad \mathbb{E}[P] = \mu_0 - \gamma U A.$$

Positive ownership is equivalent to $q_i^u > 0$, and its cross-sectional share is

$$\Omega^+ = \Phi\left(\frac{A}{\sigma_q}\right).$$

Hence, $A < 0$ if and only if $\Omega^+ < 1/2$. An increase in the fundamental variance a, holding preferences, signal and exposure technologies, supply, and the position floor fixed, leaves k, r, σ_q, and A unchanged, while $\frac{dU}{da} = \frac{U^2}{a^2} > 0$. It follows that

$$\frac{d\mathbb{E}[P]}{da} = -\gamma A \frac{U^2}{a^2}, \qquad \frac{d\mathbb{E}[P]}{da} > 0 \iff \Omega^+ < \frac{1}{2}.$$

The one half ownership threshold therefore survives with fixed outstanding supply and equilibrium inference from price.

Beliefs and holdings. In the Gaussian case, if m_i is investor i's posterior mean of X, then deviations from the cross-sectional mean satisfy

$$m_i - \bar{m} = U\left(\frac{\epsilon_i}{b} + \frac{\nu_i}{kc}\right), \qquad D = U^2\left(\frac{1}{b} + \frac{1}{k^2 c}\right) = \frac{\gamma}{k}U^2.$$

Desired traded positions instead satisfy

$$q_i^u - A = \frac{\epsilon_i - k\nu_i}{\gamma b}, \qquad \mathrm{Cov}_i(q_i^u, m_i \mid X, H) = 0.$$

The difference comes from hedging, since an investor with a high posterior mean may nevertheless choose a small traded position when that investor already carries a large positive exposure to X. Defining the valuation relevant for the traded position by $v_i = m_i - \gamma U e_i$ gives

$$q_i^u = \frac{v_i - P}{\gamma U}, \qquad v_i - \bar{v} = \frac{U}{b}(\epsilon_i - k\nu_i).$$

Thus, the short sale constraint truncates the lower tail of valuations after accounting for investors' existing exposures. Fundamental risk raises expected price under narrow ownership even though the cross-sectional ordering of holdings need not coincide with the ordering of posterior means.

B Survey Measurement and Supporting Evidence

In the Survey of Professional Forecasters (SPF), respondents submit point forecasts and probability histograms for the same macroeconomic outcome and horizon, which allows uncertainty and disagreement to be measured within a common respondent sample. For expected S&P 500 returns, disagreement in two surveys is instead compared with option-implied market uncertainty, with the caveat that the survey forecasts reflect respondents' subjective beliefs, whereas the option measure is inferred under the risk-neutral distribution.

B.1 SPF Point and Density Forecasts

The U.S. Survey of Professional Forecasters, administered by the Federal Reserve Bank of Philadelphia, reports both point forecasts and probability histograms, and the observations used here extend through 2026Q2. Since a reported point forecast need not be compatible with the accompanying histogram, centering a normal distribution on that forecast can spuriously inflate the fitted uncertainty. Each response is therefore screened for compatibility before the subjective distributions are fitted.

For respondent j in quarter t, let π_{jtk} be the reported probability in bin $[a_k, b_k]$. For the consistency screen only, inflation histograms receive the finite support $[-12, 12]$ and unemployment histograms the support $[0, 25]$, while the open tails remain open in the subsequent normal fit. Once those supports are assigned to the outer bins, the sharp set of means consistent with the histogram is

$$\mathcal{M}_{jt} = \left[\sum_k \pi_{jtk} a_k, \sum_k \pi_{jtk} b_k \right]. \tag{36}$$

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.

For each response that passes the screen, a normal distribution is centered at the reported point forecast, and its standard deviation is chosen by minimum distance, so that

$$\widehat{\sigma}_{jt} = \arg \min_{\sigma \in [10^{-4}, 100]} \sum_k \left[\pi_{jtk} - \left\{ \Phi\left(\frac{b_k - f_{jt}}{\sigma} \right) - \Phi\left(\frac{a_k - f_{jt}}{\sigma} \right) \right\} \right]^2.$$

Because the finite tail supports enter only the consistency criterion in (36), the two open tails remain open in this fit. Variable-quarter moments are computed on the identical retained respondent set:

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

Within the retained samples, median D/U is 0.232 for core CPI, 0.218 for core PCE, and 0.061 for unemployment (Table A1). Only four variable-quarter observations exceed one: core CPI in 2010Q1, 2011Q1, and 2023Q1, and core PCE in 2022Q1, with ratios of 1.302 for core CPI in 2023Q1 and 1.513 for core PCE in 2022Q1.

Table A1: SPF uncertainty and disagreement from fitted normal distributions

	(1) T	**(2)** Median N	**(3)** Median D/U	**(4)** 95th pct.	**(5)** Maximum	**(6)** $D/U > 1$
Core CPI	78	25	0.232	0.960	1.806	3
Core PCE	78	23	0.218	0.725	1.513	1
Unemployment	69	29	0.061	0.297	0.797	0

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 and U is the mean fitted posterior variance, computed on the same retained respondent set. T is the number of survey rounds and N the number of retained respondents per round.

As the forecast horizon shortens toward a fixed calendar-year outcome, mean D/U across core CPI, core PCE, and unemployment falls from 0.412 in the first survey quarter to 0.069 in the fourth, while the median falls from 0.289 to 0.057 (Table A2). Each survey round refers to the same calendar-year target, so the sequence traces the accumulation of information as the outcome approaches resolution.

B.2 Surveys of Equity Returns and Option-Implied Uncertainty

The equity comparison uses one year S&P 500 return forecasts from the Livingston Survey and the Duke CFO Survey. For the Livingston Survey, disagreement is the variance across respondents in forecasts of the S&P 500 log return over the following year. Each survey observation is matched to the nearest preceding value of the constructed index SVIX variance series at a horizon of 365 days, provided that the option observation falls within four calendar days. The ratio of forecast disagreement to index variance inferred from

Table A2: Uncertainty and disagreement as a fixed forecast event approaches

	(1) Variable-years	(2) Mean U	(3) Mean D	(4) Mean D/U	(5) 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, using all available observations for each variable and year. Respondents are screened separately in each survey round, so the number of observations differs across quarters. 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 A1.

options has a mean of 0.096 and a maximum of 0.225 across 52 observations from June 1996 through December 2021.

