Pricing News And No News With Heterogeneous Beliefs

Can Gao Brandon Yueyang Han

Abstract

We present a general-equilibrium model in which investors trade risky assets and derivatives while holding heterogeneous beliefs about news content and news intensity. Hindsight winners become richer, and their beliefs carry greater weight in equilibrium prices. Derivative positions are U-shaped across subjective expected returns: investors with extreme beliefs demand convexity exposure through options. When intensity disagreement dominates, realized news raises implied volatility by lifting the market's wealth-weighted assessment of event risk; quiet periods shift wealth toward investors expecting no news, compressing risk-neutral tail probabilities and making 'no news' effectively good news. Discrete news arrivals make intensity disagreement persistent because the wealth-share-decay penalty is small; a continuous news process eliminates it rapidly via derivative losses. When content disagreement dominates, news accelerates market selection and induces mean reversion in risk-neutral variance. The model links prices and volatility to speculation about uncertainty, not just expected returns.

Keywords: Heterogeneous beliefs, news arrival intensity, wealth dynamics, market selection, volatility persistence.

JEL Codes: D53, D83, D84, G12, G41.

Can Gao (corresponding author) is at University of St. Gallen and Swiss Finance Institute, webpage, can.gao@unisg.ch; Brandon Yueyang Han is at University of Maryland, webpage, yhan1@umd.edu. We are grateful for comments from Christian Heyerdahl-Larsen (discussant), Ralph Koijen, Peter Kondor, Pete Kyle, Roberto Marfè, Ian Martin, Albert J. Menkveld (discussant), Dimitris Papadimitriou, Rafael Repullo, Enrique Santana, Paul Whelan, Ole Wilms, and to participants in seminars at CEMFI, EPFL, Geneva, Lausanne, Maryland, St. Gallen, Vienna, and to participants in CEPR ESSFM, St. Gallen FE Workshop, SGF Conference. Can Gao is grateful to research visit funding from CEMFI via Spain's State Research Agency under the María de Maeztu Unit of Excellence Programme for the project CEX2020-001104-M.

"I am an old man and have known a great many troubles,
but most of them never happened."
— Mark Twain

What moves markets is often not news itself, but the anticipation of whether news will arrive. During periods of heightened macroeconomic stress—following a geopolitical shock or a major bank failure—uncertainty spikes, and investors speculate on the likelihood of further systemic events. Similarly, around scheduled events—earnings announcements, macro data releases—some investors position for dramatic fundamental shifts while others confidently bet on a calm post-announcement market. Often, the events unfold without surprises: implied volatility falls and prices drift upward. At other times, price and volatility detonate. These patterns suggest that investors' belief heterogeneity may be multi-dimensional: they disagree not only about whether the future is good or bad, but also about whether the future will be *eventful or calm*.

Professional forecasters' disagreement about the *entire distribution* of macroeconomic outcomes supports this point. (In Internet Appendix I.A.1, we show disagreements of individual forecasters' perceived mean and volatility of the future real GDP growth are at comparable levels.)[1] Historical evidence reinforces the distinction: Koudijs (2016) documents sharp differences in asset price volatility between periods when ships carrying financial news arrived in Amsterdam and periods when they did not—a natural experiment in which the news and no-news distinction is directly observable.

This paper formalizes this distinction. We separate belief heterogeneity about the *content* of news (what will be revealed) from heterogeneity about the *arrival intensity* of news (whether a market-moving innovation is likely to arrive over a given horizon). The distinction matters because the second dimension is naturally priced through derivatives: investors can hold similar directional views and yet take opposing positions on *convexity* if they disagree about the probability of large moves. Both the realization and non-realization of news reallocate wealth across agents with heterogeneous beliefs. These wealth shifts, in

[1] According to the Philadelphia Fed's Survey of Professional Forecasters on GDP growth, the coefficient of variation of individual perceived volatilities (0.41) exceeds that of individual mean forecasts (0.24), and the two dimensions have correlation close to zero ($\rho = -0.10$); see more details in Internet Appendix I.A.1.

turn, reshape the risk-neutral distribution of asset prices and payoffs, and therefore the equilibrium pricing of risk, which takes the form of risk-neutral variance in the main setup. In short: disagreement about *eventfulness* is not a sideshow; it is a state variable that is priced and feeds back into prices and volatility through market selection.

We study an exchange economy in discrete time. Each period, the fundamental news process moves along a trinomial tree: *no news*, *good news*, or *bad news* regarding the terminal payoff of the market. Investors agree that these are the relevant states and how they are related to the terminal payoff. But they *agree to disagree* about the probability of each outcome. Each investor's belief is indexed by two parameters: $v \in (0,1)$ governs the probability that news arrives at all (intensity), and $u \in (0,1)$ governs the conditional probability that arriving news is good (content) for the terminal payoff.[2] Investors trade three non-redundant assets that span the payoff space: the market, a risk-free bond, and a derivative security.

Two modeling choices deliver tractability without shutting down the economics we care about. First, we assume logarithmic preferences throughout this paper, which imply myopic portfolio choice and allow equilibrium prices to be expressed as wealth-weighted harmonic means of future payoffs. This assumption, however, is not essential for the core mechanism.[3] Second, we model the cross-section of beliefs as a continuum with a flexible parametric distribution. These ingredients yield closed-form expressions for (i) the evolution of the wealth-weighted belief distribution ("Mr. Market"), (ii) equilibrium prices, and (iii) risk-neutral variance, which is the measure of the pricing of risk in the trinomial framework—all as functions of two sufficient statistics: the cumulative news count and the cumulative net news.

The central mechanism is simple. Because beliefs among agents differ, investors take

[2] In our trinomial framework, disagreement about expected mean corresponds to heterogeneity in news content beliefs (u), while disagreement about perceived volatility corresponds to heterogeneity in news arrival intensity beliefs (v). The structural decomposition into content and intensity — rather than mean and variance directly — is what allows us to identify the information-theoretic mechanism behind their different persistence properties. With more nodes included, the belief heterogeneity could be extended to higher moments. We explore this in Section 3.

[3] In the Internet Appendix, we extend our framework to general CRRA preferences and show the core mechanism and main results hold accordingly.

opposite sides of trades. Optimists go long in the market and pessimists go short; those who think the market will be calm short volatility against those who think the future will be eventful. In each period, the hindsight winners gain wealth and become more influential in pricing assets and volatility. Disagreement about content primarily reallocates wealth in response to directional outcomes; disagreement about intensity reallocates wealth in response to whether high-impact payoff-relevant events occur at all. These two channels move option-implied variance in opposite directions and imply different dynamics for the market pricing of risk.

The model is extremely tractable. It gives closed-form characterization of equilibrium price and risk-neutral variance. Prices are determined by the wealth-weighted belief distribution over terminal outcomes, and the risk-neutral variance reflects the wealth-averaged belief regarding the market risk premium.[4]

Our first result contrasts the implications of content and intensity disagreements. When disagreement is primarily about *arrival intensity* and investors do not share a strong common directional view, more frequent realized news sustains dispersion in beliefs and generates persistence in risk-neutral variance; quiet periods transfer wealth toward investors who believe in lower news intensity, compress risk-neutral tail probabilities, and support valuations. In this sense, "no news" becomes good news for risky assets precisely when markets disagree about eventfulness. When disagreement is primarily about *content*, frequent news accelerates market selection in the content dimension and induces mean reversion in risk-neutral variance.[5]

Our second result establishes *when* disagreement can survive in equilibrium over a finite horizon. The wealth share of an agent with incorrect beliefs decays exponentially at a rate determined by the Kullback-Leibler divergence between the agent's beliefs and the truth. This rate decomposes additively into a component for arrival-intensity errors and a

[4]The tractability of our closed-form results benefits from log utility, which we view as a disciplined benchmark given the empirical support for the connection between risk-neutral variance and the equity premium documented by Martin (2017), Kremens and Martin (2019), Martin and Wagner (2019), and Gao and Martin (2021).

[5]This mechanism is distinct from the volatility feedback channel of Campbell and Hentschel (1992), in which a representative agent's risk aversion generates a negative price response to anticipated volatility.

component for content errors, the latter weighted by the true arrival rate. In the Poisson limit, where news is rare but impactful, the contribution of the intensity component per trade tends toward zero, so agents with incorrect intensity beliefs survive indefinitely and permanently influence prices. In the Brownian limit, where information flows continuously, the intensity component remains strictly positive, and market selection rapidly eliminates all disagreement about arrival rates.

Our third result links beliefs to derivative positions. After aggregating investors' derivative positions by their subjective beliefs on excess returns, a U-shaped pattern emerges: investors with extreme directional views—both optimistic and pessimistic—simultaneously demand exposure to large market moves, while moderate investors become their natural counterparties.

Our model can be extended so it captures realistic features of price and volatility dynamics. We extend the discrete-time environment to a continuous-time jump-diffusion limit that nests both Brownian and Poisson limits. This extension unifies continuous information flow and discrete innovations, and shows how the endogenous risk-neutral distribution generates asymmetric option prices and implied-volatility skew via state-dependent wealth transfers.

Finally, in a multi-asset economy, we show that speculation localized in one asset propagates to the aggregate market through the wealth channel. This spillover compresses the market beta of non-speculative assets and can transform them into hedges—even absent fundamental linkage—generating safe-haven behavior and excess risk-neutral correlation in equilibrium.

Our paper builds most directly on Martin and Papadimitriou (2022), who study a binary tree economy with a continuum of log-utility investors who agree to disagree about news content. Their framework generates rich dynamics for sentiment, prices, and volume through wealth redistribution. Because news in a binary tree is one-dimensional—each period is either "up" or "down"—investors can disagree about direction but not about whether something will happen. Our model introduces a "no news" state, moving beyond the two-node structure and decoupling disagreement about the mean from disagreement about the

risk or distribution of outcomes. This additional dimension separates disagreement about the content of news from disagreement about its arrival intensity. This dimension is economically distinct: it maps into demand for state-contingent convexity rather than directional trading and generates volatility persistence rather than mean reversion—dynamics that are absent in the binary setting.

A parallel literature studies disagreement about volatility directly. Atmaz and Buffa (2023) develop a dynamic equilibrium model with stochastic volatility in which two investors hold heterogeneous beliefs about the long-run mean of variance, linking volatility disagreement to the variance risk premium and the leverage effect. By contrast, our framework models disagreement over whether news arrives at all—an intensity dimension absent in diffusion economies—and, with a continuum of agents, allows belief dispersion to vary independently of the average belief. This lets us disentangle the equilibrium effects of disagreement from those of shifts in aggregate sentiment. Cao and Ou-Yang (2008) model options trading driven by differences in signal precision. The market selection literature (Sandroni, 2000; Blume and Easley, 2006; Kogan et al., 2006; Borovička, 2020) characterizes long-run survival under incorrect beliefs. We show that selection pressure is affected by the information structure of the news process and differs across disagreement dimensions: errors about arrival intensity are always penalized, while errors about content are penalized in proportion to the true arrival rate.

More broadly, the paper contributes to the literature on speculative trading under heterogeneous beliefs (Harrison and Kreps (1978); Scheinkman and Xiong (2003); Simsek (2013)) by identifying a new channel—arrival-intensity disagreement—through which speculation affects equilibrium prices and volatility dynamics. We work with a continuum of beliefs, as in Geanakoplos (2010) and Atmaz and Basak (2018), which allows us to separately identify the effects of wealth-weighted disagreement and the average belief. Ni, Pan and Poteshman (2008) provide direct empirical evidence of informed volatility trading in options markets, consistent with the mechanism we formalize.

The paper proceeds as follows. Section 1 details the setup of the trinomial tree with heterogeneous beliefs. Section 2 solves the equilibrium, characterizes how disagreement

shapes prices, volatility, and portfolio positions, analyzes the persistence of disagreement over a finite horizon, and discusses the extensions to settings with heterogeneous Bayesian learners, and with agents with CRRA preferences. Section 3 generalizes the framework to richer environments—a general M-node tree, a multi-asset economy, and a jump-diffusion limit. Section 4 concludes. Proofs are in Appendix A.

1 Setup

We consider a finite-horizon discrete-time economy with dates t ranging from 0 to T. The state of the world evolves along a trinomial tree, where each node branches into three potential outcomes for the subsequent period: 'no news', 'good news', or 'bad news' regarding the terminal payoff of the underlying asset. From any single investor's perspective, the probability distribution over these outcomes is identical and independent across periods.

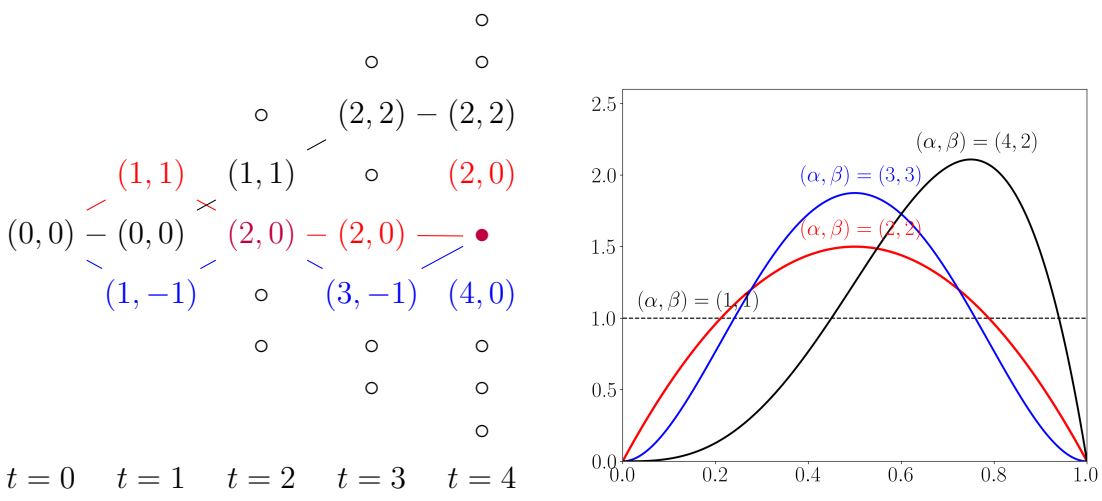
Let n_{mt}, n_{gt}, n_{bt} count no, good, and bad news from 0 to t. Since $t = n_{gt} + n_{mt} + n_{bt}$, two state variables suffice to characterize the equilibrium: news count ν_t and fundamental level (or the net news) n_t

$$\nu_t = n_{gt} + n_{bt}\,, \qquad n_t = n_{gt} - n_{bt}\,.$$

The news count ν_t proxies realized fundamental volatility, as each news event shifts n_t by ± 1. Figure 1a illustrates (ν_t, n_t) along some sample paths between $t = 0$ and $t = 4$. The red path experiences one good and one bad news event before two quiet periods at $t = 3, 4$; the black path experiences two good-news events and is otherwise quiet; the blue path, the most volatile of the three, experiences two good and two bad news events. At $t = 4$, the red and blue paths share the same net news $n_4 = 0$, but the blue path has a higher ν_4; the black and red paths share the same $\nu_4 = 2$ but the red has a lower n_4.

Agents hold heterogeneous, dogmatic beliefs about state evolution and agree to disagree.[6] Each agent's belief is parameterized by a vector $(v, u) \in (0, 1) \times (0, 1)$, where v is the probability of news arrival (good or bad), and u is the conditional probability that

[6] In Section 2.6, we extend the model to incorporate individual Bayesian learning by heterogeneous agents.



(a) Sample paths with (ν_t, n_t)

(b) Beta distributions

Figure 1: **Examples of news sample paths and beta distribution of beliefs**

arriving news is good: v captures perceived news arrival intensity and u reflects relative optimism. Disagreement in the v dimension allows us to study heterogeneity in beliefs about the full distribution of outcomes, rather than limiting attention to heterogeneity in expectations alone. The two parameters v and u yield probabilities for no news (m), good news (g), and bad news (b):

$$m = 1 - v\,, \quad g = vu\,, \quad b = v(1-u)\,, \quad \text{or conversely} \quad v = g + b\,, \quad u = \frac{g}{g+b}\,.$$

The mass of agents with belief (v, u) follows a distribution $f_0(v, u)$, and the wealth distribution across the population evolves over time t, which we call $f_t(v, u)$.[7]

We assume the population (or initial wealth) distribution over (v, u) is separable. A natural parametric choice is the product of two Beta distributions.

$$f_0(v,u) = f_0(v)f_0(u) = \frac{v^{\alpha_{v0}-1}(1-v)^{\beta_{v0}-1}}{B(\alpha_{v0},\beta_{v0})}\,\frac{u^{\alpha_{u0}-1}(1-u)^{\beta_{u0}-1}}{B(\alpha_{u0},\beta_{u0})}\,. \tag{1}$$

The Beta family, $B(\alpha, \beta)$ with $\alpha > 0$, $\beta > 0$, is flexible enough to generate uniform, U-

[7]The trinomial tree setup could be linked to the recombining binary tree (as in Martin and Papadimitriou (2022)) by mapping two binary periods to one trinomial period—for an agent with binary up-move probability $h \in [0, 1]$, good news corresponds to up-up ($g = h^2$), bad news to down-down ($b = (1-h)^2$), and no news to up-down or down-up ($m = 2h(1-h)$)—$u = h^2/[1-2h+2h^2]$ and $v = h^2 + (1-h)^2 \geq 1/2$ restrict agents' beliefs to a one-dimensional subset of $(v, u) \in [0, 1] \times [0, 1]$ with $v \geq 0.5$. We discuss the link between and difference from trinomial and binary trees in Internet Appendix I.A.7.

shaped, or highly concentrated belief distributions on $(0, 1)$. The sum $\alpha + \beta$ governs the degree of agreement—with $\alpha + \beta \to \infty$ corresponding to no disagreement—while the ratio $\alpha/(\alpha + \beta)$ controls the mean. Figure 1b illustrates several configurations.