For the Duke CFO Survey, the analysis uses the cleaned most likely return forecasts constructed by Gandhi et al. (2026), with disagreement defined as the variance across these forecasts. Each survey close date is matched to Cboe VIX1Y, which is squared to express the option measure in variance units. The resulting ratio has a mean of 0.057 and a maximum of 0.117 across 22 observations from 2020Q3 through 2025Q4. Across the two survey series, forecast disagreement lies below the corresponding variance inferred from options in all 74 observations.

Both equity surveys show that forecast disagreement is small relative to option-implied variance, a pattern consistent with information pooling through a continuously observed market price. The aligned SPF forecasts provide the direct comparison of disagreement and subjective uncertainty because their point forecasts and probability distributions concern the same outcomes and horizons.

C Variable Definitions

The stock universe consists of the ordinary common shares of U.S.-incorporated corporations in the data of Jensen et al. (2023), restricted at each formation date to primary listings on the NYSE, NYSE American, and Nasdaq with positive market equity, the counterpart of CRSP share codes 10 and 11. Predictors dated at the end of month t are matched to CRSP total returns in month $t+1$. Delisting returns are included, and a missing delisting return following a delisting for poor performance is set to -30 percent, following Shumway (1997). Unless a table states otherwise, regressions are estimated separately in each month with equal weights, and reported coefficients are time-series means of the monthly slopes with six lag Newey–West t statistics. In the monthly return regressions, next-month returns are winsorized at the monthly 1st and 99th percentiles. Portfolio returns, the daily returns in the earnings analysis, and the returns in Table 3 are not winsorized. The pooled regressions in Table 3 cluster standard errors by stock and month, the regressions in Table A3 use the stated observation weights, and the regressions forecasting overlapping future volatility in Table A5 use 12 lags.

Sample requirements. Regressions using the covered sample require all three IVOL measures (weekly returns over three years and daily returns over one month and one year), together with total and market-adjusted option variances at both 30 and 365 days. The regressions in Tables 6 and A10, which use logarithms of uncertainty, additionally require all four option variances to be positive. Each regression then requires observations for its outcome and included regressors. The benchmarks for the full stock universe and the historical portfolio sorts use the samples specified in their table and figure notes.

Idiosyncratic and realized volatility. For each stock, daily excess returns over the 252 trading days ending in month t are regressed on the market excess return, SMB, and HML, following Ang et al. (2006), with at least 200 daily observations. With n_{it} complete daily observations and fitted residuals $\widehat{\epsilon}_{id}$,

$$IVOL_{it} = \left[\frac{1}{n_{it}-1}\sum_{d}\widehat{\epsilon}_{id}^{2}\right]^{1/2}, \qquad n_{it} \geq 200.$$

The weekly alternative uses residuals from market-model regressions of weekly returns over the preceding three calendar years, with at least 52 observations, and the one-month alternative applies the daily FF3 regression to the daily returns within month t. Realized volatility is the standard deviation of daily excess returns within month t. Both measures

based on returns within month t require at least 15 daily observations and are taken from the data of Jensen et al. (2023). Volatility measures are logged, winsorized at the monthly 1st and 99th percentiles, and standardized within month unless a table specifies ranks. In Table 7, IVOL, uncertainty, and disagreement are ranked and standardized separately within each ownership quintile and month.

Option-implied uncertainty. Following Martin (2017) and Martin and Wagner (2019), each firm-month uses the final valid OptionMetrics observation in month t, with the forward price inferred from put-call parity. Out-of-the-money puts and calls are integrated across strikes on a fitted surface at a constant horizon h, with at least eight strike nodes. With spot price P_{it}, gross risk-free return $R_{f,t,h}$, fitted premium $Q_{it}(K_k)$, and midpoint strike width ΔK_k,

$$U_{it,h} = \frac{2}{R_{f,t,h} P_{it}^2} \sum_k \Delta K_k \, Q_{it}(K_k),$$

the risk-neutral variance of the return over horizon h divided by the square of the gross risk-free return. The principal measure uses $h = 365$ days, so that it is an annual variance, and the comparison with Martin and Wagner (2019) in Table 3 uses the variance over $h = 30$ days. The market-adjusted measure is $\max\{U_{it,h} - \beta_{it}^2 U_{t,h}^{mkt}, 0\}$, where β_{it} is estimated before formation and $U_{t,h}^{mkt}$ is the corresponding S&P 500 index variance. In Table 3, firm variance and returns are measured net of their value-weighted sample averages, returns are divided by the gross risk-free return, and standard errors are clustered by stock and month.

Variance premium proxy. For each past formation month, we regress annualized realized variance over the next 252 trading days on a constant, annualized variance over the preceding year, standardized log IVOL, and the seven firm characteristics. Annualized variance is 252 times the variance of daily stock returns. At each current formation date, we average the coefficients from all past monthly regressions whose subsequent 252 trading days have been fully observed, weighting each month equally. Applying these coefficients to current predictors gives the forecast used in Table A6. The proxy is 365-day option variance less this forecast, winsorized and standardized within month.

Disagreement. Disagreement uses annual EPS forecasts in U.S. dollars from at least three analysts, with a fiscal endpoint 180 to 550 days after the forecast date, choosing

the endpoint nearest one year when two coexist. With $\mathrm{SD}(EPS)_{it}$ the standard deviation across analysts and P_{it} the formation price,

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

and observations with zero dispersion are excluded.

Firm characteristics. Missing values of the book-to-market ratio, momentum, reversal, operating profitability, and asset growth are replaced with the corresponding monthly medians among stocks in the Jensen et al. (2023) universe with positive market equity, before matching to the belief measures. These five characteristics enter in levels; market equity enters in logarithms, and analyst coverage as the logarithm of one plus the I/B/E/S analyst count. Each control is winsorized at the monthly 1st and 99th percentiles and standardized on the matched panel before applying the common sample requirements and the restrictions specific to each regression. Market equity and analyst coverage are not imputed.

Borrowing cost and ownership breadth. Implied borrowing cost is the stock-specific rate supplied by OptionMetrics, taken from the final valid observation in month t at the maturity nearest 365 days among maturities of 30 to 550 days. Following Chen et al. (2002), ownership breadth is the fraction of reporting S12 mutual funds holding the stock, lagged one quarter. Both enter interactions as standardized within-month percentile ranks. For a characteristic M_{it}, the interaction regressions in Table 5 are

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

where $y_{i,t+1}$ is the next-month return, or contemporaneous IVOL in the regressions explaining IVOL, and X_{it} contains the characteristics. In these regressions U and D enter as standardized within-month ranks. The separate IVOL return regressions replace U and D by IVOL. In Table A15, borrowing cost is first projected within month on size and liquidity, where the liquidity composite averages the within-month ranks of the bid–ask spread, Amihud illiquidity, and zero trading, and the residual is reranked and standardized.

Earnings announcements. Compustat announcement dates are mapped to the CRSP trading calendar, and the event window contains trading days -1 through $+1$. Table 9 and columns (1)–(3) of Table A11 compare average daily returns within the window and

on other trading days, retaining months with observations in both periods. Table A12 uses average daily returns within the window. Columns (4)–(6) of Table A11 instead allocate the compounded monthly excess return across days. Starting with one dollar in the stock and one dollar in the risk-free asset at the beginning of the month, day d contributes $W^s_{d-1} r_d - W^f_{d-1} r_{f,d}$, where the two wealth terms record the respective values at the beginning of that day. Summing over announcement days and other days gives contributions that add exactly to the monthly excess return. This decomposition also retains valid months without an announcement, with zero contribution from announcement days.

The regressions in Table A13 use compounded market-adjusted returns over days -1 through $+1$. The belief measures and firm characteristics are dated at the last month-end strictly before trading day -10 relative to the announcement. The earnings surprise is I/B/E/S actual EPS minus the most recent consensus available before the return window, on a common share basis, divided by the closing price on day -2 and winsorized at the monthly 1st and 99th percentiles.

IVOL decomposition. Table 8 uses standardized log IVOL, ranked U, and log D, retaining the monthly transformations calculated before dividing stocks into ownership groups. Within each group, we demean the transformed variables and returns. Let $P_{\mathcal{A}}$ denote the cross-sectional projection onto a subset $\mathcal{A}$ of $\{U, D\}$ and define

$$v(\mathcal{A}) = \frac{\mathrm{Cov}(P_{\mathcal{A}} IVOL,\, P_{\mathcal{A}} r)}{\mathrm{Var}(IVOL)}.$$

The uncertainty contribution in Table 8 is $\phi_U = \frac{1}{2}\{v(\{U\}) - v(\varnothing) + v(\{U, D\}) - v(\{D\})\}$, the disagreement contribution ϕ_D is defined symmetrically, and the residual is the univariate IVOL slope less $v(\{U, D\})$, so the three components sum to the IVOL slope.

Reweighting. In Table A3, a covered stock in size and liquidity group g in month t receives weight $w_{it} = N^{all}_{g,t} / N^{cov}_{g,t}$, the ratio of the number of all stocks to covered stocks in the group, normalized to a monthly mean of one. The reference distribution is formed before requiring option or analyst data, IVOL, or subsequent returns.

D Additional Empirical Results

This appendix reports the supporting estimates cited in Section 5.

D.1 Sample Coverage and Missing Characteristics

Table A3 weights covered stocks each month by the relative frequencies of size and liquidity groups in the full and covered samples. We form either 2×2 or 5×5 groups each month by sorting stocks on size and liquidity. In the finer partition, groups represented in the covered sample account for 87.1 percent of all observations in the full sample.

Table A4 compares the principal return regression on the stock-months for which every characteristic is observed, with and without the replacement of missing characteristics by monthly medians.

D.2 Uncertainty

Table A5 relates uncertainty to realized volatility over the following year, and Table A6 adds the volatility spread and a variance risk premium proxy to the return regressions of Table 4 on identical observations.

D.3 Constraints

Table A7 repeats columns (1)–(6) of Table 5 with market-adjusted uncertainty, which subtracts beta-scaled index option variance at the same horizon. Table A8 repeats columns (7)–(9) with the weekly three-year and one-month IVOL measures. Table A10 repeats Table 6 with groups formed on market equity rather than ownership breadth.

Table A9 compares separate and joint interactions of IVOL with firm size and ownership breadth. All three specifications use identical observations and retain the standard characteristics.

D.4 Earnings Announcements

Table A11 decomposes the IVOL relation in the full stock sample into the parts earned in earnings windows and on other days. Table A12 adds lottery characteristics to the earnings-window IVOL regressions and repeats the comparison with weekly three-year IVOL. Table A13 relates the belief measures to earnings surprises, measured from quarterly I/B/E/S actual earnings and the last consensus before the announcement and scaled by price, and to announcement returns conditional on the surprise.

Table A3: Reweighting the covered sample toward all stocks

	(1)	(2)	(3)	(4)	(5)
Weights:	Unweighted	2×2	2×2 capped	5×5	5×5 capped
Panel A: Dependent variable is IVOL					
U	0.822^{***}	0.810^{***}	0.809^{***}	0.810^{***}	0.810^{***}
	(92.50)	(77.46)	(78.21)	(83.26)	(83.45)
D	0.079^{***}	0.107^{***}	0.106^{***}	0.111^{***}	0.105^{***}
	(9.15)	(10.40)	(10.40)	(10.35)	(10.13)
Mean R^2	0.768	0.780	0.779	0.793	0.785
Panel B: Dependent variable is next-month return					
U	-0.079	-0.038	-0.079	-0.084	-0.083
	(-0.45)	(-0.21)	(-0.42)	(-0.47)	(-0.47)
D	-0.186^{**}	-0.328^{***}	-0.331^{***}	-0.459^{***}	-0.334^{***}
	(-2.01)	(-2.95)	(-2.94)	(-4.00)	(-3.39)
Mean R^2	0.059	0.054	0.053	0.061	0.052
Panel C: Dependent variable is next-month return					
U	-0.089	-0.039	-0.074	-0.084	-0.080
	(-0.51)	(-0.21)	(-0.39)	(-0.48)	(-0.46)
D	-0.130	-0.173^{*}	-0.158	-0.208^{**}	-0.165^{*}
	(-1.44)	(-1.82)	(-1.62)	(-2.13)	(-1.80)
$U \times D$	-0.186^{***}	-0.202^{***}	-0.222^{***}	-0.279^{***}	-0.221^{***}
	(-3.96)	(-3.25)	(-3.69)	(-3.87)	(-3.92)
Mean R^2	0.064	0.061	0.059	0.070	0.057
Panel D: Dependent variable is next-month return					
IVOL	-0.288^{**}	-0.671^{***}	-0.588^{***}	-0.751^{***}	-0.510^{***}
	(-2.57)	(-2.95)	(-2.93)	(-3.60)	(-3.66)
U	0.166	0.555^{***}	0.438^{***}	0.570^{***}	0.349^{***}
	(1.32)	(3.36)	(2.90)	(3.68)	(3.00)
D	-0.169^{*}	-0.295^{***}	-0.301^{***}	-0.411^{***}	-0.299^{***}
	(-1.83)	(-2.74)	(-2.77)	(-3.68)	(-3.19)
Mean R^2	0.063	0.063	0.060	0.070	0.057
Months	356	356	356	356	356
Stocks	7,593	7,593	7,593	7,593	7,593