Asset payoffs are structured to ensure dynamic market completeness. Agents trade three non-redundant assets: (1) a risk-free bond with zero net supply and gross return normalized to one, (2) a long-lived underlying asset in unit supply, and (3) a derivative on the underlying in zero net supply.[8]

All agents exhibit logarithmic preferences over terminal wealth. We use p_t to denote the time-t price of the underlying market asset, and q_t the price of the one-period derivative with payoff x_{t+1} at $t + 1$. At the subsequent nodes, underlying prices are $(p_{t+1,m}, p_{t+1,g}, p_{t+1,b})$ and derivative payoffs are $(x_{t+1,m}, x_{t+1,g}, x_{t+1,b})$. The underlying's terminal payoff is $p_T(n_T)$.

In a dynamically complete market, derivative payoffs do not alter wealth allocations across states. To take advantage of this, we work with a derivative contract that has a payoff uncorrelated with the underlying under the risk-neutral measure, period by period:

$$\mathrm{Cov}_t^*[x_{t+1}, p_{t+1}] = 0\,. \tag{2}$$

This condition also implies that the derivative has *zero delta*. Although this choice is abstract, it is generic. Any non-bond payoff vector satisfying (2), together with the underlying payoff vector p_{t+1} and the risk-free bond $(1, 1, 1)$, forms a basis for the payoff space.[9] Any other derivative payoff can therefore be decomposed into a position in the bond, the underlying, and the derivative satisfying (2). As a result, the cross-sectional distribution of option demands is invariant to the choice of spanning derivative, up to a multiplicative scaling factor. Importantly, (2) does not constrain the derivative to be uncorrelated with the market under agents' subjective beliefs—only under the risk-neutral measure.

[8] Because consumption occurs only at the terminal date T, there is no intertemporal consumption choice to determine the interest rate. The risk-free rate is treated as exogenous. This assumption is commonly adopted in models without intermediate consumption; see, for example Pastor and Veronesi (2012), Atmaz and Basak (2018), and Martin and Papadimitriou (2022).

[9] The bond payoff vector trivially satisfies (2). Any non-bond derivative payoff x_{t+1} cannot be linear in p_{t+1}, since the risk-neutral covariance in (2) would then be proportional to the risk-neutral variance of p_{t+1}, which is strictly positive.

2 Equilibrium

Logarithmic utility implies myopic optimization: an agent with belief (v, u) allocates wealth to $\theta_t^{(v,u)}$ units of the underlying, $\phi_t^{(v,u)}$ of the derivative, and the rest in the risk-free bond, solving $\max_{\theta_t,\phi_t} \mathbb{E}_t^{(v,u)} \log\left[w_t^{(v,u)} + \theta_t^{(v,u)}(p_{t+1} - p_t) + \phi_t^{(v,u)}(x_{t+1} - q_t)\right]$, where $w_t^{(v,u)}$ represents the time t wealth of the individual agent with beliefs (v, u).

In equilibrium, each agent's subjective probability-weighted marginal rate of substitution pins down the risk-neutral probabilities:

$$m_t^* = m \cdot \frac{w_t^{(v,u)}}{w_{t+1,m}^{(v,u)}}, \quad g_t^* = g \cdot \frac{w_t^{(v,u)}}{w_{t+1,g}^{(v,u)}}, \quad b_t^* = b \cdot \frac{w_t^{(v,u)}}{w_{t+1,b}^{(v,u)}}, \tag{3}$$

where (m_t^*, g_t^*, b_t^*) are risk-neutral probabilities; $w_t^{(v,u)}$ is the wealth of the agent with beliefs (v, u) at time t; and $w_{t+1,m}^{(v,u)}$, $w_{t+1,g}^{(v,u)}$, and $w_{t+1,b}^{(v,u)}$ denote this agent's wealth at the respective nodes at time $t+1$. Because the market is dynamically complete, the risk-neutral measure is unique. Despite their differing subjective beliefs, each agent's first-order condition in (3) remains consistent with this unique pricing measure.

Optimal allocations align Arrow-Debreu weights with subjective beliefs (m, g, b). One-period wealth growth is the ratio of subjective belief to the risk-neutral measure: in the no news state, wealth grows by m/m_t^*; in the good state, by g/g_t^*; and in the bad state, by b/b_t^*. Aggregate wealth growth equals the market portfolio change, so one-period state-dependent returns satisfy

$$\frac{p_{t+1,m}}{p_t} = \frac{M_t}{m_t^*}, \quad \frac{p_{t+1,g}}{p_t} = \frac{G_t}{g_t^*}, \quad \frac{p_{t+1,b}}{p_t} = \frac{B_t}{b_t^*}. \tag{4}$$

Here, (M_t, G_t, B_t) are wealth-weighted average beliefs—investors with those beliefs at t allocate fully to the underlying, which we call *Mr. Market*, adopting the terminology of Martin and Papadimitriou (2022).[10]

Equation (4) also implies p_t^{-1} is a sentiment-weighted harmonic mean of next-period

[10]The term originates with Graham (1949), who personified the market as an emotional counterparty whose quotes reflect shifting sentiment rather than fundamental value. In our setting, Mr. Market's beliefs are not irrational but instead aggregate the wealth-weighted beliefs of a heterogeneous population.

prices:

$$p_t^{-1} = M_t p_{t+1,m}^{-1} + G_t p_{t+1,g}^{-1} + B_t p_{t+1,b}^{-1} \,. \tag{5}$$

Thus, sentiment—the wealth-weighted (m, g, b)—drives pricing intertemporally.

2.1 Asset Prices

Although aggregate beliefs are path-dependent, (5) yields tractable pricing. The market price and risk-neutral variance are determined by the wealth-weighted distribution over terminal states, $\Pr(n_T|n_t, f_t)$, which aggregates over all news trajectories from t to T:

Result 1 (Market Price and Risk-Neutral Variance)**.** *Assuming the trinomial tree has T periods and the payoffs are $p_T(n_T)$ with $n_T \in [-T, T]$, the value of the market at time t is given by*

$$p_t^{-1} = \mathbb{E}_t \left[p_T^{-1}(n_T) \right] \,, \tag{6}$$

and the risk-neutral variance of the return from time t to T is given by:

$$\mathrm{Var}_t^* [p_T/p_t - 1] = \mathbb{E}_t \left[p_T(n_T) \right] \mathbb{E}_t \left[p_T^{-1}(n_T) \right] - 1 \,, \tag{7}$$

where the expectation is taken under the wealth-weighted belief distribution over terminal states n_T, given by:

$$\Pr(n_T|n_t, f_t) = \sum_{\substack{\forall\, n_g+n_m+n_b=T-t \\ n_g-n_b=n_T-n_t}} \frac{(T-t)!}{n_m! n_g! n_b!} \int v^{T-t-n_m} (1-v)^{n_m} u^{n_g} (1-u)^{n_b} f_t(v,u) \, \mathrm{d}v\mathrm{d}u \,, \tag{8}$$

conditional on the current fundamental n_t and the wealth distribution f_t.

In limits with no v- or u-disagreement, markets ignore arrivals (n_m) or content (n_g, n_b), and $f(v, u)$ reduces to a marginal distribution. Result 1 holds generally for any initial wealth distribution.

We now specialize the initial wealth distribution to the separable Beta form in (1).[11]

[11] This setting nests the Dirichlet distribution on (m, g, b) when $\alpha_{v0} = \alpha_{u0} + \beta_{u0}$, yielding parameters $(\beta_{v0}, \alpha_{u0}, \beta_{u0})$. The Dirichlet distribution's rotational symmetry implies a weak linkage between content and arrival disagreement; our framework relaxes this.

Comparing (3)–(4), the time-t wealth distribution is

$$f_t(v,u) = \frac{v^{\alpha_{vt}-1}(1-v)^{\beta_{vt}-1}}{B(\alpha_{vt},\beta_{vt})} \frac{u^{\alpha_{ut}-1}(1-u)^{\beta_{ut}-1}}{B(\alpha_{ut},\beta_{ut})}, \tag{9}$$

where $t = n_{gt} + n_{mt} + n_{bt}$ and $(\alpha_{vt}, \beta_{vt}, \alpha_{ut}, \beta_{ut})$ evolve according to the following:

$$\alpha_{vt} = \alpha_{v0} + \nu_t, \quad \beta_{vt} = \beta_{v0} + t - \nu_t, \quad \alpha_{ut} = \alpha_{u0} + \frac{n_t + \nu_t}{2}, \quad \beta_{ut} = \beta_{u0} + \frac{\nu_t - n_t}{2}. \tag{10}$$

The derivations of (9) and (10) are in Appendix A. Under the Beta specification, (8) takes the closed form:

$$\Pr(n_T|n_t, f_t) = \sum_{\substack{\forall n_g - n_b = n_T - n_t, \\ n_g + n_m + n_b = T - t}} \frac{(T-t)!}{n_m! n_g! n_b!} \frac{B(\alpha_{vt} + T - t - n_m, \beta_{vt} + n_m)}{B(\alpha_{vt}, \beta_{vt})} \frac{B(\alpha_{ut} + n_g, \beta_{ut} + n_b)}{B(\alpha_{ut}, \beta_{ut})}. \tag{11}$$

Under the separable Beta specification, heterogeneous agents' pricing is equivalent to that of a representative log-utility agent whose prior is (1) and who updates via Bayes' rule. The representative-agent economy is not dynamically identical—without heterogeneous beliefs there is no trade—but the equivalence highlights a key feature: markets learn collectively through wealth redistribution, with hindsight winners amplifying their beliefs in pricing.

Result 2 (Wisdom of the Crowd)**.** *Pricing in the heterogeneous-agent economy is identical to pricing in an economy with a single log-utility representative agent whose prior over the probability of future news processes coincides with the initial wealth distribution of the heterogeneous economy, and who updates beliefs over time via Bayes' rule.*

The crowd's wisdom can never surpass its most informed members. If the belief distribution includes the truth $(v_{\text{truth}}, u_{\text{truth}})$, prolonged data enable Mr. Market to converge to the truth for a large number of trading periods.[12] In contrast, a common bias in v or u across all investors prevents convergence to the truth. Under Beta-distributed beliefs, the representative agent's sentiment takes a simple closed form:

$$V_t(\nu_t, n_t) = \frac{\nu_t + \alpha_{v0}}{t + \alpha_{v0} + \beta_{v0}}, \quad U_t(\nu_t, n_t) = \frac{n_{gt} + \alpha_{u0}}{\nu_t + \alpha_{u0} + \beta_{u0}}. \tag{12}$$

[12] A large number of trading periods does not necessarily correspond to a long time interval until the terminal payoff T, but may arise from an increased frequency of trading opportunities.

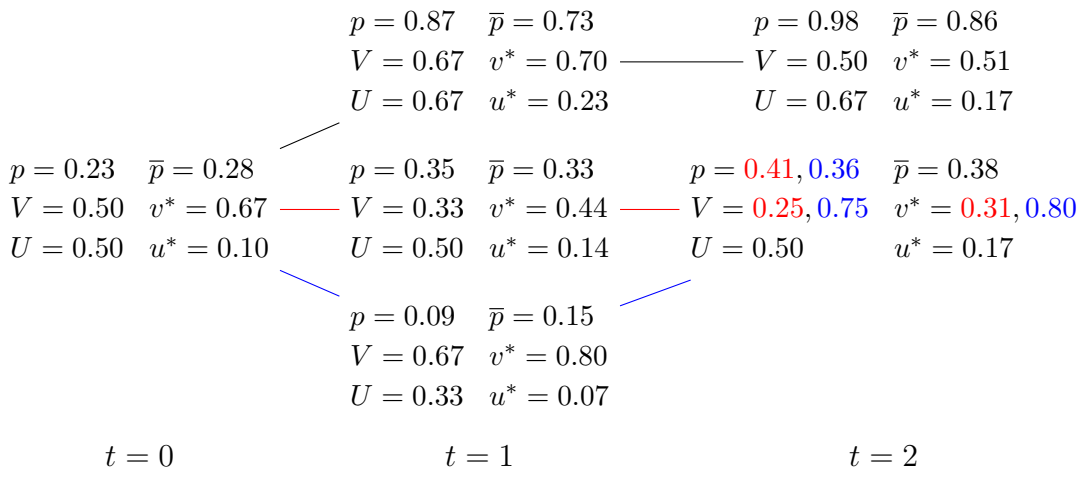


Figure 2: **An example with** $T = 3$. In this example agents' beliefs u and v are both uniformly distributed. p is the market price, $\overline{p}$ is the price in the homogeneous economy where all agents hold the median belief $(v, u) = (1/2, 1/2)$. (v, u) represents Mr. Market's beliefs and (v^*, u^*) is the risk-neutral measure. The nodes at $T = 3$ are omitted from the figure to make it compact: top to bottom nodes are $p = (4.95, 2.22, 1.00, 0.45, 0.20, 0.09, 0.04)$.

When disagreement about v is prominent—i.e., when $\alpha_{v0} + \beta_{v0}$ is not too large—V_t is sensitive to the ratio ν_t/t: higher realized arrival intensity elevates the wealth-weighted v. Conversely, when disagreement is nearly absent ($\alpha_{v0} + \beta_{v0} \to \infty$), finite values of ν_t and t have negligible effect on V_t, analogous to a Bayesian representative agent with a highly precise prior. When disagreement about u is high, $n_{gt} = \frac{1}{2}(\nu_t + n_t)$—the count of good-news arrivals—increases the value of U_t. Disagreement in the heterogeneous economy thus maps directly into the learning speed of the Bayesian representative agent.

Example—A trinomial tree with $T = 3$**.** A three-period example illustrates the core mechanism. We assume agents' beliefs are uniformly distributed in both u and v. Figure 2 reports the pricing and sentiment on several news sample paths between $t = 0$ and $t = 2$: the black path experiences good news followed by no news; the red path is the calmest outcome where no news occurs; the blue path is the most volatile and leads to the same net news as the red one—bad news at $t = 1$ followed by good news. At $t = 1$, good or bad news raises V from 0.5 to 0.67, as high-v agents outperform low-v ones; good news also

shifts wealth to high-u agents, raising U, while bad news lowers it. No news preserves U (no content update) but depresses V, elevating prices—"no news is good news". At $t = 2$, no news on both the black path and the red path further depresses V, as volatility sellers gain more. The blue path's good news offsets the $t = 1$ bad news ($n_2 = 0$, like red), but a higher ν_2 on the blue path yields $V_2 = 0.75$ (vs. red's 0.25), depressing prices relative to the red path despite the good news at $t = 2$.