Monthly Fama–MacBeth regressions weighted by the relative frequencies of size and liquidity groups in the full and covered samples. Groups are formed by 2×2 or 5×5 sorts; the represented 5×5 groups cover 87.1 percent of observations in the full sample. Capped weights are limited to the monthly 99th percentile. Stock counts refer to Panels B–D; Panel A does not require a subsequent return and includes 7,594 stocks. No firm characteristics are included. U is 365-day total uncertainty and enters as a standardized monthly rank; D and IVOL enter as standardized logarithms. These transformations are retained in the weighted regressions. t statistics use Newey–West standard errors with six lags.

Table A4: Missing characteristics and return slopes

	Main sample	Complete-case sample	
	(1) All observations	**(2)** Filled controls	**(3)** Unfilled controls
Panel A: 365-day total uncertainty			
U	0.253^{**}	0.288^{**}	0.285^{**}
	(2.25)	(2.52)	(2.48)
D	-0.177^{***}	-0.165^{**}	-0.165^{***}
	(-2.68)	(-2.57)	(-2.58)
IVOL	-0.132	-0.178^{**}	-0.179^{**}
	(-1.40)	(-2.15)	(-2.15)
Realized volatility	-0.225^{***}	-0.184^{***}	-0.184^{***}
	(-3.77)	(-2.95)	(-2.95)
Panel B: 365-day market-adjusted uncertainty			
U	0.174^{***}	0.203^{***}	0.202^{***}
	(2.63)	(3.04)	(3.00)
D	-0.163^{**}	-0.150^{**}	-0.150^{**}
	(-2.41)	(-2.28)	(-2.29)
IVOL	-0.094	-0.133	-0.135
	(-0.77)	(-1.19)	(-1.21)
Realized volatility	-0.206^{***}	-0.163^{**}	-0.163^{**}
	(-3.02)	(-2.27)	(-2.28)
Characteristics	✓	✓	✓
IVOL and realized volatility	✓	✓	✓
Months	356	356	356
Stocks	7,593	6,766	6,766

Column (1) repeats column (3) of Table 4, in which missing values of book-to-market, momentum, reversal, operating profitability, and asset growth are replaced by the cross-sectional median of the month. Columns (2) and (3) restrict the sample to stock-months with all characteristics observed. In column (3), the five financial controls are winsorized and standardized using their observed values on the matched panel before selecting these observations; column (2) retains the transformations calculated after replacing missing values. Appendix C describes the population used for the median replacements. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A5: Uncertainty and future volatility

	(1)	(2)	(3)	(4)
Panel A: Idiosyncratic volatility over the next 252 trading days				
U	0.839***		0.364***	0.282***
	(141.91)		(19.68)	(17.81)
IVOL		0.729***	0.461***	0.434***
		(62.74)	(20.20)	(20.94)
Realized vol.		0.154***	0.080***	0.075***
		(19.34)	(15.69)	(15.22)
D				0.066***
				(7.73)
Mean R^2	0.694	0.727	0.752	0.771
Panel B: Total volatility over the next 252 trading days				
U	0.844***		0.442***	0.375***
	(97.81)		(21.44)	(21.18)
IVOL		0.635***	0.309***	0.290***
		(36.54)	(11.08)	(11.70)
Realized vol.		0.252***	0.163***	0.141***
		(19.69)	(19.46)	(21.05)
D				0.080***
				(7.45)
Mean R^2	0.705	0.714	0.750	0.772
Characteristics	✗	✗	✗	✓
Months	343	343	343	343
Stocks	7,213	7,213	7,213	7,213

Monthly Fama–MacBeth regressions of future realized volatility on predictors measured at the end of the formation month. Idiosyncratic volatility is the residual standard deviation from daily FF3 regressions over the window; total volatility is the standard deviation of daily returns. Each outcome is logged, winsorized, and standardized within month. U is 365-day total uncertainty and enters as a standardized monthly rank; D, IVOL, and realized volatility enter as standardized logarithms. IVOL and realized volatility use past returns. Months and stocks refer to Panel A. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with 12 lags, since the outcome windows overlap.

Table A6: Uncertainty, the volatility spread, and the variance risk premium

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: With the standard characteristics						
U	−0.012	−0.020	0.000	−0.025		
	(−0.07)	(−0.11)	(0.00)	(−0.14)		
D	−0.201***	−0.204***	−0.190***	−0.174***	−0.190**	−0.176*
	(−2.87)	(−2.93)	(−2.64)	(−2.59)	(−2.08)	(−1.95)
Volatility spread		0.023		0.323**	0.033	
		(0.40)		(2.05)	(0.61)	
Variance risk premium			−0.080	−0.370**		−0.071
			(−1.46)	(−2.40)		(−1.00)
Mean R^2	0.096	0.100	0.100	0.105	0.086	0.087
Panel B: With the standard characteristics, IVOL, and realized volatility						
U	0.229**	0.318**	0.336***	0.145		
	(1.98)	(2.09)	(2.93)	(0.67)		
D	−0.173**	−0.161**	−0.150**	−0.158**	−0.150**	−0.136*
	(−2.56)	(−2.30)	(−2.13)	(−2.31)	(−2.08)	(−1.88)
Volatility spread		−0.078		0.195	0.015	
		(−1.16)		(0.98)	(0.26)	
Variance risk premium			−0.139**	−0.283		−0.071
			(−2.32)	(−1.62)		(−1.13)
IVOL	−0.120	−0.205*	−0.221**	−0.023	−0.022	−0.044
	(−1.24)	(−1.65)	(−2.00)	(−0.16)	(−0.15)	(−0.31)
Realized vol.	−0.218***	−0.217***	−0.216***	−0.202***	−0.170**	−0.164**
	(−3.54)	(−3.64)	(−3.51)	(−3.72)	(−2.48)	(−2.44)
Mean R^2	0.102	0.107	0.106	0.111	0.102	0.103
Characteristics	✓	✓	✓	✓	✓	✓
Months	344	344	344	344	344	344
Stocks	7,473	7,473	7,473	7,473	7,473	7,473

Monthly Fama–MacBeth regressions of next-month returns (percentage points). The volatility spread is the square root of annualized 365-day implied variance minus trailing one-year realized volatility. The proxy for the variance risk premium is 365-day implied variance minus the forecast defined in Appendix C, which averages coefficients from past monthly regressions only after their subsequent returns have been observed. Both are winsorized and standardized within month. All columns use identical observations, on which the average monthly cross-sectional correlation between the two proxies is 0.845. U is 365-day total uncertainty and enters as a standardized monthly rank; D, IVOL, and realized volatility enter as standardized logarithms. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A7: Constraints and the composition of IVOL with market-adjusted uncertainty

	Dependent variable: IVOL			Dependent variable: next-month return		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
M:	Borrow.	Breadth	Size	Borrow.	Breadth	Size
U	0.522^{***}	0.509^{***}	0.517^{***}	0.008	0.009	-0.010
	(31.96)	(30.81)	(31.31)	(0.07)	(0.07)	(-0.09)
D	0.202^{***}	0.206^{***}	0.208^{***}	-0.137^{*}	-0.136^{*}	-0.147^{**}
	(16.13)	(16.48)	(16.18)	(-1.85)	(-1.85)	(-1.98)
$U \times M$	-0.015^{***}	0.013^{***}	0.026^{***}	-0.076^{***}	0.085^{***}	0.042
	(-5.02)	(3.32)	(4.37)	(-3.68)	(2.82)	(1.37)
$D \times M$	0.032^{***}	-0.019^{***}	-0.018^{***}	-0.105^{***}	0.112^{***}	0.106^{***}
	(7.75)	(-3.68)	(-3.36)	(-3.46)	(3.61)	(3.16)
Characteristics	✓	✓	✓	✓	✓	✓
Months	356	356	356	356	356	356
Stocks	7,575	7,592	7,594	7,573	7,591	7,593
Mean R^2	0.732	0.727	0.728	0.094	0.095	0.095

The table repeats columns (1)–(6) of Table 5 with 365-day market-adjusted uncertainty. As in that table, U, D, and M enter as standardized within-month ranks. Columns (7)–(9) of that table do not involve U and are not repeated. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

D.5 Alternative Explanations

Table A14 adds expected idiosyncratic skewness and MAX to the return regressions on a common sample. Table A15 varies the reversal and liquidity controls in the borrowing cost interactions and replaces borrowing cost with its component orthogonal to size and liquidity. Tables A16 and A17 use the corrected mispricing score released by Stambaugh et al. (2015). Its quintiles are formed before matching option and analyst forecast data. Because the released score uses revised accounting-data timing, the value-weighted sorts in Table A17 establish the direction of the asymmetry over the sample of their paper rather than numerical identity with the published alphas. Table A18 allows DIS and the belief measures to vary with borrowing cost and size.

Table A8: IVOL return slopes under alternative IVOL measures

	(1)	(2)	(3)
M:	Borrowing cost	Ownership breadth	Size
Panel A: Weekly three-year IVOL			
IVOL	−0.141	−0.135	−0.168
	(−0.98)	(−0.94)	(−1.16)
IVOL×M	−0.128***	0.132***	0.123***
	(−5.28)	(4.28)	(3.62)
Mean R^2	0.092	0.093	0.093
Panel B: One-month daily FF3 IVOL			
IVOL	−0.187*	−0.192**	−0.204**
	(−1.92)	(−1.98)	(−2.09)
IVOL×M	−0.121***	0.139***	0.114***
	(−4.91)	(4.93)	(3.65)
Mean R^2	0.086	0.088	0.087
Characteristics	✓	✓	✓
Months	356	356	356
Stocks	7,573	7,591	7,593

The table repeats columns (7)–(9) of Table 5 for the two alternative IVOL measures. These regressions do not involve U; the level of M is included but not shown. Months and stocks refer to Panel A. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A9: IVOL return slopes, firm size, and ownership breadth

	Dependent variable: next-month return		
	(1) Size	**(2)** Breadth	**(3)** Joint
IVOL	-0.218	-0.185	-0.194
	(-1.39)	(-1.20)	(-1.25)
IVOL$\times$Size	0.135^{***}		-0.070
	(3.75)		(-0.93)
IVOL$\times$Breadth		0.159^{***}	0.235^{***}
		(4.94)	(3.48)
Size rank	-0.281^{**}		-0.386^{***}
	(-2.43)		(-3.07)
Breadth rank		0.125^{*}	0.163^{**}
		(1.71)	(2.08)
Characteristics	✓	✓	✓
Months	356	356	356
Stocks	7,591	7,591	7,591
Mean R^2	0.094	0.094	0.098

Monthly Fama–MacBeth regressions of next-month returns, in percentage points. Columns (1) and (2) repeat the size and ownership-breadth specifications in Table 5 on common observations; column (3) includes both interactions and their lower-order terms. IVOL is logged, winsorized at the monthly 1st and 99th percentiles, and standardized within month. Size and breadth enter as standardized within-month ranks. The seven standard characteristics from Table 5 are included; the log market equity control is retained alongside the size rank where shown. Their transformations and the replacement of missing financial characteristics by monthly medians are unchanged. Returns retain the monthly 1st/99th percentile winsorization. All columns use the same 695,597 stock-months from February 1996 through September 2025. An intercept is included. t statistics in parentheses use Newey–West standard errors with six lags. Stars denote significance at the 10, 5, and 1 percent levels.