2.2 Asset Positions

For any derivative security payoff satisfying the zero-delta assumption (2), we prove (details in appendix A) that it must be spanned by the risk-free payoff $(1, 1, 1)$ and the vector $(0, b_t^* - B_t, G_t - g_t^*)$:

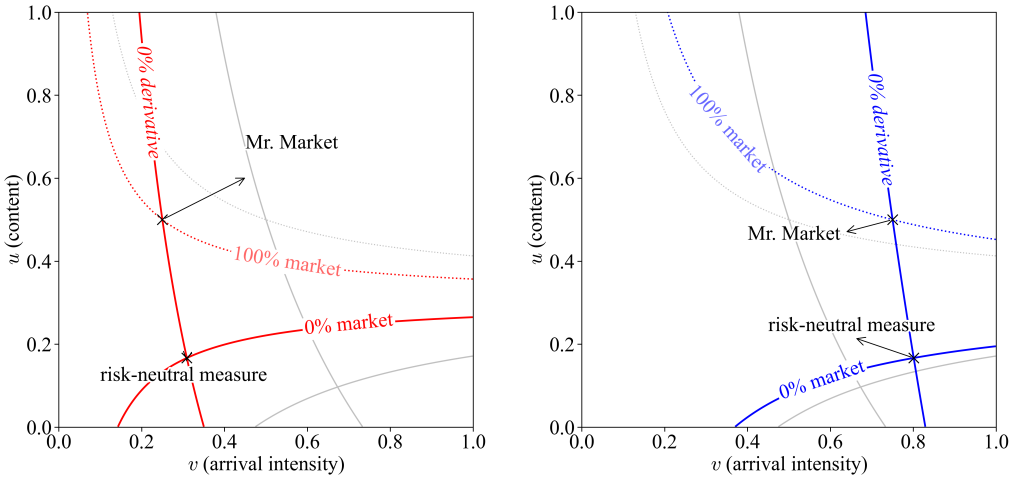
$$(x_{t+1,m}, x_{t+1,g}, x_{t+1,b}) \in \text{span}\left\{(1, 1, 1), (0, b_t^* - B_t, G_t - g_t^*)\right\}. \tag{13}$$

This means any component of the derivative payoff proportional to $(1, 1, 1)$ can be replicated independently without affecting the demand for the derivative itself. Consequently, without loss of generality, we may restrict attention to the remaining component and represent the derivative as $x_{t+1,m} = 0$, $x_{t+1,g} = p_t(b_t^* - B_t)$, $x_{t+1,b} = p_t(G_t - g_t^*)$. The return of this derivative is uncorrelated with the market under the risk-neutral measure—its delta is zero. Therefore, the derivative can be interpreted as a delta-neutral combination of one-period call and put options on the market portfolio, with a strike price at the "no news" realization $p_{t+1,m}$. Since $(0, b_t^* - B_t, G_t - g_t^*) = (0, 1, 1)(b_t^* - B_t) - (0, 0, 1)(v_t^* - V_t)$, the trading of x_{t+1} at t reflects agents' two motives in trading options: speculation on both the arrival intensity and content of news.

Because of market completeness, agents' holdings in Arrow-Debreu securities are isomorphic to positions in the underlying asset and the derivative security. For an agent with beliefs (v, u), the optimal positions in stock and derivative are given by:

$$\frac{\theta_t^{(v,u)} p_t}{w_t^{(v,u)}} = \frac{\mathbb{E}_t^{(v,u)}[p_{t+1}/p_t] - 1}{\text{Var}_t^*[p_{t+1}/p_t]}, \quad \frac{\phi_t^{(v,u)} q_t}{w_t^{(v,u)}} = \frac{\mathbb{E}_t^{(v,u)}[x_{t+1}/q_t] - 1}{\text{Var}_t^*[x_{t+1}/q_t]}. \tag{14}$$

The expressions in (14) resemble the Merton (1969) formula for risky share in a homoge-



(a) Low realised intensity (b) High realised intensity

Figure 3: **Agents' range of beliefs.** As a continuation of Figure 2, the two panels plot the boundaries of agents' asset allocation on the belief domain (v, u) corresponding to (14). The coloured curves partition the belief space into regions of long/short stock and long/short derivative positions. Panel (a) corresponds to the red path (low realised intensity, $\nu_2 = 0$) and panel (b) to the blue path (high realised intensity, $\nu_2 = 2$) from Figure 2. Gray lines show the corresponding boundaries at $t = 0$ for reference.

neous agent setting: risky share $= \mu/\gamma\sigma^2$, where μ is the expected risk premium, $\gamma > 0$ is the agent's risk aversion, and σ^2 is the variance of return. In Merton (1969), physical and risk-neutral variances coincide. Here, agents disagree about physical variance but agree on its risk-neutral counterpart—so the risk-neutral variance, not the physical one, determines portfolio weights.

Martin (2017) showed, with log utility, the simple risk-neutral variance of the market return equals the representative agent's expectation of the risk premium: $\mathrm{Var}_t^*[p_{t+1}/p_t] = \mathbb{E}_t^{(V_t,U_t)}[p_{t+1}/p_t] - 1$. This allows us to express agents' stock positions as

$$\frac{\theta_t^{(v,u)} p_t}{w_t^{(v,u)}} = \frac{\mathbb{E}_t^{(v,u)}[p_{t+1}/p_t] - 1}{\mathbb{E}_t^{(V_t,U_t)}[p_{t+1}/p_t] - 1}\,. \tag{15}$$

Agents who are more bullish than Mr. Market deploy more than 100% of their wealth into the stock market, while those who believe the expected excess return is negative go short. Positions in the underlying scale linearly with an investor's expected return (14), but derivative positions do not.

Figure 3 maps these allocations onto the belief domain (v, u), continuing the $T = 3$ example of Figure 2. The solid line labeled "0% market" separates agents who go long the stock (above) from those who short it (below); agents above the "100% market" boundary invest more than their entire wealth in equity. The solid curve labeled "0% derivative" divides long-derivative (right) from short-derivative (left) agents. Mr. Market's location—the wealth-weighted average belief—and the risk-neutral measure are marked for reference.

Comparing the two panels reveals how realized news reshapes these boundaries. Panel (a) corresponds to a low-intensity history (the red path in Figure 2, with no news arrival by $t = 2$), while panel (b) corresponds to a high-intensity history (the blue path, with two news arrivals). After more news, Mr. Market moves to a higher V, reflecting the wealth transfer toward agents who anticipated eventfulness. However, the population of agents who would stay long in derivatives (right panel) is smaller compared with the case in which no news arrives (left panel)—the blue path has 100% news arrival rate while the red has 0%. The stock position boundaries shift for the same reason. After high realized intensity (panel b), the elevated V raises the one-period risk-neutral variance—equivalently, Mr. Market's expected excess return. By (15), each agent's stock allocation scales inversely with this risk premium: for a given directional view, agents allocate a smaller fraction of wealth to the underlying. The "100% market" boundary therefore moves toward more extreme bullish beliefs, requiring stronger optimism to justify full equity exposure. Conversely, the "0% market" boundary shifts because the conditional state-prices at the successor nodes have changed—the no-news node carries a lower price (calm-world agents have lost wealth), while the news nodes reflect the updated wealth distribution—altering the locus of beliefs at which the expected one-period return equals zero.[13]

2.3 The Effects of Disagreements

News arrival exerts two distinct effects on the market. First, a high historical count of news increases the market's average belief about the future news arrival rate, V_t. Second,

[13]The corresponding plot of portfolio choice boundary for the path with history 'good news' and 'no news' from Figure 2 is in Internet Appendix I.A.4.

frequent news arrivals lead to a concentration of wealth among agents whose beliefs are correct ex-post, and reduce disagreement over news content. The level of "agreement" among investors, $\alpha_{ut} + \beta_{ut} = \alpha_{u0} + \beta_{u0} + \nu_t$, increases with ν_t.

When investors hold a strong common directional view on news content, for instance when the beta distribution on u satisfies $\alpha_{u0} \gg \beta_{u0}$, indicating that investors believe news is almost certainly good, the increase in fundamentals over the remaining periods is approximately given by the number of future news arrivals. Effectively, the belief heterogeneity reduces to a one-dimensional disagreement about v, and price increases with intensity sentiment V_t, which increases in realized news ν_t (see (12)). Symmetrically, when $\alpha_{u0} \ll \beta_{u0}$, investors view news as almost surely bad — price decreases with realized news. Those two hypothetical cases also point to the fact that disagreements in news intensity and content should not be taken literally as disagreements in mean and variance about returns: disagreement about all moments is embedded in this framework and it is more natural to relate them to odd (news content) and even (news intensity) moments of price and returns. To cleanly distinguish disagreements about higher moment theoretically, we would need model more nodes on the tree. In section 3, we explore this general framework and examples with four nodes.

Nevertheless, when investor disagreement on news content does not reflect a strong directional view, the two types of disagreement introduce opposite forces on market sentiment about returns: the arrival of news increases intensity sentiment, producing a more dispersed wealth-weighted belief distribution over future returns, while it also introduces concentration of the wealth-weighted belief distribution through news content updates. Which force dominates depends on the relative level of the two types of disagreement about news.

These opposing effects on market sentiment about returns also translate into opposing effects on the pricing of risk. By changing the wealth-weighted distribution of beliefs over future return realizations, news arrival changes the equilibrium compensation required for bearing future uncertainty, which is summarized by the risk-neutral variance in this setting.

This contrast is best illustrated along sample paths at balanced market sentiment of

news content, i.e. when Mr. Market assigns a 50-50 chance to good and bad future news.[14] We specialize to an exponential terminal payoff, $p_T(n_T) = e^{\sigma n_T}$, under which the distinction between different dimensions of disagreement takes a sharp form.[15]

Result 3 (News Arrival Implications with Intensity and Content Disagreement)**.** *The impact of the realized news count ν_t on asset valuations and risk-neutral variance takes two distinct forms:*

i) Along realized histories with balanced sentiment in news content, when disagreement about news content is not disproportionately large relative to disagreement about the news-arrival rate, i.e., $\alpha_{u0} + \beta_{u0} \geqslant \alpha_{v0}$, a higher realized news count ν_t strictly increases the risk-neutral variance $\mathrm{Var}_t^*[p_T/p_t - 1]$ and lowers the normalized asset price $p_t/e^{\sigma n_t}$.

ii) When all agents agree on the news arrival rate, holding the content sentiment U_t and $t < T - 1$ fixed, a higher realized news count ν_t strictly reduces the risk-neutral variance $\mathrm{Var}_t^*[p_T/p_t - 1]$ and raises the normalized asset price $p_t/e^{\sigma n_t}$.

The intuition is straightforward. Frequent news raises the wealth-weighted average arrival intensity V_t while simultaneously concentrating the cross-sectional distribution of content beliefs. When disagreement about arrival intensity is high, each news event generates substantial wealth transfers, particularly away from derivative sellers whose short-convexity positions are exposed to discrete price movements. In this case, the rise in intensity sentiment dominates, so the market places more weight on extreme future return realizations and demands greater compensation for bearing that risk, which appears as higher risk-neutral variance. When disagreement is lower, by contrast, the concentration of beliefs dominates: news reallocates wealth toward agents with more accurate content beliefs, compressing the wealth-weighted distribution of future returns and lowering the compensation required to

[14] In Internet Appendix I.A.2 we show that part *i*) of Result 3 holds under a wide range of specifications in which intensity disagreement dominates content disagreement.

[15] This specification lends itself naturally to a continuous-time limit, where the interval between trading opportunities shrinks to zero. Unlike an arithmetic payoff, the exponential (or geometric) form guarantees that the payoff remains strictly positive and cannot fall below zero.

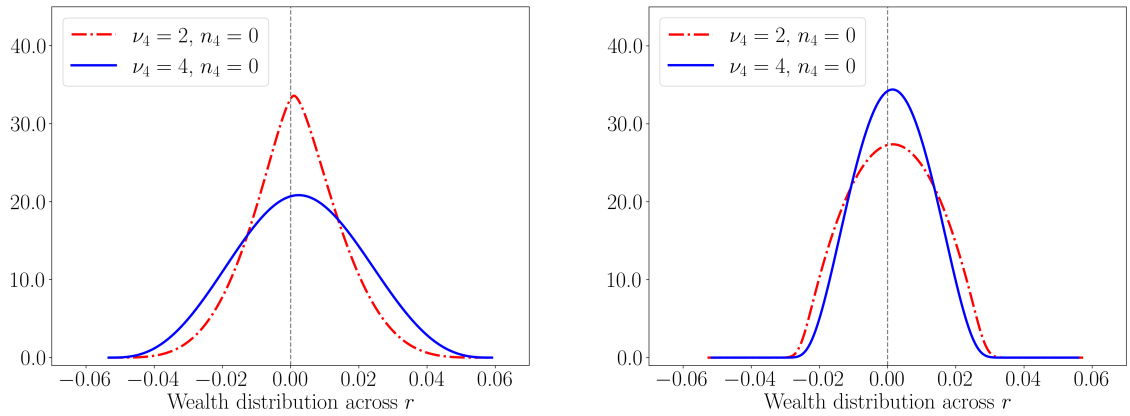
bear future risk. This distinction provides a structural mechanism for the asymmetric response of implied volatility around scheduled events. [16]

No news becomes good news when disagreement about news arrival intensity is high. The market contains agents who assign large probability to extreme outcomes and demand substantial compensation to hold the market. These agents are marginal price-setters when they control significant wealth. A quiet period—one in which no news arrives—is more surprising to high-v types than to low-v types, so wealth flows disproportionately from volatile-world believers to calm-world believers. As calm agents gain pricing influence, the market's effective assessment of tail risk falls, risk-neutral variance compresses, and prices rise. "No news is good news" is thus a statement about the endogenous evolution of risk-bearing capacity: quiet periods relax the aggregate demand for tail-risk compensation by shifting wealth toward agents who perceive less tail risk.

The mechanism is visible in the cross-section of wealth. Figure 4 plots the wealth distribution across investor cohorts indexed by their expected return r, comparing the volatile blue path and the quiet red path from Figure 1a under high and low arrival-intensity disagreement (closed-form expressions are in Internet Appendix I.A.3, Result IA2).

The same realized volatility generates opposite effects on the wealth distribution depending on the regime. Under high arrival-intensity disagreement (Figure 4a), the volatile blue path ($\nu_4 = 4$) disperses wealth toward agents with extreme views, widening the distribution relative to the quiet red path ($\nu_4 = 2$)—consistent with point i) of Result 3. Under low arrival-intensity disagreement (Figure 4b), the same volatile path instead concentrates wealth among agents with accurate directional beliefs, narrowing it—consistent with point ii). This reversal is the core empirical prediction: whether news amplifies or dampens implied volatility depends entirely on the prevailing regime of arrival-intensity disagreement.

[16] An empirical literature documents that option-implied variance reliably increases in the days prior to major announcements (such as FOMC meetings) and collapses immediately afterward, accompanied by a positive drift in the underlying index—often summarized as the "pre-FOMC announcement drift" (Lucca and Moench, 2015; Hu, Pan and Wang, 2022).



(a) High v-disagreement

(b) Low v-disagreement

Figure 4: **Wealth distribution across investor cohorts by expected return r.** The figure illustrates the wealth distribution across investor cohorts as a function of the expected return r. The left panel (4a) assumes a uniform initial wealth distribution across agents, with $(\alpha_{v0}, \beta_{v0}, \alpha_{u0}, \beta_{u0}) = (1, 1, 1, 1)$. In contrast, the right panel (4b) corresponds to low disagreement regarding the news arrival rate, with parameters $(\alpha_{v0}, \beta_{v0}, \alpha_{u0}, \beta_{u0}) = (100, 100, 1, 1)$. The terminal payoff at $T = 6$ is given by the exponential function $e^{0.05 n_T}$. See Figure 1a for an illustration of the possible sample paths leading to the corresponding wealth distributions and gamma positions at $t = 4$.

A Nokia Earnings Event Study. A concrete illustration is provided by the Nokia earnings event of April 22, 2010, when Nokia announced earnings far below expectations and its share price dropped approximately 15% within minutes. Huang, Menkveld and Yu (2021) decompose the central counterparty's intraday exposure on that day and find that return volatility dominated the initial exposure jump—consistent with wealth transferring toward high-v agents upon news arrival. More strikingly, traders *expanded* rather than reduced their positions throughout the afternoon, with portfolio concentration across clearing members contributing substantially to the exposure buildup. Although the positions tracked are unsettled equity trades rather than derivative contracts, the pattern mirrors the wealth-dispersion channel in Figure 4a: under high arrival-intensity disagreement, a news realization disperses wealth toward agents with extreme views, who take larger offsetting positions.[17]

[17] We thank Albert J. Menkveld for pointing out this event study to us.

2.4 How Persistent Are The Disagreements?

While Result 2 shows that prices *aggregate* as if a single Bayesian agent were learning, the individual wealth dynamics driving this aggregation are governed by the KL divergence between each agent's beliefs and the truth. This raises a natural question: how quickly does market selection concentrate wealth among agents with correct beliefs, and does the answer differ across the two dimensions of disagreement?

A key object governing belief selection in this environment is the Kullback–Leibler (KL) divergence. For Bernoulli distributions with parameters a and b, it is defined as $KL(a,b) = a\log\frac{a}{b} + (1-a)\log\frac{1-a}{1-b}$. The KL divergence determines the exponential rate at which agents with incorrect beliefs lose wealth in the limit of frequent trading.