Table A10: Return slopes by firm size and borrowing cost

	Low borrowing cost		High borrowing cost		Contrasts		
	(1) Large	**(2)** Small	**(3)** Large	**(4)** Small	**(5)** (4)−(2)	**(6)** (4)−(3)	**(7)** Double difference
Panel A: IVOL							
IVOL	−0.102	−0.063	−0.251	−0.417**	−0.354***	−0.166	−0.204*
	(−0.69)	(−0.43)	(−1.44)	(−2.16)	(−3.50)	(−1.53)	(−1.65)
Panel B: Uncertainty and disagreement							
U	−0.018	−0.008	−0.160	0.073	0.081	0.233**	0.223
	(−0.09)	(−0.05)	(−0.87)	(0.41)	(0.62)	(2.02)	(1.45)
D	−0.054	−0.068	−0.129	−0.316***	−0.249**	−0.188*	−0.174
	(−0.46)	(−0.70)	(−1.51)	(−2.82)	(−2.39)	(−1.74)	(−1.25)
$U \times D$	−0.075	0.018	−0.124*	−0.147**	−0.165	−0.023	−0.116
	(−0.87)	(0.21)	(−1.89)	(−2.24)	(−1.54)	(−0.32)	(−0.91)
Panel C: Uncertainty and disagreement with volatility controls							
U	0.311	0.078	0.076	0.408**	0.331*	0.333**	0.566**
	(1.27)	(0.50)	(0.53)	(2.47)	(1.78)	(2.33)	(2.49)
D	−0.043	−0.061	−0.113	−0.221*	−0.160	−0.108	−0.090
	(−0.36)	(−0.64)	(−1.36)	(−1.92)	(−1.50)	(−0.88)	(−0.60)
$U \times D$	−0.048	0.002	−0.128**	−0.180**	−0.182*	−0.052	−0.102
	(−0.55)	(0.02)	(−2.02)	(−2.57)	(−1.67)	(−0.70)	(−0.79)
IVOL	−0.217*	0.111	−0.158	−0.177	−0.288*	−0.019	−0.347*
	(−1.76)	(0.93)	(−1.20)	(−1.06)	(−1.83)	(−0.12)	(−1.92)
Realized vol.	−0.076	−0.231***	−0.130	−0.399***	−0.168*	−0.269***	−0.115
	(−1.10)	(−2.96)	(−1.53)	(−4.37)	(−1.71)	(−2.79)	(−0.99)
Characteristics	✓	✓	✓	✓	—	—	—
Months	356	356	356	356	356	356	356
Stocks	2,206	6,218	5,186	5,423	6,947	7,437	7,545

The table repeats Table 6 with groups formed on market equity rather than ownership breadth. Separate monthly Fama–MacBeth regressions of next-month returns within each group. Groups are halves of market equity within the lowest and highest terciles of implied borrowing cost. Column (7) equals (4)−(2)−[(3)−(1)]. Differences are formed month by month before inference. U is 365-day total uncertainty. U and D enter as standardized logarithms. Stocks in the contrast columns count the union of the groups involved. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4 other than analyst coverage. t statistics in parentheses use Newey–West standard errors with six lags.

Table A11: Earnings windows and the IVOL relation in the full stock sample

Dependent variable:	Average daily return			Monthly excess return		
	(1) Window	**(2)** Other	**(3)** Difference	**(4)** Window	**(5)** Other	**(6)** Total
IVOL	-0.0483^{**}	0.0165	-0.0648^{***}	-0.0979^{***}	-0.1490	-0.2470
	(-2.44)	(1.62)	(-3.65)	(-5.59)	(-0.83)	(-1.31)
Characteristics	✓	✓	✓	✓	✓	✓
Months	355	355	355	355	355	355
Stocks	14,749	14,749	14,749	15,695	15,695	15,695
Mean R^2	0.022	0.063	0.020	0.006	0.047	0.047

Monthly Fama–MacBeth regressions on standardized log one-year daily FF3 IVOL and the firm characteristics described in Section 4 other than analyst coverage, which is not available for all stocks, February 1996–August 2025. Columns (1)–(3) use the average daily return on earnings-window days $[-1, +1]$, on other trading days, and their difference, for stock-months with an announcement. Columns (4)–(6) report the contributions of announcement days and other days to the compounded monthly excess return, followed by the total. The contributions use wealth at the beginning of each day, as defined in Appendix C. Months without an announcement are retained with zero contribution from announcement days. Earnings windows account for fewer than 5 percent of stock trading days. t statistics in parentheses use Newey–West standard errors with six lags.

Table A12: Earnings-window IVOL slopes with lottery characteristics and an alternative IVOL measure

	(1) Baseline IVOL	(2) Add D	(3) Add lottery characteristics
Panel A: One-year daily FF3 IVOL, sample with lottery characteristics			
IVOL	-0.0661^{***}	-0.0389^{*}	-0.0654^{***}
	(-3.14)	(-1.70)	(-2.91)
D		-0.0664^{***}	
		(-3.01)	
MAX			0.0185
			(1.03)
Expected skewness			-0.0279
			(-1.32)
Characteristics	✓	✓	✓
Months	336	336	336
Stocks	7,235	7,235	7,235
Mean R^2	0.045	0.050	0.053
Panel B: Weekly three-year IVOL			
IVOL	-0.0403^{**}	-0.0082	
	(-1.96)	(-0.34)	
D		-0.0777^{***}	
		(-3.33)	
Characteristics	✓	✓	—
Months	355	355	—
Stocks	7,453	7,453	—
Mean R^2	0.044	0.049	—

Monthly Fama–MacBeth regressions of the average daily return on earnings-window days $[-1, +1]$ (percentage points). All estimated columns include the standard firm characteristics. Panel A uses the common sample with MAX and expected idiosyncratic skewness; column (2) adds disagreement to the baseline IVOL regression, and column (3) adds both lottery characteristics. Panel B repeats columns (1) and (2) of Panel A of Table 9 with weekly three-year IVOL in place of one-year daily FF3 IVOL. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A13: Earnings surprises and announcement returns

Dependent variable:	Earnings surprise		Announcement return			
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
U	-0.078^{***}	-0.100^{***}	-0.013	0.033	-0.145^{**}	-0.098
	(-3.92)	(-4.28)	(-0.21)	(0.56)	(-2.25)	(-1.62)
D	-0.038^{***}	-0.016^{*}	-0.201^{***}	-0.165^{***}	-0.213^{***}	-0.372^{***}
	(-3.60)	(-1.89)	(-3.93)	(-3.13)	(-3.44)	(-6.93)
$U \times D$		-0.129^{***}				
		(-4.58)				
Earnings surprise				292.199^{***}		
				(18.75)		
Characteristics	✓	✓	✓	✓	✓	✓
Surprise decile indicators	✗	✗	✗	✗	✓	✗
Sign-specific surprise slopes	✗	✗	✗	✗	✗	✓
Months	355	355	355	355	355	355
Stocks	7,364	7,364	7,364	7,364	7,364	7,364
Mean R^2	0.096	0.110	0.051	0.110	0.174	0.124

Columns (1) and (2) regress the price-scaled earnings surprise, measured from quarterly I/B/E/S actual earnings and the last consensus before the announcement, on the predictors. Columns (3)–(6) regress the three-day market-adjusted announcement return on the predictors with the indicated surprise controls, and column (6) allows separate slopes on negative and positive surprises. Column (5) includes nine indicators for deciles of the price-scaled earnings surprise, formed separately within each announcement month among eligible observations. The lowest decile is the reference category. Coefficients are in percentage points of the surprise or the return. The surprise in column (4) enters as a ratio to price, so its coefficient of 292.199 corresponds to 2.92 percentage points of announcement return for a surprise of 1 percent of price. The belief measures and firm characteristics are dated at the last month-end strictly before trading day -10 relative to the announcement; the consensus used for the earnings surprise is dated separately before the return window. Regressions are estimated across announcements within each month and averaged over months. U is 365-day total uncertainty. U and D enter as standardized monthly ranks. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A14: Uncertainty, disagreement, and lottery characteristics

	(1)	(2)	(3)	(4)
U	-0.059	-0.031	0.043	0.067
	(-0.35)	(-0.19)	(0.27)	(0.44)
D	-0.158^{**}	-0.161^{**}	-0.153^{**}	-0.155^{**}
	(-2.55)	(-2.54)	(-2.51)	(-2.50)
Expected skewness		-0.100		-0.094
		(-1.57)		(-1.48)
MAX			-0.160^{***}	-0.158^{***}
			(-3.62)	(-3.65)
Characteristics	✓	✓	✓	✓
Months	336	336	336	336
Stocks	7,400	7,400	7,400	7,400
Mean R^2	0.097	0.101	0.100	0.103

Monthly Fama–MacBeth regressions of next-month returns (percentage points) in the common sample with expected idiosyncratic skewness, through January 2024 returns. U, D, and expected skewness are standardized ranks; MAX is winsorized, ranked, and standardized. U is 365-day total uncertainty. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A15: Borrowing-cost interactions with reversal and liquidity controls

Main regressors:	U and D			IVOL			IVOL, residualized B	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
U	0.009	−0.006	0.041					
	(0.06)	(−0.03)	(0.28)					
D	−0.125**	−0.141**	−0.128**					
	(−2.02)	(−2.26)	(−2.12)					
$U \times B$	−0.057***	−0.060***	−0.067***					
	(−2.81)	(−2.88)	(−3.17)					
$D \times B$	−0.103***	−0.106***	−0.102***					
	(−3.29)	(−3.42)	(−3.46)					
IVOL				−0.196	−0.181	−0.176	−0.190	−0.196
				(−1.28)	(−1.17)	(−1.43)	(−1.22)	(−1.25)
IVOL$\times B$				−0.137***	−0.139***	−0.145***		
				(−5.27)	(−5.18)	(−5.35)		
$B^{\perp}$							−0.162***	−0.164***
							(−7.96)	(−8.14)
IVOL$\times B^{\perp}$							−0.115***	−0.117***
							(−4.88)	(−4.96)
Characteristics	✓	✓	✓	✓	✓	✓	✓	✓
One-month reversal	✗	✓	✓	✗	✓	✓	✓	✓
Liquidity controls	✗	✗	✓	✗	✗	✓	✗	✗
Months	356	356	356	356	356	356	356	356
Stocks	7,573	7,573	7,573	7,573	7,573	7,573	7,573	7,573
Mean R^2	0.094	0.100	0.106	0.086	0.093	0.100	0.092	0.092

Monthly Fama–MacBeth regressions of next-month returns (percentage points). B is the standardized rank of implied borrowing cost; its level is included but not shown. Columns (1)–(3) and (4)–(6) use identical observations within each block, with the standard characteristics excluding one-month reversal, including it, and adding the bid–ask spread, Amihud illiquidity, and the incidence of zero trading. Columns (7) and (8) replace B with implied borrowing cost residualized on size and the liquidity composite, or on size and the three liquidity measures ($B^{\perp}$). U is 365-day total uncertainty. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

Table A16: Uncertainty, disagreement, and IVOL across mispricing groups

	Mispricing quintile				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
	1 (most underpriced)	2	3	4	5 (most overpriced)
Panel A: IVOL					
IVOL	0.027	0.058	−0.036	−0.009	−0.440**
	(0.15)	(0.37)	(−0.19)	(−0.04)	(−2.00)
Panel B: Uncertainty and disagreement					
U	0.139	0.141	0.092	0.152	−0.223
	(0.75)	(0.81)	(0.46)	(0.71)	(−1.05)
D	0.000	−0.016	−0.042	−0.096	−0.105
	(0.00)	(−0.17)	(−0.38)	(−0.87)	(−0.84)
Panel C: Uncertainty, disagreement, and IVOL					
U	0.426**	0.290*	0.230	0.238	0.148
	(2.42)	(1.74)	(1.33)	(1.18)	(0.76)
D	0.008	−0.013	−0.029	−0.069	−0.091
	(0.10)	(−0.14)	(−0.27)	(−0.63)	(−0.75)
IVOL	−0.340**	−0.177	−0.182	−0.124	−0.456***
	(−2.51)	(−1.57)	(−1.47)	(−0.66)	(−3.21)
Panel D: Uncertainty and disagreement with volatility controls					
U	0.451***	0.327**	0.303*	0.327*	0.244
	(2.68)	(2.07)	(1.89)	(1.80)	(1.36)
D	0.002	−0.020	−0.028	−0.062	−0.082
	(0.03)	(−0.22)	(−0.26)	(−0.57)	(−0.69)
IVOL	−0.278**	−0.092	−0.102	−0.011	−0.240*
	(−2.10)	(−0.80)	(−0.78)	(−0.06)	(−1.73)
Realized vol.	−0.097	−0.141*	−0.199**	−0.255**	−0.387***
	(−1.25)	(−1.72)	(−2.41)	(−2.44)	(−2.79)
MAX, market beta, mispricing polynomial	✗	✗	✗	✗	✗
Months	252	252	252	252	252
Stocks	3,052	3,882	4,300	4,449	4,184

Monthly Fama–MacBeth regressions of next-month returns (percentage points) within quintiles of the corrected mispricing score released by Stambaugh et al. (2015), February 1996–January 2017. *U* is 365-day total uncertainty. *U* and *D* enter as standardized monthly ranks and IVOL and realized volatility as standardized logarithms, calculated before the division into mispricing groups. The regressions include no other controls. *t* statistics in parentheses use Newey–West standard errors with six lags.