Result 4 (Decay Rate of Individual Wealth Share)**.** *Assume there exists a true data-generating process for fundamental news characterized by* (v_{true}, u_{true})*. The wealth share of an agent with beliefs* (v,u) *decays exponentially at rate* $KL(v_{true},v) + v_{true}\,KL(u_{true},u)$ *per trading period.*

Result 4 characterizes wealth decay *per trading period*. Whether agents with incorrect beliefs are driven out over a finite horizon therefore depends on how this decay rate behaves as trading opportunities increase. Over a fixed calendar horizon $[0,T]$ with N trading periods, the total log-wealth penalty is the sum of N per-period KL terms. In the Poisson limit, the per-period $\mathrm{KL}(v_{\text{true}},v) = O(1/N)$ while $N\to\infty$, so the cumulative penalty remains $O(1)$—bounded but nonzero—allowing incorrect-v agents to retain nontrivial wealth shares over any finite $[0,T]$.

The asymmetric structure of the decay rate has a natural interpretation. Errors in beliefs about the arrival rate of news (v) are always penalized, whereas errors about the direction of news (u) matter only when news actually arrives and are therefore weighted by v_{true}. When the true arrival rate of news is low, accurate beliefs about volatility are thus far more important for survival than accurate beliefs about direction.

A Poisson Limit. The Poisson limit both sharpens the implied volatility dynamics distinction in Result 3 and illustrates the survival mechanism: prices and risk-neutral variance admit closed-form expressions, and the KL penalty for intensity errors tends toward zero per trading period. We partition $[0, T]$ into $N = T/\delta$ periods. Each piece of news is assumed to have a substantial impact on fundamentals, and the terminal payoff remains identical to the discrete-time case: $p_T = e^{\sigma n_T}$. With α_{v0}, α_{u0}, and β_{u0} held constant and independent of N, and β_{v0} defined as $\beta_{v0} = \alpha_{v0}\left(\lambda^{-1}\,\frac{N}{T} - 1\right)$, the cross-sectional averages of (v, u) are given by $\mathbb{E}\, v = \lambda\delta$ and $\mathbb{E}\, u = \alpha_{u0}/(\alpha_{u0} + \beta_{u0})$. In the limit as $N \to \infty$ and $\delta \to 0$, the counts of "good news" and "bad news" events follow Poisson distributions. Taking the limit $N \to \infty$, we obtain the market price and risk-neutral variance.

Result 5 (Poisson Limit of Trinomial Tree with Heterogeneous Beliefs)**.** *In the Poisson limit, the market price at time t is given by*

$$p_t = e^{\sigma n_t}\left[1 - \Lambda_t(e^{\sigma} - 1)\right]^{\alpha_{vt}} {}_2F_1\left(\alpha_{vt}, \alpha_{ut}; \alpha_{ut} + \beta_{ut}; \frac{e^{\sigma} - e^{-\sigma}}{e^{\sigma} - 1 - \Lambda_t^{-1}}\right)^{-1}, \tag{16}$$

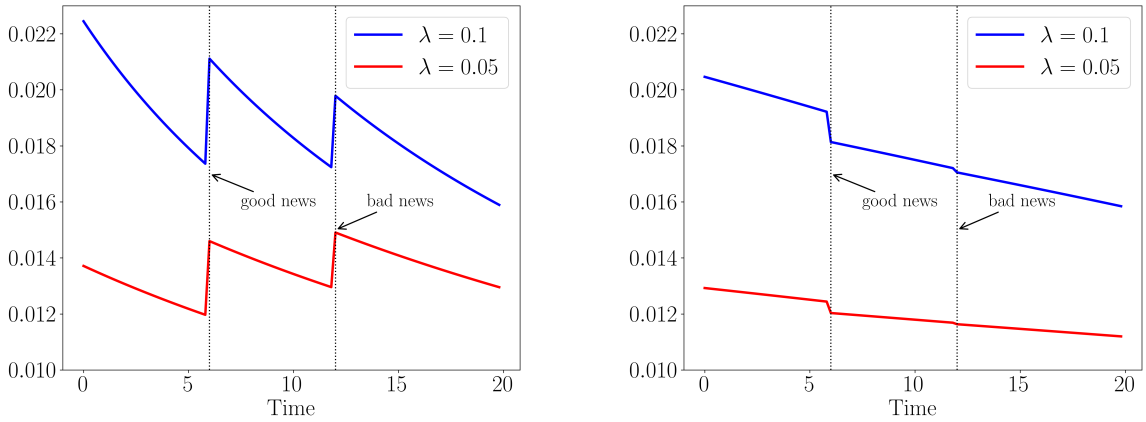
where $\Lambda_t = (T - t)/(\alpha_{v0}\lambda^{-1} + t)$, ${}_2F_1(\ldots)$ denotes the Gauss hypergeometric function, and the expressions for α_{vt}, α_{ut}, and β_{ut} are provided in (10).

The risk-neutral variance of the return from t to T is:

$$\mathrm{Var}_t^*[p_T/p_t] = p_t^{-1} e^{\sigma n_t}\left[1 + \Lambda_t(1 - e^{-\sigma})\right]^{-\alpha_{vt}} {}_2F_1\left(\alpha_{vt}, \alpha_{ut}; \alpha_{ut} + \beta_{ut}; \frac{e^{\sigma} - e^{-\sigma}}{1 + \Lambda_t^{-1} - e^{-\sigma}}\right) - 1\,. \tag{17}$$

Figure 5 illustrates Result 5 for $T = 20$, $\alpha_{u0} = \beta_{u0} = 1$, and terminal payoff $e^{0.05 n_T}$. It presents trajectories for the implied volatility $\sqrt{\mathrm{Var}_t^*[p_T/p_t]/(T - t)}$ following a "good news" event at $t = 6$ and a "bad news" event at $t = 12$, under high and low disagreement about the news arrival rate v. The implied volatility also equals the strike of the simple variance swap of Martin (2011).[18] In the left panel, implied volatility rises after news arrivals, following point *i*) from Result 3. In contrast, the right panel shows implied volatility falling after news events, reflecting point *ii*) of Result 3.

[18]The simple variance swap pays $\mathcal{V}_t/p_t^2$, where $\mathcal{V}_t = \sum_{s=t}^{T-1}(p_{s+1} - p_s)^2$. Its time-$t$ price equals the risk-neutral variance $\mathrm{Var}_t^*[p_T/p_t - 1]$.



(a) High v-disagreement

(b) Low v-disagreement

Figure 5: **Implied volatility trajectories.** The figure plots the implied volatility from t to T under high disagreement ($\alpha_{v0} = 2$) and low disagreement ($\alpha_{v0} = 1000$) about news arrival rate, following a "good news" event at $t = 6$ and a "bad news" event at $t = 12$. The parameters are $\alpha_{u0} = \beta_{u0} = 1$, and the terminal payoff at $T = 20$ is $e^{0.05 n_T}$. Blue lines correspond to a higher expected news arrival rate $\lambda = 0.1$, while red lines correspond to a lower rate $\lambda = 0.05$. We compute this per-period implied volatility as $\sqrt{\mathrm{Var}_t^*[p_T/p_t]/(T-t)}$.

A Diffusion Limit. The interval $[0, T]$ is partitioned into $N = T/\delta$ periods, each of length δ. The terminal payoff is defined as $p_T(n_T) = e^{\sigma\sqrt{\frac{T}{N}} n_T}$, reflecting that each piece of news has a smaller effect on the fundamental, which is scaled as $\sqrt{T/N} n_T$. This is the same setting as in Cox, Ross and Rubinstein (1979). As N grows, we scale the wealth distribution parameters—$\alpha_{v,0} = \lambda\theta N$, $\beta_{v,0} = \alpha_{v,0}(\lambda^{-1} - 1)$—to concentrate v-beliefs around the true arrival rate. We impose this not merely for analytical convenience, but as an equilibrium necessity: if v-dispersion were maintained in the continuous limit, the *infinite frequency* of trading opportunities would, by Result 4, instantly impoverish all agents except those holding the true belief v_{true}. Thus, the Brownian limit isolates an economy where only disagreement about continuous news content (θ) survives to affect risk-neutral variance.

We set $\alpha_{u0} = \frac{1}{2}\lambda\theta N + \frac{1}{2}\eta\sqrt{N}$ and $\beta_{u0} = \frac{1}{2}\lambda\theta N - \frac{1}{2}\eta\sqrt{N}$. The cross-sectional average of (v, u) becomes: $\mathbb{E}\, v = \lambda$, $\mathbb{E}\, u = \frac{1}{2} + \frac{\eta}{2\lambda\theta} N^{-\frac{1}{2}}$, where θ captures disagreement about news content, and the model reduces to the homogeneous belief case when $\theta \to \infty$. In the large-N limit, prices and risk-neutral variance take closed form.

Result 6 (Brownian Limit of Trinomial Tree with Heterogeneous Beliefs)**.** *In the Brownian limit, the market price at time t is*

$$p_t = \exp\left[\frac{1-\phi}{\theta+\phi}\eta\sigma\sqrt{T} + \frac{\theta+1}{\theta+\phi}\eta_t\sigma\sqrt{\phi T} - \frac{1}{2}\frac{\theta+1}{\theta+\phi}(1-\phi)\lambda\sigma^2 T\right]. \tag{18}$$

where $\phi = t/T$ and $\eta_t = n_t/\sqrt{\phi N}$.

The risk-neutral variance of the return from t to T is given by

$$\mathrm{Var}_t^*[p_T/p_t] = \exp\left(\frac{\theta+1}{\theta+\phi}(1-\phi)\lambda\sigma^2 T\right) - 1\,.$$

The Brownian limit corresponds to a binary tree model with only content disagreement: everyone agrees that news arrives with probability λ.[19] Risk-neutral variance is deterministic and unaffected by news realizations—precisely because the arrival-rate disagreement has been eliminated by market selection.

Comparison of the two limits. These two limits yield opposite conclusions for the survival of intensity disagreement. In the Poisson limit, both v_{true} and v tend to zero, and the KL divergence governing expected wealth growth,

$$KL(v_{true}, v) = v_{true}\log\left(\frac{v_{true}}{v}\right) + (1-v_{true})\log\left(\frac{1-v_{true}}{1-v}\right), \tag{19}$$

approaches zero as $N \to \infty$. Although wealth shares decay exponentially in trading time, the per-period decay rate converges to zero. As a result, agents with misspecified beliefs can retain non-negligible wealth over any fixed calendar horizon, even as the number of trading opportunities approaches infinity. This mechanism supports persistent disagreement over news arrival intensity and allows such disagreement to materially affect prices.

By contrast, in the Brownian limit, persistent disagreement about the frequency of news arrivals cannot be sustained. The per-period KL divergence remains strictly positive whenever $v \neq v_{\text{true}}$. As trading opportunities increase, agents with incorrect beliefs repeatedly make small but systematic mistakes, driving their wealth shares to zero. Short-maturity volatility derivatives play a central role in this process: by permitting frequent bets on continuous variance, they turn volatility disagreement into repeated trading gains and losses,

[19] The case when $\lambda = 1$ was studied extensively in Martin and Papadimitriou (2022). At time $t = 0$, equation (18) has a simple expression $p_0 = \exp\left[\frac{\eta}{\theta}\sigma\sqrt{T} - \frac{1+\theta}{2\theta}\lambda\sigma^2 T\right]$.

rapidly reallocating wealth toward agents with more accurate beliefs and rendering residual disagreements irrelevant for prices.

The contrast between the two limits has a broader implication for the role of financial markets in aggregating dispersed beliefs. When fundamentals evolve continuously—the Brownian limit—markets are highly effective information aggregators: frequent trading rapidly concentrates wealth on agents with accurate beliefs, and risk-neutral distributions converge toward fundamentals. When information arrives in discrete lumps—the Poisson limit—the same wealth-dynamics mechanism that disciplines content beliefs is too slow to eliminate intensity disagreement over any finite horizon. Markets fail to reach consensus on *whether an event is likely to happen*. This is not a market failure—agents trade optimally given their beliefs—but it implies that the risk-neutral distribution, and therefore the prices of all state-contingent claims, persistently reflects the wealth-weighted average of potentially incorrect beliefs about eventfulness.

2.5 Convexity Demand

An agent's derivative (convexity) exposure, which is *delta-neutral* by construction in our model, can be related to their beliefs about news intensity and delta exposure:[20]

$$\frac{\phi_t^{(v,u)} q_t}{w_t^{(v,u)}} = \frac{v - v_t^*}{1 - v_t^*} + \frac{v_t^* - V_t}{1 - v_t^*} \frac{\theta_t^{(v,u)} p_t}{w_t^{(v,u)}} .$$

The first component of the gamma portfolio weight comes from the agent's speculation about news arrival intensity, and the second component reflects the agent's speculation about news content through options trading. $\frac{\theta_t^{(v,u)} p_t}{w_t^{(v,u)}}$ scales linearly with an investor's expected return (14), while $\frac{\phi_t^{(v,u)} q_t}{w_t^{(v,u)}}$ scales nonlinearly because the agent's individual belief about v cannot be expressed through a linear stock position.

[20] Agents with a Gamma-neutral position have beliefs that satisfy

$$\begin{aligned}(m, g, b) &= \kappa(M_t, G_t, B_t) + (1 - \kappa)(m_t^*, g_t^*, b_t^*) \\ (1 - v, vu, v(1 - u)) &= \kappa(1 - V_t, V_t U_t, V_t(1 - U_t)) + (1 - \kappa)(1 - v_t^*, v_t^* u_t^*, v_t^*(1 - u_t^*)) ,\end{aligned}$$

where κ is a scalar. This can be verified by substituting subjective beliefs back into (14).

Result 7 (U-Shaped Convexity Demand)**.** *For a cohort of investors sharing the same subjective expected market return r, the wealth-weighted average portfolio weight in the derivative security is determined by their average belief regarding the news arrival rate, $\mathbb{E}[v|r]$:*

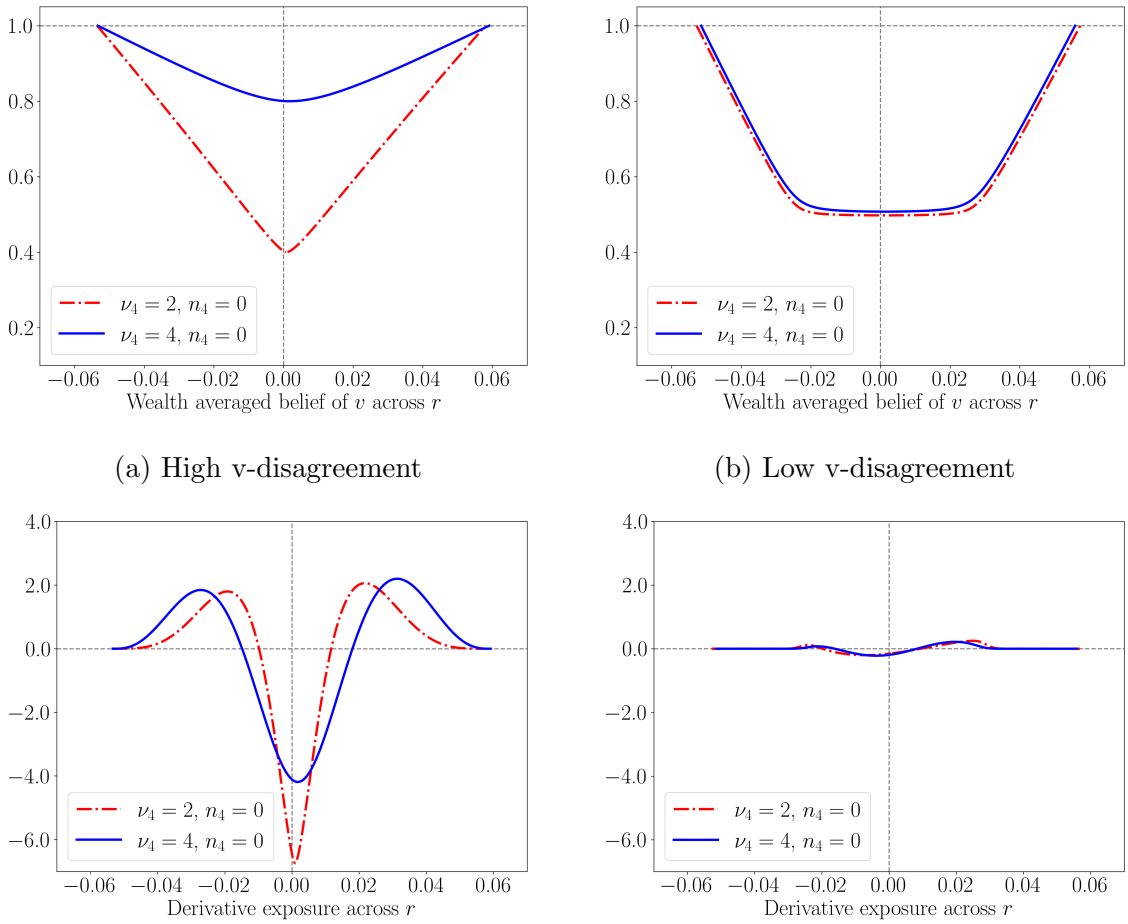
$$\mathbb{E}\left[\frac{\phi_t^{(v,u)} q_t}{w_t^{(v,u)}}|r\right] = \frac{1}{1-v_t^*}\left(\mathbb{E}[v|r] - v_t^* + \frac{v_t^* - V_t}{\mathrm{Var}_t^*[p_{t+1}/p_t]}\, r\right). \tag{20}$$

Because the cohort's average eventfulness belief, $\mathbb{E}[v|r]$, exhibits a pronounced U-shaped relationship with its directional expectation r, derivative positioning is highly asymmetric across the population (closed-form expressions for $\mathbb{E}[v|r]$ are provided in Internet Appendix I.A.3, Result IA2). Investors with extreme directional expectations—both hyper-optimists anticipating large rallies and deep pessimists anticipating crashes—simultaneously hold extreme expectations about the overall news arrival rate ($\mathbb{E}[v|r] \to 1$ at the boundaries). Investors in both tails of the directional distribution therefore hold large long-derivative positions, resembling out-of-the-money call and put buyers respectively. Market clearing requires that investors with moderate expectations—who assign lower probability to news arrivals—supply the offsetting short-derivative positions.

Figures 6a and 6b illustrate this relationship under high and low disagreement in news arrival rates, using $(\alpha_v, \beta_v) = (1, 1)$ and $(100, 100)$, respectively. While both exhibit the U-shape, the low-disagreement case produces a noticeably flatter curve near the center, reflecting stronger consensus among investors. Differences in news arrival rate expectations naturally lead to variation in derivative positions. Figures 6c and 6d show the derivative holdings of investor cohorts, expressed as fractions of total wealth in the economy. These holdings are obtained by multiplying the wealth distribution in Result IA2 with the corresponding portfolio weights from (20).

Since derivative contracts have zero net supply, the long-convexity positions of extreme optimists and pessimists must be absorbed by moderates who sell convexity. The equilibrium derivative price thus reflects the balance between speculative demand from the tails and liquidity supply from the center of the belief distribution.[21]

[21] Consistent with Result 7, Lakonishok et al. (2007) find that sophisticated institutional activity ac-



(a) High v-disagreement

(b) Low v-disagreement

(c) High v-disagreement

(d) Low v-disagreement

Figure 6: **Beliefs of news intensity, market return and derivative position.** The figure reports wealth-weighted average news arrival rate v and derivative exposures for the cohort of investors with expected return r. The left panels (6a and 6c) assume a uniform initial wealth distribution across agents, with $(\alpha_{v0}, \beta_{v0}, \alpha_{u0}, \beta_{u0}) = (1, 1, 1, 1)$. In contrast, in the right panels (6b and 6d), the initial wealth distribution corresponds to low disagreement in the news arrival rate dimension, characterized by $(\alpha_{v0}, \beta_{v0}, \alpha_{u0}, \beta_{u0}) = (100, 100, 1, 1)$. Terminal payoff: $e^{0.05 n_T}$ with $T = 6$. See Figure 1a for the sample paths leading to these wealth distributions at $t = 4$.

Derivative positions offer an empirically observable metric to distinguish between cases of high and low disagreement about the news arrival rate. When disagreement about the news arrival rate is high, investors take significantly larger derivative positions, as shown in Figure 6c. These sizable derivative investments, together with the resulting wealth redistribution, are the central drivers of volatility persistence in Result 3. In contrast, when disagreement is low, $\mathbb{E}[v|r]$ remains close to the market average V_t across most investor cohorts. Derivative portfolio weights are therefore small, leading to limited activity in derivative markets, as illustrated in Figure 6d.

2.6 Extensions

Individual Bayesian Learning. The model can be extended to allow heterogeneous individuals to update their beliefs over time according to Bayes' rule. In this setting, (v, u) serve as agent indices, while the actual content and arrival-intensity beliefs are characterized by $(\tilde{v}, \tilde{u})$. The initial mass of the population across heterogeneous agents is given by (1), and agent (v, u) holds prior beliefs $\tilde{v} \sim \text{Beta}(\zeta v, \zeta(1-v))$ and $\tilde{u} \sim \text{Beta}(\zeta u, \zeta(1-u))$, where the Beta distribution is defined on $(0, 1)$ and its parameters determine the expected value and confidence of the belief. In this parametrization, the prior mean equals (v, u) for $(\tilde{v}, \tilde{u})$, while the concentration parameter ζ governs the dispersion of beliefs: higher ζ corresponds to more precise priors and lower perceived uncertainty. The value of the market at time t is given by the same expression as in (6), where the expectation is taken under a generalized distribution of (11), $\Pr(n_T)$:

$$\sum_{\substack{\forall n_g - n_b = n_T - n_t, \\ n_g + n_m + n_b = T - t}} \frac{(T-t)!}{n_m! n_g! n_b!} \int \frac{B(\zeta v + \nu_t + T - t - n_m, \zeta(1-v) + t - \nu_t + n_m)}{B(\zeta v + \nu_t, \zeta(1-v) + t - \nu_t)} \frac{B(\zeta u + n_{gt} + n_g, \zeta(1-u) + n_{bt} + n_b)}{B(\zeta u + n_{gt}, \zeta(1-u) + n_{bt})} f_t(v, u) \mathrm{d}v \mathrm{d}u.$$

In the limit as $\zeta \to \infty$, agents hold dogmatic beliefs, the terminal states distribution converges to (11), and the value of the market converges accordingly. In Internet Ap-

counts for the bulk of structured convexity positioning, while Ni, Pan and Poteshman (2008) show that this volatility-specific demand carries predictive power distinct from directional order flow.

pendix I.A.4, we illustrate the quantitative effects of finite ζ on equilibrium prices and portfolio boundaries. Individual learning compounds with market selection, amplifying the pricing effects documented in Results 3–4; when ζ is small enough, the effective belief domain compresses sufficiently that the risk-neutral measure can fall outside the support of any agent's posterior beliefs.

CRRA Preferences. Under CRRA preferences with risk aversion γ, the competitive equilibrium is Pareto optimal and characterized by *Negishi weights* $\psi_t(v,u)$ that aggregate heterogeneous beliefs into a social welfare function. The key structural result, proved in Internet Appendix I.A.5, is that Negishi weights evolve as pure likelihood ratios,

$$\frac{\psi_t(v,u)}{\psi_t(v',u')} = \frac{\psi_0(v,u)}{\psi_0(v',u')} \prod_{\tau=1}^{t} \frac{\Pr^{(v,u)}(s_\tau)}{\Pr^{(v',u')}(s_\tau)},$$

independent of γ, the risk-neutral measure, and all equilibrium prices. When $\psi_0(v,u)$ is initialized with the same functional form as the log-utility wealth distribution $f_0(v,u)$, the Negishi weights $\psi_t(v,u)$ track $f_t(v,u)$ exactly for all t, γ, and sample paths.

This preference invariance carries the core results directly. Because the selection rate against any agent is determined by the KL divergence of their beliefs from the truth—independently of γ—the selection dichotomy of Result 4 holds verbatim under CRRA: intensity disagreement survives in the Poisson limit and is eliminated in the Brownian limit. The CRRA consensus measure $\widetilde{\pi}_t(n_T)$ inherits the same news-driven dispersion dynamics as its log-utility counterpart, so the qualitative distinction in volatility dynamics is preserved. One qualification applies to the "no news is good news" valuation result: it requires $\gamma > 1$, the empirically relevant range for equity investors, since the additional curvature of $p_T^{-\gamma}$ relative to $p_T^{1-\gamma}$ is what translates a compressed $\widetilde{\pi}_t$ into a higher price.

3 Richer Environments

The trinomial tree isolates the content–intensity distinction in its most transparent form. Three extensions demonstrate that the same economic forces—wealth-driven selection that

disciplines content beliefs but not intensity beliefs, and the resulting persistence of volatility disagreement—operate in substantially richer settings. In these settings, the same mechanism is reflected in pricing of risk through risk-neutral skewness, volatility skew, and correlation. First, we generalize to an M-node tree that accommodates multi-dimensional fundamental news and multiple assets. Second, we study a two-asset economy in which speculation localized in one asset propagates to the broader market through the wealth channel, endogenously generating safe-haven behavior and excess risk-neutral correlation without any fundamental linkage. Third, we develop a continuous-time jump-diffusion limit that nests both the Poisson and Brownian limits within a single framework, unifying gradual information flow and discrete innovations.

3.1 The M Tree

We begin by extending the trinomial tree to M nodes per period, accommodating richer state spaces—multiple fundamental variables, asset-specific news, or combinations thereof—within the same analytical framework. An agent's subjective belief is defined as $\mathbf{h} = (h_1, h_2, \ldots, h_M)$, where h_i indicates the probability of the i-th node of the tree occurring from top to bottom.[22]

The market is dynamically complete. There are K non-redundant risky assets with one unit of net supply, $M - K - 1$ non-redundant derivatives with zero net supply, and one risk-free bond with zero net supply. The payoff structure of all risky assets at the terminal date T is common knowledge, defined as functions of the terminal cumulative news $\mathbf{s}_T = (s_{1,T}, s_{2,T}, \ldots, s_{M,T})$, where $s_{i,T} \in \{0, 1, \ldots, T\}$ denotes the realization count of news i (node i) between time 0 and T.

Logarithmic utility implies myopic optimization: an agent with belief $\mathbf{h} = (h_1, h_2, \ldots, h_M)$ allocates wealth to $\theta_t^{(\mathbf{h})} = (\theta_{1,t}^{(\mathbf{h})}, \theta_{2,t}^{(\mathbf{h})}, \ldots, \theta_{K,t}^{(\mathbf{h})})$ units of the risky assets, $\phi_t^{(\mathbf{h})} = (\phi_{1,t}^{(\mathbf{h})}, \phi_{2,t}^{(\mathbf{h})}, \ldots, \phi_{M-K-1,t}^{(\mathbf{h})})$

[22] Alternatively, it is possible to follow our formulation of (v, u) in the trinomial tree case and parameterize $\mathbf{h}$ as $h_1 = (1 - v_1)$, $h_2 = v_1(1 - v_2)$, $\ldots$, $h_M = v_1 v_2 \cdots v_{M-1}$, where $v_i \in (0, 1)$.

units of the zero-supply derivatives, and the remainder in the risk-free bond, solving:

$$\max_{\theta_t^{(\mathbf{h})},\phi_t^{(\mathbf{h})}} \mathbb{E}_t^{(\mathbf{h})} \log \left[w_t^{(\mathbf{h})} + \sum_{i=1}^{K} \theta_{i,t}^{(\mathbf{h})}(p_{i,t+1} - p_{i,t}) + \sum_{i=1}^{M-K-1} \phi_{i,t}^{(\mathbf{h})}(x_{i,t+1} - q_{i,t}) \right], \tag{21}$$

where $w_t^{(\mathbf{h})}$ represents the time t wealth of individual agent $\mathbf{h}$.

The total market value $p_t = \sum_{i=1}^{K} p_{i,t}$, representing the sum of all assets with positive supply, satisfies:

$$p_t^{-1} = \sum_{j=1}^{M} H_{j,t} p_{t+1}^{-1}(j, \mathbf{s}_t), \tag{22}$$

where $H_{j,t}$ is the wealth-weighted consensus probability of node j occurring in the next period, and $p_{t+1}(j, \mathbf{s}_t)$ is the market value if node j is realized at time $t+1$.

Similarly to (4), the risk-neutral measures are given by:

$$h_{j,t}^{*} = H_{j,t}\, p_t(\mathbf{s}_t) p_{t+1}^{-1}(j, \mathbf{s}_t). \tag{23}$$

Given the aggregate belief $H_{j,t}$, we can solve for all asset prices backward from the terminal date T using equations (22) and (23).

Result 8 (Sentiment and Risk-Neutral Measure)**.** *Assuming the news path history from time 0 to time t is $\mathbf{s}_t = (s_{1,t}, \ldots, s_{M,t})$ and agents' beliefs follow a Dirichlet distribution $Dir(\beta_{1,0}, \ldots, \beta_{M,0})$, the aggregate belief for node j is:*

$$H_{j,t} = \beta_{j,t} \left(\sum_{i=1}^{M} \beta_{i,t} \right)^{-1}, \tag{24}$$

where $\beta_{j,t} = \beta_{j,0} + s_{j,t}$. The path-dependent risk-neutral measure between t and T is:

$$h_{t\to T}^{*}(\mathbf{s}_T) = \frac{p_t(\mathbf{s}_t)}{p_T(\mathbf{s}_T)} \frac{(T-t)!}{(s_{1,T} - s_{1,t})! \cdots (s_{M,T} - s_{M,t})!} \frac{B(\beta_{1,T}, \ldots, \beta_{M,T})}{B(\beta_{1,t}, \ldots, \beta_{M,t})}. \tag{25}$$

A Two Asset Economy. In Internet Appendix I.A.6.2, we develop continues time limit of the two-asset economy in which three types of jumps can occur: idiosyncratic jumps in asset A, idiosyncratic jumps in asset B, and simultaneous co-jumps affecting both. Investors hold heterogeneous beliefs over all three intensities, governed by a Dirichlet distribution. The resulting continuous-time limit yields closed-form expressions for market prices, market shares, and risk-neutral variances via a single one-dimensional Fourier integral.

This environment delivers an additional prediction that the Brownian limit cannot generate: the risk-neutral correlation between asset returns strictly exceeds the physical correlation, $|\rho_t^*| > |\rho_t|$. The amplification arises because the log-utility pricing kernel disproportionately weights the co-jump state, in which the aggregate market suffers the largest loss. Each realized co-jump further increases the wealth-weighted co-jump intensity $\lambda_{AB,t}^*$, reinforcing the effect over time. Thus the same mechanism that generates volatility persistence through intensity disagreement in the single-asset economy (Result 3) produces *correlation persistence* in the multi-asset economy: a history of co-jumps elevates future risk-neutral correlation, while quiet periods gradually compress it.[23]

This result suggests that belief heterogeneity can simultaneously shape both the level and the co-movement structure of risk-neutral distributions across assets—linking the dispersion of directional and volatility views to observable quantities in index and single-name option markets.[24]

3.2 A Market with Jump and Diffusion

We previously distinguished between news and no-news outcomes, with news reflecting rare, large fundamental changes. In practice, however, fundamentals may still adjust marginally on no-news dates, even when such changes are too small to be newsworthy. We consider a model with $M = 4$ nodes and a single asset ($K = 1$), where nodes 1 and 4 correspond to rare, news-type jumps, while nodes 2 and 3 capture frequent but non-news incremental adjustments. This richer state space allows disagreement to shape the distribution of future outcomes more flexibly, separating beliefs about rare jump risk from beliefs about more continuous fluctuations.

Taking the continuous-time limit, we partition the interval $[0, T]$ into $N = T/\delta$ periods.

[23] Driessen, Maenhout and Vilkov (2009) show that implied correlations systematically exceed realized correlations, driving a structural wedge between index and individual variance risk premia.

[24] Investors who speculate that dispersion will rise are effectively shorting market volatility. If they emerge as hindsight winners, the implied volatility of the market is suppressed as the market index drifts upward. For recent empirical examples of dispersion trading, see "Wall Street seems calm. A closer look shows something else." (New York Times, June, 2024) and "A $12bn dispersion trade" (Financial Times, June, 2017).

The terminal payoff is given by $p_T = \exp\left[\omega\left(s_{1,T} - s_{4,T}\right) + \sigma\sqrt{\frac{T}{N}}\left(s_{2,T} - s_{3,T}\right)\right]$, where ω denotes the jump size and σ the scaled diffusion volatility. We scale the parameters of the Dirichlet distribution for the agent's initial beliefs as $\beta_{2,0} = \frac{1}{2}\theta N + \frac{1}{2}\eta\sqrt{N}$, $\beta_{3,0} = \frac{1}{2}\theta N - \frac{1}{2}\eta\sqrt{N}$, with $\beta_{1,0}$ and $\beta_{4,0}$ held constant and independent of N. As in Result 8 and Result 5, we express agents' time-t beliefs by $\beta_{j,t} = \beta_{j,0} + s_{j,t}$ and use $\lambda_j = \beta_j/\theta T$ to represent the corresponding expected jump intensity. Result 8 yields the market price and the risk-neutral MGF of the log return. Together, these objects characterize not only the market's valuation of the asset, but also how disagreement over the distribution of future outcomes is reflected in the pricing of variance, skewness, and other higher-order risks.