Table A17: IVOL return spreads across the mispricing classification

	Mispricing quintile				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
	1 (most underpriced)	2	3	4	5 (most overpriced)
High minus low IVOL	0.546^{*}	0.473^{*}	0.293	-0.126	-1.026^{***}
	(1.85)	(1.84)	(0.86)	(-0.32)	(-2.86)
Months	546	546	546	546	546

Monthly return spreads (percentage points) between high and low IVOL portfolios within mispricing quintiles, using the corrected mispricing score released by Stambaugh et al. (2015) over their calendar window, August 1965 through January 2011, with dependent value-weighted sorts: mispricing quintiles are formed first and IVOL quintiles within each. IVOL is one-year daily FF3 IVOL, and the sample includes all stocks with a valid value, without the option and analyst requirement. t statistics in parentheses use Newey–West standard errors with six lags.

Table A18: DIS and its interactions with borrowing cost and size

	(1)	(2)
U		0.025
		(0.15)
D		-0.081
		(-1.14)
$U \times D$		-0.051
		(-1.08)
DIS	-0.209^{**}	-0.163^{***}
	(-2.02)	(-2.77)
$U \times B$		-0.015
		(-0.46)
$D \times B$		-0.063^{**}
		(-2.13)
DIS$\times B$	-0.079^{***}	-0.033
	(-3.23)	(-0.97)
$U \times S$		-0.039
		(-0.72)
$D \times S$		-0.013
		(-0.32)
DIS$\times S$	0.033	0.050
	(0.91)	(1.17)
Characteristics	✓	✓
B and S main effects	✓	✓
Months	324	324
Stocks	6,651	6,651
Mean R^2	0.093	0.115

Monthly Fama–MacBeth regressions of next-month returns. Column (1) includes DIS, $DIS \times B$, $DIS \times S$, the levels of B and S, and the standard characteristics. Column (2) adds U, D, $U \times D$, $U \times B$, $D \times B$, $U \times S$, and $D \times S$. Here B is implied borrowing cost and S is market equity. U, D, DIS, B, and S enter as standardized monthly ranks; DIS includes IVOL in its construction. Column (5) of Table 10 reports the same specification without the DIS interactions. A tick under Characteristics indicates that the regression includes the firm characteristics described in Section 4. t statistics in parentheses use Newey–West standard errors with six lags.

References

Albagli, E., Hellwig, C., and Tsyvinski, A. (2024). Information aggregation with asymmetric asset payoffs. *Journal of Finance*, 79(4):2715–2758.

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

Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2009). High idiosyncratic volatility and low returns: International and further U.S. evidence. *Journal of Financial Economics*, 91(1):1–23.

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

Bali, T. G. and Cakici, N. (2008). Idiosyncratic volatility and the cross section of expected returns. *Journal of Financial and Quantitative Analysis*, 43(1):29–58.

Bali, T. G., Cakici, N., and Whitelaw, R. F. (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.

Banerjee, S. and Kremer, I. (2010). Disagreement and learning: Dynamic patterns of trade. *Journal of Finance*, 65(4):1269–1302.

Berkman, H., Dimitrov, V., Jain, P. C., Koch, P. D., and Tice, S. (2009). Sell on the news: Differences of opinion, short-sales constraints, and returns around earnings announcements. *Journal of Financial Economics*, 92(3):376–399.

Boehme, R. D., Danielsen, B. R., and Sorescu, S. M. (2006). Short-sale constraints, differences of opinion, and overvaluation. *Journal of Financial and Quantitative Analysis*, 41(2):455–487.

Boyer, B. H., Mitton, T., and Vorkink, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies*, 23(1):169–202.

Breon-Drish, B. (2015). On existence and uniqueness of equilibrium in a class of noisy rational expectations models. *Review of Economic Studies*, 82(3):868–921.

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

Chen, Z. and Petkova, R. (2012). Does idiosyncratic volatility proxy for risk exposure? *Review of Financial Studies*, 25(9):2745–2787.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

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

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

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

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

Engelberg, J., McLean, R. D., and Pontiff, J. (2018). Anomalies and news. *Journal of Finance*, 73(5):1971–2001.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

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

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

Glebkin, S., Gondhi, N., and Kuong, J. (2021). Funding constraints and informational efficiency. *Review of Financial Studies*, 34(9):4269–4322.

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

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

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

Han, Y. and Lesmond, D. A. (2011). Liquidity biases and the pricing of cross-sectional idiosyncratic volatility. *Review of Financial Studies*, 24(5):1590–1629.

Harrison, J. M. and Kreps, D. M. (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.

Herskovic, B., Kelly, B., Lustig, H., and Van Nieuwerburgh, S. (2016). The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics*, 119(2):249–283.

Hou, K. and Loh, R. K. (2016). Have we solved the idiosyncratic volatility puzzle? *Journal of Financial Economics*, 121(1):167–194.

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

Jensen, T. I., Kelly, B., and Pedersen, L. H. (2023). Is there a replication crisis in finance? *Journal of Finance*, 78(5):2465–2518.

Jiang, G. J., Xu, D., and Yao, T. (2009). The information content of idiosyncratic volatility. *Journal of Financial and Quantitative Analysis*, 44(1):1–28.

Johnson, T. C. (2004). Forecast dispersion and the cross section of expected returns. *Journal of Finance*, 59(5):1957–1978.

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

Lehmann, E. L. (1988). Comparing location experiments. *Annals of Statistics*, 16(2):521–533.

Mankiw, N. G. and Reis, R. (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 Wagner, C. (2019). What is the expected return on a stock? *Journal of Finance*, 74(4):1887–1929.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *Journal of Finance*, 42(3):483–510.

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

Nezafat, M., Schroder, M., and Wang, Q. (2017). Short-sale constraints, information acquisition, and asset prices. *Journal of Economic Theory*, 172:273–312.

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

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

Shumway, T. (1997). The delisting bias in CRSP data. *Journal of Finance*, 52(1):327–340.

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

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

Vuolteenaho, T. (2002). What drives firm-level stock returns? *Journal of Finance*, 57(1):233–264.

Wang, J. (1993). A model of intertemporal asset prices under asymmetric information. *Review of Economic Studies*, 60(2):249–282.

Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. In Aghion, P., Frydman, R., Stiglitz, J. E., and Woodford, M., editors, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pages 25–58. Princeton University Press, Princeton, NJ.

Yuan, K. (2006). The price impact of borrowing and short-sale constraints. Working paper, London School of Economics.

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