Result 9 (A Jump-Diffusion Limit of Market)**.** *Let $\phi = t/T$ and $\eta_t = (s_{2,t} - s_{3,t})/\sqrt{\phi N}$. In the Jump-Diffusion limit, the market price is given by*

$$\begin{aligned} p_t =& e^{\omega\left[(\beta_{1,t}-\beta_{4,t})-(\beta_{1,0}-\beta_{4,0})\right]}\,\mathscr{J}(-1;\phi,\theta)^{\beta_{1,t}}\,\mathscr{J}(1;\phi,\theta)^{\beta_{4,t}} \\ & \exp\left[\frac{1-\phi}{\theta+\phi}\eta\sigma\sqrt{T} + \frac{\theta+1}{\theta+\phi}\eta_t\sigma\sqrt{\phi T} - \frac{1}{2}\frac{\theta+1}{\theta+\phi}(1-\phi)\sigma^2 T\right], \end{aligned} \tag{26}$$

where $\mathscr{J}(\chi;\phi,\theta) = 1 - (1-\phi)(e^{\chi\omega} - 1)/(\theta+\phi)$.

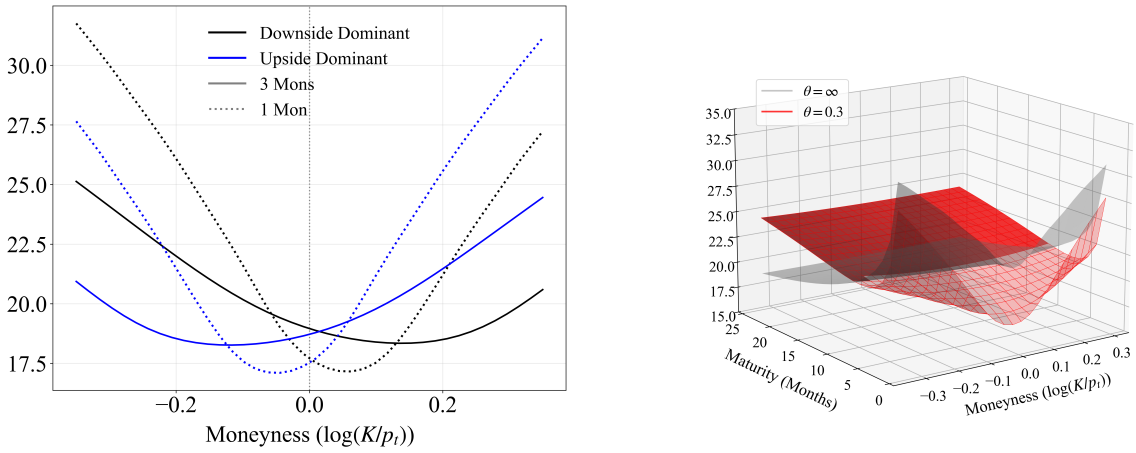
The moment-generating function of the log return under the risk-neutral measure is

$$\begin{aligned} \mathbb{E}_t^*[e^{\chi\log(p_T/p_t)}] = \mathbb{E}_t^*\left[\left(\frac{p_T}{p_t}\right)^{\chi}\right] =& \left[\frac{\mathscr{J}(\chi-1;\phi,\theta)}{\mathscr{J}(-1;\phi,\theta)^{1-\chi}}\right]^{-\beta_{1,t}}\left[\frac{\mathscr{J}(1-\chi;\phi,\theta)}{\mathscr{J}(1;\phi,\theta)^{1-\chi}}\right]^{-\beta_{4,t}} \\ & \exp\left[\frac{\chi(\chi-1)}{2}\left(\frac{\theta+1}{\theta+\phi}\right)(1-\phi)\sigma^2 T\right]. \end{aligned} \tag{27}$$

Evaluating (27) at $\chi = 2$ and subtracting one yields the risk-neutral variance of the simple return, $\mathrm{Var}_t^*[p_T/p_t]$. Risk-neutral skewness and higher-order moments follow analogously from this moment-generating function. [25]

Result 9 unifies the two polar cases we have shown: the continuous Gaussian diffusion in Result 6 and the pure Poisson jump process in Result 5. Setting the jump size $\omega = 0$ yields the Brownian limit, while setting the volatility $\sigma = 0$ yields the Poisson limit, with $\alpha_{ut} = \beta_{1,t}$, $\beta_{ut} = \beta_{4,t}$, $\alpha_{vt} = \beta_{1,t} + \beta_{4,t}$, and $\beta_{vt} = \theta N$. It provides a natural extension of

[25] For the risk-neutral variance to be well defined, the jump size and the degree of disagreement must satisfy $\mathscr{J}(1;0,\theta) = 1 - (e^{\omega} - 1)/\theta > 0$, i.e. $\theta > e^{\omega} - 1$. Too much disagreement (when θ is too small) could result in an unbounded second risk-neutral moment, or equivalently an unbounded second moment of the SDF (see Martin and Shi (2026)).



(a) Volatility Smile (Smirk)

(b) Implied Volatility Surface

Figure 7: **Implied Volatility Dynamics.** This figure illustrates the Black-Scholes implied volatility generated by the jump-diffusion model with heterogeneous agents, evaluated with continuous diffusion volatility $\sigma = 0.15$, jump size $\omega = 0.10$ and $t = 12$ months.

homogeneous-belief models, such as Cox and Ross (1976), Merton (1976), to settings with heterogeneous beliefs.

Wealth is endogenously reallocated toward the ex post correct agent. Following a jump realization, the market's wealth-weighted belief assigns a higher likelihood to future jumps, reflecting the same mechanism that generates volatility persistence under disagreement about news arrival rates, as characterized in Result 3.

The risk-neutral MGF (27) maps directly into an implied volatility surface via the fast Fourier transform (Carr and Madan, 1999). A large literature, including Campbell and Hentschel (1992), documents the asymmetric pricing of downside jump risk; Figure 7 shows how belief heterogeneity shapes these asymmetries.

Figure 7a isolates the volatility skew for short-dated options (1 and 3 months maturities) with high disagreement ($\theta = 0.3$). Holding total expected jump intensity constant ($\lambda_{total} = 4.0$), it compares two cases: bad news dominant ($\lambda_1 = 0.5, \lambda_4 = 3.5$, black) and good news dominant ($\lambda_1 = 3.5, \lambda_4 = 0.5$, blue). The left-skew observed in stock option markets (Bates (1991), Rubinstein (1994)) and the right-skew seen in VIX and commodity option markets (Mencia and Sentana (2013), Jia, Ruan and Zhang (2021)) emerge because wealth shifts toward agents whose jump predictions are realized, amplifying their influence on

market prices. The 1-month dashed curves are steeper in the tails than the 3-month solid curves, reflecting the greater sensitivity of short-dated options to discrete news arrivals.[26] Figure 7b plots the implied volatility surface against moneyness $\log(K/p_t)$ and remaining maturity (1 to 24 months) under balanced jump expectations ($\lambda_1 = \lambda_4 = 2.0$). It contrasts a homogeneous baseline with no disagreement ($\theta = \infty$, transparent gray) against high disagreement ($\theta = 0.3$, red).

While Result 9 captures the essential dynamics of asymmetric skewness and time-varying variance, it does so within a specific parametric scope. The model endogenizes disagreement over the *intensity* of the jumps ($\beta_{1,t}$ vs. $\beta_{4,t}$), but assumes the physical jump size ω is fixed and universally agreed upon. It generates realistic skewness, but the kurtosis is structurally bounded by the fixed jump magnitude. Extending the framework to allow for heterogeneous beliefs over continuous jump-size distributions (e.g., a double-exponential jump limit) represents a natural direction for future research to capture richer tail behavior.

4 Conclusion

This paper develops a tractable general-equilibrium framework for pricing *news* and *no news* when investors disagree not only about the direction of fundamentals, but also about the likelihood of new information's arrival. The distinction between disagreement about news *content* and news *arrival intensity* is economically critical. It naturally separates directional trading from volatility speculation, driving distinct price dynamics and distinct implications for the pricing of risk through endogenous wealth redistribution.

Our central theoretical result is a sharp distinction in how disagreement and speculation shape the dynamics of risk pricing. When disagreement about eventfulness is high, realized news sustains belief dispersion, generating a positive response of implied volatility to the contemporaneous volatility induced by news realizations. Quiet periods, in contrast,

[26] Andersen, Fusari and Todorov (2015) decompose implied volatility into continuous and jump components, finding that the jump risk premium is highly episodic and heavily concentrated at short maturities. Kelly, Pástor and Veronesi (2016) and Dubinsky et al. (2019) document an inverted term structure of implied volatility immediately preceding anticipated jump events like earnings or political elections.

transfer wealth toward low-intensity believers, compressing risk-neutral tail probabilities and supporting valuations—making "no news" effectively good news. Conversely, when disagreement is primarily about content, frequent news accelerates market selection, concentrates wealth among accurate directional forecasters, and induces mean reversion in implied variance. The framework explains why volatility disagreement can survive in equilibrium: because wealth transfer among agents caused by discrete jumps is episodic, errors regarding event arrival intensity are penalized at a rate that tends toward zero per trading period in the continuous-time limit. This allows heterogeneity over eventfulness to persist and shape prices over extended horizons, whereas disagreement over continuous diffusion is rapidly eliminated by continuous price observation.

These endogenous wealth dynamics generate four testable predictions for future empirical work. First, the response of implied volatility to scheduled events hinges entirely on the prevailing regime of arrival intensity disagreement. High disagreement implies volatility persistence and low disagreement implies mean reversion. In particular, when option volume and open interest are elevated—reflecting greater disagreement about implied event probabilities—the VIX should display increased persistence. Second, the cross-section of convexity demand is U-shaped: extreme directional optimists and pessimists aggressively demand convexity (long-gamma), while moderates supply it. Third, intense localized speculation in a single asset propagates through the market portfolio, endogenously inflating risk-neutral correlations and compressing the market betas of non-speculative assets. Finally, short-dated options exhibit pronounced skewness and heightened sensitivity around discrete, jump-type events compared to periods of continuous diffusion.

Our framework highlights a broader message: markets price not only beliefs about *what will happen*, but also beliefs about *whether anything will happen*.

A Mathematical Appendix

Proof of equations (9) and (10). Comparing the individual wealth growth and aggregate wealth growth equations in (3) and (4), we derive the following equation for the

change of wealth distribution:

$$\frac{f_{t+1}(m,g,b)}{f_t(m,g,b)} = \begin{cases} \frac{m}{M_t}, & \text{'no news' at } t+1 \\ \frac{g}{G_t}, & \text{'good news' at } t+1 \\ \frac{b}{B_t}, & \text{'bad news' at } t+1 \end{cases}$$

Solving this recursively backward from t to 0, the wealth distribution $f_t(v,u)$ must satisfy $f_t(v,u) = \lambda_t\, f_0(v,u)\,(1-v)^{n_{mt}}(vu)^{n_{gt}}(v(1-u))^{n_{bt}}$, where the constant λ_t is path-dependent.

Integrating both sides over $(v,u) \in (0,1)^2$, we have $1 = \int_0^1\int_0^1 \lambda_t\, f_0(v,u)\,(1-v)^{n_{mt}}(vu)^{n_{gt}}(v(1-u))^{n_{bt}}\,\mathrm{d}v\,\mathrm{d}u$, which implies $\lambda_t = B(\alpha_{v0},\beta_{v0})B(\alpha_{u0},\beta_{u0})(B(\alpha_{vt},\beta_{vt})B(\alpha_{ut},\beta_{ut}))^{-1}$, where $\alpha_{vt} = \alpha_{v0} + n_{gt} + n_{bt}$, $\beta_{vt} = \beta_{v0} + n_{mt}$, $\alpha_{ut} = \alpha_{u0} + n_{gt}$, and $\beta_{ut} = \beta_{u0} + n_{bt}$.

Converting the notation of (n_{mt}, n_{gt}, n_{bt}) into (ν_t, n_t), we obtain (9) and (10). □

Proof of Result 1. By (5), we compute $1/p_t$ by solving it backward from T to t. It is the linear combination of wealth-weighted aggregate beliefs across all possible future sample paths between t and T multiplied by the inverse terminal payoff $1/p_T(n_T)$.

To formally prove this statement, let (a,b,c) be the number of moves that correspond to (high, middle, low) on a sample path from t to T. They must satisfy $a - c = n_T - n_t$ and $a + b + c = T - t$. This can be solved in terms of b as $a = \frac{1}{2}(T-t+n_T-n_t-b)$ and $c = \frac{1}{2}(T-t-n_T+n_t-b)$, and (a,c) are integers only when b and $T-t-(n_T-n_t)$ are both odd or even numbers ($T-t+(n_T-n_t)$ always shares parity with $T-t-(n_T-n_t)$).

Since the arguments for the positive and negative cases of $n_T - n_t$ are symmetric, we focus on $n_T-n_t > 0$. If $T-t-(n_T-n_t)$ is an odd number, b starts from 1; otherwise, it starts from 0. The size of b cannot exceed $T-t-(n_T-n_t)$ for c to stay non-negative. Given b, a and c are determined. There are $\frac{(T-t)!}{a!b!c!}$ possible paths that satisfy this characteristic, i.e., a times high, b times middle, and c times low. These paths can be merged. To show this, it suffices to show the three basic cases: 1) two trips of 'middle-low' and 'low-middle' give the same risk-neutral probability given the same initial condition; 2) two trips of 'middle-high' and 'high-middle' are equivalent; 3) two trips of 'low-high' and 'high-low' are equivalent.

The risk-neutral variance of the return from t to T, p_T/p_t-1, is given by $p_t^{-2}\,\mathbb{E}_t^*[p_T^2-p_t^2] = \mathbb{E}_t[p_T(n_T)]/p_t - 1$. Substituting into (6), we obtain (7). □

Proof of Result 2. Let $f_0(v,u)$ denote the prior distribution of (v,u) at time 0. Given

n_{gt}, n_{mt}, and n_{bt} realizations of good, no, and bad news up to time t, the posterior is

$$f_t(v,u) = \frac{(1-v)^{n_{mt}}(uv)^{n_{gt}}((1-u)v)^{n_{bt}} f_0(v,u)}{\int_0^1\int_0^1 (1-v)^{n_{mt}}(uv)^{n_{gt}}((1-u)v)^{n_{bt}} f_0(v,u)\,\mathrm{d}u\,\mathrm{d}v} \tag{A1}$$

When the prior distribution takes the 2D Beta form in (1), the corresponding time-t posterior distribution simplifies to (9). In general, the representative agent's belief distribution over terminal states n_T follows (8). With log utility, the inverse price is a martingale $p_t^{-1} = \sum_{n_T} \Pr(n_T|n_t, f_t) p_T^{-1}(n_T)$, which is consistent with (6). □

Proof of Result 3. *i)* Under balanced sentiment $\alpha_{ut} = \beta_{ut}$, the remaining increments $(n_g, n_m, n_b) \sim \text{Multinomial}(T-t;\, vu,\, 1-v,\, v(1-u))$ with $v \sim \text{Beta}(\alpha_{vt}, \beta_{vt})$ and $u \sim \text{Beta}(\alpha_{ut}, \alpha_{ut})$ independent. With $p_T = e^{\sigma n_T}$,

$$p_t^{-1} = M(-\sigma)\, e^{-\sigma n_t}, \qquad \mathbb{E}_t[p_T] = M(\sigma)\, e^{\sigma n_t}, \qquad \text{Var}_t^*[p_T/p_t - 1] = M(\sigma)M(-\sigma) - 1,$$

where $M(s) = \mathbb{E}[e^{s(n_g - n_b)}]$. An increase of δ in news arrival along a balanced path induces the shifts $\alpha_{vt} \mapsto \alpha_{vt} + \delta$, $\beta_{vt} \mapsto \beta_{vt} - \delta$, $\alpha_{ut} \mapsto \alpha_{ut} + \frac{\delta}{2}$, and $\beta_{ut} \mapsto \beta_{ut} + \frac{\delta}{2}$, preserving the symmetry of u. It therefore suffices to establish the following.

Lemma A1. *Parametrize $M_\delta(s)$ by the shift $(\alpha_{vt}+\delta,\ \beta_{vt}-\delta,\ \alpha_{ut}+\delta/2)$ with $\delta \geqslant 0$. If $\alpha_{vt} \leqslant 2\alpha_{ut}$, then $M_\delta(s)$ is strictly increasing in δ for every $s \neq 0$ and $T - t \geqslant 1$.*

Proof. Let $z = 2u-1$, $a = \cosh s - 1 \geqslant 0$, and $b = \sinh s$. The single-draw MGF conditional on (v,u) is $1 - v + vue^s + v(1-u)e^{-s} = 1 + v(a + bz)$, so

$$M_\delta(s) = \sum_{k=0}^{T-t} \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{T-t}{k} \binom{k}{2j} a^{k-2j} b^{2j}\ \mathbb{E}[v_\delta^k]\ \mathbb{E}[z_\delta^{2j}],$$

where we expand the $(T-t)$-th power and use the symmetry $z_\delta \overset{d}{=} -z_\delta$ to eliminate odd moments. All coefficients are nonnegative. By Beta moments, with $(x)_m$ the rising factorial,

$$\mathbb{E}[v_\delta^k] = \frac{(\alpha_{vt}+\delta)_k}{(\alpha_{vt}+\beta_{vt})_k}, \qquad \mathbb{E}[z_\delta^{2j}] = \frac{(\frac{1}{2})_j}{(\alpha_{ut}+\frac{\delta}{2}+\frac{1}{2})_j}.$$

The δ-dependent factor in each summand is therefore $F_{k,j}(\delta) = (\alpha_{vt}+\delta)_k / (\alpha_{ut}+\frac{\delta}{2}+\frac{1}{2})_j$. Differentiating $\log F_{k,j}$,

$$\frac{d}{d\delta} \log F_{k,j} = \sum_{r=0}^{k-1} \frac{1}{\alpha_{vt} + \delta + r} \ - \ \frac{1}{2} \sum_{s=0}^{j-1} \frac{1}{\alpha_{ut} + \frac{\delta}{2} + \frac{1}{2} + s}.$$

Since $j \leqslant k/2$, pair the first $2j$ terms of the first sum against the second sum:

$$\sum_{s=0}^{j-1}\left[\frac{1}{\alpha_{vt}+\delta+2s}+\frac{1}{\alpha_{vt}+\delta+2s+1}-\frac{1}{2(\alpha_{ut}+\frac{\delta}{2}+\frac{1}{2}+s)}\right].$$

The hypothesis $\alpha_{vt} \leqslant 2\alpha_{ut}$ gives $\alpha_{vt}+\delta+2s+1 \leqslant 2(\alpha_{ut}+\frac{\delta}{2}+\frac{1}{2}+s)$, so each subtracted term is at most $1/(\alpha_{vt}+\delta+2s+1)$, leaving a remainder $\geqslant 1/(\alpha_{vt}+\delta+2s) > 0$. The leftover terms $r \geqslant 2j$ in the first sum are positive. When $j = 0$, $(\log F_{k,0})' = \sum_r (\alpha_{vt}+\delta+r)^{-1} > 0$. □

The condition $\alpha_{u0} + \beta_{u0} \geqslant \alpha_{v0}$ gives $\alpha_{vt} \leqslant \alpha_{ut} + \beta_{ut} = 2\alpha_{ut}$, so Lemma A1 applies. Since $M_\delta(\pm\sigma)$ each strictly increase in δ, a higher ν_t strictly raises Var_t^* and lowers $p_t/e^{\sigma n_t}$.

ii) When all agents agree on v, only $u \sim \mathrm{Beta}(\alpha_{ut}, \beta_{ut})$ is heterogeneous. Holding $U_t = \alpha_{ut}/(\alpha_{ut}+\beta_{ut})$ fixed, a higher ν_t raises $\alpha_{ut} + \beta_{ut}$, which strictly concentrates the Beta distribution of u around U_t. This reduces the dispersion of $(n_g - n_b)|v$, so $M(\pm\sigma)$ strictly decreases by convexity of $e^{\pm\sigma x}$, giving lower Var_t^* and higher $p_t/e^{\sigma n_t}$. □

Proof of Result 4. Recall the true probability measure (u_{true}, v_{true}) governs the data generating process of fundamental news. Define the wealth share of agent (v, u) at time T with sample history (ν_T, n_T) as

$$\Omega(u, v; n_T, \nu_T) = \frac{v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}f_0(v,u)}{\int v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}f_0(v,u)\,\mathrm{d}u\,\mathrm{d}v}$$

The asymptotic decay rate (exponentially) of wealth share satisfies

$$\lim_{N\to\infty} -\frac{\log\Omega(v,u;\nu_T,n_T)}{N} = \lim_{N\to\infty} -\frac{1}{N}\log\frac{v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}f_0(v,u)}{\int v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}f_0(v,u)\,\mathrm{d}v\,\mathrm{d}u}$$

Note that in the large N limit, we can write $\nu_T = Nv_{true}$, $n_T = n_{gT} - n_{bT} = N(u_{true}v_{true} - (1-u_{true})v_{true}) = N(2u_{true}v_{true} - v_{true})$. This means we can write

$$\begin{aligned}
&-\frac{1}{N}\log\left(v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}\right)\\
&= -v_{true}\log v - (1-v_{true})\log(1-v) - v_{true}\left(u_{true}\log u + (1-u_{true})\log(1-u)\right),
\end{aligned}$$

and also compute $\lim_{N\to\infty}\frac{1}{N}\log\int v^{\nu_T}(1-v)^{N-\nu_T}u^{\frac{1}{2}n_T+\frac{1}{2}\nu_T}(1-u)^{\frac{1}{2}\nu_T-\frac{1}{2}n_T}f_0(v,u)\,\mathrm{d}v\,\mathrm{d}u = \lim_{N\to\infty}\frac{1}{N}\log\int\exp(N\,h(v_{true},u_{true};v,u)+\log f_0(v,u))\,\mathrm{d}v\,\mathrm{d}u = \sup_{v,u}\{h(v_{true},u_{true};v,u)\}$. The maximum of $h(v_{true}, u_{true}; v, u)$ is achieved at (v_{true}, u_{true}) because the following quantity is always positive when $(v, u) \neq (v_{true}, u_{true})$ (this is *Gibbs's inequality* in information

theory, which guarantees the relative entropy is always positive between two probability measures): $h(v_{true}, u_{true}; v_{true}, u_{true}) - h(v_{true}, u_{true}; v, u) = v_{true} \log \frac{v_{true}}{v} + (1-v_{true}) \log \frac{1-v_{true}}{1-v} + v_{true} \left(u_{true} \log \frac{u_{true}}{u} + (1-u_{true}) \log \frac{1-u_{true}}{1-u} \right) = \mathrm{KL}(v_{true}, v) + v_{true}\, \mathrm{KL}(u_{true}, u) > 0.$

Thus $\lim_{t\to\infty} -\frac{\log \Omega(v,u;\nu_t,n_t)}{t}$ is the sum of $v_{true} \log \frac{v_{true}}{v} + (1-v_{true}) \log \frac{1-v_{true}}{1-v}$ and $v_{true} \left(u_{true} \log \frac{u_{true}}{u} + (1-u_{true}) \log \frac{1-u_{true}}{1-u} \right)$, which is $\mathrm{KL}(v_{true}, v) + v_{true}\mathrm{KL}(u_{true}, u)$. □

Proof of Result 5. Let $\xi = v/\delta$ denote the Poisson arrival rate of news. As $\delta \to 0$, the distribution of ξ at time t converges to a Gamma distribution $\xi \sim \Gamma(\alpha_{vt}\,, \alpha_{v0}\lambda^{-1} + t)$.

For an agent with belief $(u, \xi\delta)$, the arrival rate of "good news" is ξu and bad news $\xi(1-u)$. The counts of "good news" and "bad news" for the interval $[t, T]$ respectively follow Poisson distributions $n_{gT} - n_{gt} \sim \text{Poisson}(\,\xi u(T-t)\,)$, $n_{bT} n_{bt} \sim \text{Poisson}(\,\xi(1-u)(T-t)\,)$.

Using the moment generating functions for $n_{gT} - n_{gt}$, $n_{bT} - n_{bt}$ and ξ, we find that

$$\begin{aligned}
p_t^{-1} &= e^{-\sigma n_t}\, \mathbb{E}[e^{-\sigma(n_T - n_t)} | f_t] = e^{-\sigma n_t} \int e^{-\sigma(n_{gT}-n_{gt}) + \sigma(n_{bT}-n_{bt})} f_t(u,\xi) d\xi du \\
&= e^{-\sigma n_t} \int \exp\left(\xi u(T-t)(e^{-\sigma}-1) + \xi(1-u)(T-t)(e^{\sigma}-1)\right) f_t(u,\xi) d\xi du \\
&= e^{-\sigma n_t} \int \left(1 - \frac{u(T-t)(e^{-\sigma}-1) + (1-u)(T-t)(e^{\sigma}-1)}{\alpha_{v0}\lambda^{-1} + t}\right)^{-\alpha_{vt}} \frac{u^{\alpha_{ut}-1}(1-u)^{\beta_{ut}-1}}{B(\alpha_{ut}, \beta_{ut})} du \\
&= e^{-\sigma n_t} \left(1 - \frac{(T-t)(e^{\sigma}-1)}{\alpha_{v0}\lambda^{-1} + t}\right)^{-\alpha_{vt}} {}_2F_1\left(\alpha_{vt}, \alpha_{ut}; \alpha_{ut} + \beta_{ut}; -\frac{e^{\sigma} - e^{-\sigma}}{(\alpha_{v0}\lambda^{-1} + t)/(T-t) - (e^{\sigma}-1)}\right).
\end{aligned}$$

The final step uses Euler's integral representation for hypergeometric functions. Similarly,

$$\begin{aligned}
\mathbb{E}[p_T] &= e^{\sigma n_t}\, \mathbb{E}[e^{\sigma(n_T - n_t)} | f_t] \\
&= e^{\sigma n_t} \int \left(1 - \frac{u(T-t)(e^{\sigma}-1) + (1-u)(T-t)(e^{-\sigma}-1)}{\alpha_{v0}\lambda^{-1} + t}\right)^{-\alpha_{vt}} \frac{u^{\alpha_{ut}-1}(1-u)^{\beta_{ut}-1}}{B(\alpha_{ut}, \beta_{ut})} du \\
&= e^{\sigma n_t} \left(1 + \frac{(T-t)(1-e^{-\sigma})}{\alpha_{v0}\lambda^{-1} + t}\right)^{-\alpha_{vt}} {}_2F_1\left(\alpha_{vt}, \alpha_{ut}; \alpha_{ut} + \beta_{ut}; \frac{e^{\sigma} - e^{-\sigma}}{(\alpha_{v0}\lambda^{-1} + t)/(T-t) + (1-e^{-\sigma})}\right).
\end{aligned}$$

Substituting the above expression into (7), we obtain (17). □

Proof of Result 6. By construction, from Mr. Market's point of view at time 0, the variables (n_{mT}, n_{gT}, n_{bT}) follow a Dirichlet distribution $B(\beta, \alpha, \gamma) = B(\theta N(1-\lambda), \frac{1}{2}\theta\lambda N + \frac{1}{2}\eta\sqrt{N}, \frac{1}{2}\theta\lambda N - \frac{1}{2}\eta\sqrt{N})$. It is then straightforward to compute the reciprocal of the initial price using Result 1. To obtain a closed-form formula, we use Paul and Plackett (1978).

In the limit as $N \to \infty$, the variable $n_T = n_{gT} - n_{bT}$ follows a normal distribution with the mean $\mathbb{E}\, n_T = N\left(\frac{\alpha-\gamma}{\alpha+\beta+\gamma}\right) = \frac{\eta}{\theta}\sqrt{N}$ and variance $\mathrm{var}[n_T] = N\left(\lambda - \frac{\eta^2}{\theta^2}N^{-1}\right)\frac{1+\theta}{\theta} \to \lambda\frac{1+\theta}{\theta}N$. Thus the random variable $n_T/\sqrt{N}$ follows a normal distribution in the large-N limit: $\frac{n_T}{\sqrt{N}} \sim \mathcal{N}\left(\frac{\eta}{\theta}, \frac{\lambda(1+\theta)}{\theta}\right)$. Consequently, the reciprocal of price is given by $p_0^{-1} = \mathbb{E}\exp\left(-\sigma\sqrt{T}\frac{n_T}{\sqrt{N}}\right)$ and this implies $p_0 = \exp\left(\frac{\eta}{\theta}\sigma\sqrt{T} - \frac{\lambda(1+\theta)}{2\theta}\sigma^2 T\right)$.

For the market price at time t, we need to adjust the wealth distribution and the number of periods. The number of periods left is $N(1-t/T) = N(1-\phi)$, where $\phi = t/T$. The random variables $(n_{mT} - n_{mt}, n_{gT} - n_{gt}, n_{bT} - n_{bt})$ follow a Dirichlet distribution at time t: $B_t(\beta', \alpha', \gamma') = B_t((\theta+\phi)N(1-\lambda), \frac{1}{2}\lambda(\theta+\phi)N + \frac{1}{2}\eta\sqrt{N} + \frac{1}{2}\eta_t\sqrt{\phi N}, \frac{1}{2}\lambda(\theta+\phi)N - \frac{1}{2}\eta\sqrt{N} - \frac{1}{2}\eta_t\sqrt{\phi N})$. The parameter η_t captures the wealth transfer between 0 and t among agents because of heterogeneous beliefs about news content $\eta_t = \frac{n_t}{\sqrt{\phi N}}$. The price at time t can then be written as $p_t^{-1} = \mathbb{E}_t\, p_T^{-1} = e^{-\eta_t\sigma\sqrt{\phi T}}\,\mathbb{E}_t\left(\exp\sigma\sqrt{T/N}(n_t - n_T)\right)$.

The large N limits for the mean and variance of $n_T - n_t$ (which is Gaussian distributed), according to Paul and Plackett (1978), are $\mathbb{E}_t[n_T - n_t] = (1-\phi)N\frac{\alpha'-\gamma'}{\alpha'+\beta'+\gamma'} = \frac{1-\phi}{\theta+\phi}(\eta+\eta_t\sqrt{\phi})\sqrt{N}$ and $\mathrm{Var}_t[n_T - n_t] = (1-\phi)N\left(\lambda - \left(\frac{\eta+\eta_t\sqrt{\phi}}{\theta+\phi}\right)^2 N^{-1}\right)\frac{(1-\phi)N+(\theta+\phi)N}{1+(\theta+\phi)N} \to \lambda\frac{\theta+1}{\theta+\phi}(1-\phi)N$. We can then compute p_t as:

$$p_t = \exp\left(\frac{1-\phi}{\theta+\phi}\eta\sigma\sqrt{T} + \frac{\theta+1}{\theta+\phi}\eta_t\sigma\sqrt{\phi T} - \frac{\lambda(\theta+1)}{2(\theta+\phi)}(1-\phi)\sigma^2 T\right).$$

The risk-neutral variance. Similarly,

$$\mathbb{E}_t\, p_T(n_T) = \lim_{N\to\infty}\mathbb{E}_t\exp\left(\sigma\sqrt{T}\frac{n_T}{\sqrt{N}}\right) = \exp\left(\frac{1-\phi}{\theta+\phi}\eta\sigma\sqrt{T} + \frac{\theta+1}{\theta+\phi}\eta_t\sigma\sqrt{\phi T} + \frac{\lambda(\theta+1)}{2(\theta+\phi)}(1-\phi)\sigma^2 T\right).$$

Thus, we have $\mathcal{V}_{t\to T} = \exp\left(\frac{\lambda(1+\theta)}{\theta+\phi}(1-\phi)\sigma^2 T\right)$. □

Derivation for Equation (13). Assumption (2) is mathematically equivalent to:

$$m_t^* x_{t+1,m} + g_t^* x_{t+1,g} + b_t^* x_{t+1,b} = M_t x_{t+1,m} + G_t x_{t+1,g} + B_t x_{t+1,b}\,. \tag{A2}$$

which can be rewritten as $(g_t^* - G_t)x_{t+1,g} + (m_t^* - M_t)x_{t+1,m} + (b_t^* - B_t)x_{t+1,b} = 0$. Noting that $(m_t^* - M_t) + (g_t^* - G_t) + (b_t^* - B_t) = 0$, this implies that any vector $(x_{t+1,m}, x_{t+1,g}, x_{t+1,b})$ satisfying the zero-delta condition must lie in the span of the risk-free vector $(1,1,1)$, and a second vector that is linearly independent of it, which we may take as $(0, b_t^* - B_t, G_t - g_t^*)$. □

Derivation for Equation (14). The subjective expectation of the price change $p_{t+1} - p_t$, for an agent characterized by belief parameters (u, v) and represented by probabilities (m, g, b), is given by:

$$\mathbb{E}_t^{(v,u)}[p_{t+1} - p_t] = m(p_{t+1,m} - p_t) + g(p_{t+1,g} - p_t) + b(p_{t+1,b} - p_t)\,. \tag{A3}$$

Substituting (3) into (A3) and simplifying:

$$\mathbb{E}_t^{(v,u)}[p_{t+1} - p_t] = \frac{\theta_t^{(v,u)}}{w_t^{(v,u)}}\left[m_t^*(p_{t+1,m} - p_t)^2 + g_t^*(p_{t+1,g} - p_t)^2 + b_t^*(p_{t+1,b} - p_t)^2\right],$$

which equalizes $\frac{\theta_t^{(v,u)}}{w_t^{(v,u)}} \operatorname{Var}_t^*[p_{t+1} - p_t]$ and the equality follows because the risk-neutral covariance between the market price and derivative payoff is zero. This concludes the proof for the market demand $\theta_t^{(v,u)}$. The proof for the derivative demand $\phi_t^{(v,u)}$ follows analogously. □

Proof of Result 7. For an agent with belief (v, u), the expected excess payoff of the derivative is given by:

$$\begin{aligned}
\mathbb{E}_t^{(v,u)}[x_{t+1} - q_t] &= p_t(b_t^* - B_t)uv + p_t(G_t - g_t^*)(1-u)v - p_t(b_t^* G_t - g_t^* B_t) \\
&= \frac{\operatorname{Var}_t^*[p_{t+1}]}{(p_{t+1,g} - p_{t+1,b})}\left(v - v_t^*\right) + \frac{p_t^2(v_t^* - V_t)}{p_{t+1,g} - p_{t+1,b}} r \\
&= \frac{\operatorname{Var}_t^*[x_{t+1} - q_t]}{m_t^* q_t}\left(v - v_t^* + \frac{(v_t^* - V_t)}{\mathbb{E}_t^{(V_t,U_t)}[p_{t+1}/p_t - 1]}\, r\right).
\end{aligned} \tag{A4}$$

Substituting this expression into (14), we obtain (20). □

Proof of Result 8. First, we derive the aggregate belief $H_{n,t}$ for node n. As established in the proofs for the trinomial tree (e.g., Result 2), an agent with log utility and dogmatic beliefs $\mathbf{h} = (h_1, \ldots, h_M)$ optimally acts as a Kelly bettor. Consequently, their wealth share in the economy at time t, denoted as $f_t(\mathbf{h})$, is proportional to their prior belief density multiplied by the likelihood of the observed news path $\mathbf{s}_t$.

Given that the prior distribution is a Dirichlet distribution $f_0(\mathbf{h}) = \frac{1}{B(\beta_1,\ldots,\beta_M)} \prod_{n=1}^M h_n^{\beta_n - 1}$, the wealth-weighted posterior distribution at time t is updated via Bayes' rule: $f_t(\mathbf{h}) \propto \prod_{n=1}^M h_n^{\beta_n - 1} \prod_{n=1}^M h_n^{s_{n,t}} = \prod_{n=1}^M h_n^{\beta_{n,t} - 1}$, where $\beta_{n,t} = \beta_n + s_{n,t}$. Thus, $f_t(\mathbf{h})$ is also a Dirichlet distribution with parameters $(\beta_{1,t}, \ldots, \beta_{M,t})$. The aggregate sentiment $H_{n,t}$ for node n is the wealth-weighted average of the subjective probabilities, which corresponds to the mean

of this updated Dirichlet distribution: $H_{n,t} = \mathbb{E}_{f_t}[h_n] = \frac{\beta_{n,t}}{\sum_{i=1}^{M}\beta_{i,t}} = \frac{s_{n,t}+\beta_n}{t+\sum_{i=1}^{M}\beta_i}$, where we used the identity $\sum_{i=1}^{M} s_{i,t} = t$.

Second, we derive the path-dependent risk-neutral measure $h^*_{t\to T}(\mathbf{s}_T)$. With a representative log-utility investor (whose beliefs correspond to the wealth-weighted distribution f_t), the stochastic discount factor (SDF) from t to T for a specific path ending at $\mathbf{s}_T$ is exactly the inverse of the market return: $M_{t\to T} = \left(\frac{p_T(\mathbf{s}_T)}{p_t}\right)^{-1} = \frac{p_t}{p_T(\mathbf{s}_T)}$. The physical probability of reaching state $\mathbf{s}_T$ from $\mathbf{s}_t$ under the aggregate belief f_t is the expectation of the multinomial likelihood over the Dirichlet posterior. This yields the standard Dirichlet-multinomial predictive distribution: $P(\mathbf{s}_T|\mathbf{s}_t) = \frac{(T-t)!}{\prod_{n=1}^{M}(s_{n,T}-s_{n,t})!}\frac{B(\beta_{1,T},\ldots,\beta_{M,T})}{B(\beta_{1,t},\ldots,\beta_{M,t})}$, where $B(\ldots)$ is the multivariate Beta function.

The risk-neutral probability $h^*_{t\to T}(\mathbf{s}_T)$ equals the physical probability times the SDF, since the gross risk-free rate is normalized to 1. Therefore $h^*_{t\to T}(\mathbf{s}_T) = P(\mathbf{s}_T|\mathbf{s}_t)\frac{p_t}{p_T(\mathbf{s}_T)}$. Substituting the expression for $P(\mathbf{s}_T|\mathbf{s}_t)$ yields the expression of $h^*_{t\to T}(\mathbf{s}_T)$ in Result 8. □

Proof of Result 9. The log payoff, $\log(p_T)$, decomposes into a two jump components, $\omega s_{1,T}$ and $-\omega s_{4,T}$, and a diffusion component, $\sigma\sqrt{T/N}(s_{2,T} - s_{3,T})$. In the continuous-time limit, the distribution of upward jump intensity at time t converges to a Gamma distribution $\Gamma(\beta_{1,t}, (\theta+\phi)N)$, so the wealth-weighted average beliefs of the upward jump count in $[t, T]$ converges to Negative-binomial distribution $NB(\beta_{1,t}, \frac{\theta+\phi}{\theta+1})$. Similarly, the wealth-weighted average beliefs of the downward jump count converges to $NB(\beta_{4,t}, \frac{\theta+\phi}{\theta+1})$. The belief corresponding to the diffusion component is derived in the proofs of Result 6. Combining these results, we obtain the inverse of market price,

$$p_t^{-1} = \mathbb{E}[e^{-\log(p_T)}|f_t] = e^{-\omega(s_{1,t}-s_{4,t})-\sigma\sqrt{T/N}(s_{2,t}-s_{3,t})}\left(1-\frac{(1-\phi)(e^{-\omega}-1)}{\theta+\phi}\right)^{-\beta_{1,t}}$$
$$\left(1-\frac{(1-\phi)(e^{\omega}-1)}{\theta+\phi}\right)^{-\beta_{4,t}}\exp\left[-\frac{1-\phi}{\theta+\phi}\eta\sigma\sqrt{T}-\frac{1-\phi}{\theta+\phi}\eta_t\sigma\sqrt{\phi T}+\frac{1}{2}\frac{\theta+1}{\theta+\phi}(1-\phi)\sigma^2 T\right]. \quad \text{(A5)}$$

Taking the reciprocal of this expression gives (26). The moment-generating function of the log return is $\mathbb{E}^*_t\left[\left(\frac{p_T}{p_t}\right)^{\chi}\right] = p_t^{1-\chi}\,\mathbb{E}_t\left[e^{(\chi-1)\log(p_T)}|f_t\right]$. Substituting into (A5) and using the moment-generating function of the Negative Binomial distribution yields (27). □

References

Andersen, Torben G, Nicola Fusari, and Viktor Todorov. 2015. "The risk premia embedded in index options." *Journal of Financial Economics*, 117(3): 558–584.

Atmaz, Adem, and Andrea M. Buffa. 2023. "Volatility Disagreement and Asset Prices." *Working Paper*.

Atmaz, Adem, and Suleyman Basak. 2018. "Belief Dispersion in the Stock Market." *Journal of Finance*, 73(3): 1225–1279.

Bates, David S. 1991. "The crash of '87: was it expected? The evidence from options markets." *The Journal of Finance*, 46(3): 1009–1044.

Blume, Lawrence, and David Easley. 2006. "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica*, 74(4): 929–966.

Borovička, Jaroslav. 2020. "Survival and Long-Run Dynamics with Heterogeneous Beliefs Under Recursive Preferences." *Journal of Political Economy*, 128(1): 206–251.

Campbell, John Y, and Ludger Hentschel. 1992. "No news is good news: An asymmetric model of changing volatility in stock returns." *Journal of Financial Economics*, 31(3): 281–318.

Cao, H Henry, and Hui Ou-Yang. 2008. "Differences of opinion of public information and speculative trading in stocks and options." *The Review of Financial Studies*, 22(1): 299–335.

Carr, Peter, and Dilip Madan. 1999. "Option valuation using the fast Fourier transform." *Journal of Computational Finance*, 2(4): 61–73.

Cox, John C, and Stephen A Ross. 1976. "The valuation of options for alternative stochastic processes." *Journal of Financial Economics*, 3(1-2): 145–166.

Cox, John C, Stephen A Ross, and Mark Rubinstein. 1979. "Option pricing: A simplified approach." *Journal of Financial Economics*, 7(3): 229–263.

Driessen, Joost, Pascal Maenhout, and Grigory Vilkov. 2009. "The price of correlation risk: Evidence from equity options." *The Journal of Finance*, 64(3): 1377–1406.

Dubinsky, Andrew, Michael Johannes, Andreas Kaeck, and Norman J Seeger. 2019. "Option pricing of earnings announcement risks." *The Review of Financial Studies*, 32(2): 646–687.

Gao, Can, and Ian W. R. Martin. 2021. "Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment." *Journal of Finance*, 76(6): 3211–3254.

Geanakoplos, John. 2010. "The Leverage Cycle." In *NBER Macroeconomic Annual 2009*. Vol. 24, , ed. Daron Acemoglu, Kenneth Rogoff and Michael Woodford, 1–65. Chicago:University of Chicago Press.

Graham, Benjamin. 1949. *The Intelligent Investor.* New York:Harper & Brothers.

Harrison, J. Michael, and David M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Quarterly Journal of Economics*, 92(2): 323–336.

Huang, Wenqian, Albert J. Menkveld, and Shihao Yu. 2021. "Central Counterparty Exposure in Stressed Markets." *Management Science*, 67(6): 3596–3617.

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2022. "Premium for heightened uncertainty: Explaining pre-announcement market returns." *Journal of Financial Economics*, 145(1): 64–87.

Jia, Xiaolan, Xinfeng Ruan, and Jin E Zhang. 2021. "The implied volatility smirk of commodity options." *Journal of Futures Markets*, 41(1): 72–104.

Kelly, Bryan, Ľuboš Pástor, and Pietro Veronesi. 2016. "The price of political uncertainty: Theory and evidence from the option market." *The Journal of Finance*, 71(5): 2417–2480.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield. 2006. "The Price Impact and Survival of Irrational Traders." *Journal of Finance*, 61(1): 195–229.

Koudijs, Peter. 2016. "The Boats That Did Not Sail: Asset Price Volatility in a Natural Experiment." *Journal of Finance*, 71(3): 1185–1226.

Kremens, Lukas, and Ian W. R. Martin. 2019. "The Quanto Theory of Exchange Rates." *American Economic Review*, 109(3): 810–843.

Lakonishok, Josef, Inmoo Lee, Neil D Pearson, and Allen M Poteshman. 2007. "Option market activity." *The Review of Financial Studies*, 20(3): 813–857.

Lucca, David O, and Emanuel Moench. 2015. "The pre-FOMC announcement drift." *The Journal of Finance*, 70(1): 329–371.

Martin, Ian. 2011. "Simple variance swaps." National Bureau of Economic Research.

Martin, Ian, and Dimitris Papadimitriou. 2022. "Sentiment and speculation in a market with heterogeneous beliefs." *American Economic Review*, 112(8): 2465–2517.

Martin, Ian W. R. 2017. "What is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132(1): 367–433.

Martin, Ian W. R., and Christian Wagner. 2019. "What is the Expected Return on a Stock?" *Journal of Finance*, 74(4): 1887–1929.

Martin, Ian WR, and Ran Shi. 2026. "On the Moments of the Stochastic Discount Factor." *Working Paper*.

Mencia, Javier, and Enrique Sentana. 2013. "Valuation of VIX derivatives." *Journal of Financial Economics*, 108(2): 367–391.

Merton, Robert C. 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *The Review of Economics and Statistics*, 247–257.

Merton, Robert C. 1976. "Option pricing when underlying stock returns are discontinuous." *Journal of Financial Economics*, 3(1-2): 125–144.

Ni, Sophie X, Jun Pan, and Allen M Poteshman. 2008. "Volatility information trading in the option market." *The Journal of Finance*, 63(3): 1059–1091.

Pastor, Lubos, and Pietro Veronesi. 2012. "Uncertainty about Government Policy and Stock Prices." *The Journal of Finance*, 67(4): 1219–1264.

Paul, S. R., and R. L. Plackett. 1978. "Inference Sensitivity for Poisson Mixtures." *Biometrika*, 65(3): 591–602.

Rubinstein, Mark. 1994. "Implied binomial trees." *The Journal of Finance*, 49(3): 771–818.

Sandroni, Alvaro. 2000. "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica*, 68(6): 1303–1341.

Scheinkman, José A., and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." *Journal of Political Economy*, 111(6): 1183–1219.

Simsek, Alp. 2013. "Belief Disagreements and Collateral Constraints." *Econometrica*, 81(1): 1–53.

Internet Appendix to

Pricing News and No News with Heterogeneous Beliefs

Can Gao Brandon Yueyang Han[1]

[not for publication]

Abstract

This Internet Appendix presents supplementary material and results not included in the main body of the paper.

[1]Can Gao is at University of St. Gallen and Swiss Finance Institute, webpage, can.gao@unisg.ch; Brandon Yueyang Han is at University of Maryland, webpage, yhan1@umd.edu.

I.A.1 Disagreements among macro forecasters

We document that professional forecasters disagree not only about the direction of macroeconomic fundamentals, but also—and to a comparable degree—about their perceived volatility. We use individual-level probability histogram forecasts from the Philadelphia Fed's Survey of Professional Forecasters (SPF), which elicits each panelist's full predictive distribution for next-year real GDP growth. From each individual histogram, we extract the implied mean μ_i and standard deviation σ_i, then compute cross-sectional statistics each quarter.

Data construction. The SPF worksheet PRGDP reports each panelist's probability assignment (summing to 100) across a set of growth-rate bins for annual-average over annual-average real GDP growth. Each survey elicits histograms for the current year and the following year; we use the *next-year* forecast throughout. Following Table 7 of the SPF Documentation (March 2024), the bin structure for the next-year forecast has changed as follows: 1) 1992:Q1–2009:Q1 (variables PRGDP11–20): 10 bins of approximately 1 pp width, ranging from "≥ 6" (PRGDP11) to "< -2" (PRGDP20); 2) 2009:Q2–2020:Q1 (variables PRGDP12–22): 11 bins of approximately 1 pp width. The additional bin extends the left tail from "< -2" to "-3 to -2.1" and "< -3."; 3) 2020:Q2–2024:Q1 (variables PRGDP12–22): 11 bins with substantially widened ranges, reflecting the unprecedented uncertainty of the COVID-19 pandemic and its aftermath (e.g., the top bin is "≥ 16"); 4) 2024:Q2–present (variables PRGDP12–22): 11 bins with intermediate ranges (top bin "≥ 9").

For each bin we assign a midpoint equal to the centre of interior bins and a representative value for the open-ended tails (e.g., 7.0 for "≥ 6" and -3.0 for "< -2" in the 10-bin regime; -4.0 for "< -3" in the 11-bin regime). We drop individual forecasts whose bin probabilities do not sum to approximately 100 (outside the 95–105 range). The implied mean and standard deviation of forecaster i's distribution are $\mu_i = \sum_k p_k m_k$ and $\sigma_i = [\sum_k p_k (m_k - \mu_i)^2]^{1/2}$, where p_k and m_k are the normalised probability and midpoint of bin k. The resulting sample covers 137 quarters from 1992:Q1 to 2026:Q1, with 4,707

Table I.A.1: Cross-sectional disagreement in the Survey of Professional Forecasters.

	1992:Q1–2026:Q1
$\mathrm{SD}(\mu_i)$ (pp)	0.60
$\mathrm{SD}(\sigma_i)$ (pp)	0.45
$\overline{\mu_i}$ (pp)	2.45
$\overline{\sigma_i}$ (pp)	1.09
$\mathrm{CV}(\mu) \equiv \overline{\mathrm{SD}(\mu_i)} \,/\, \overline{\mu_i}$	0.24
$\mathrm{CV}(\sigma) \equiv \overline{\mathrm{SD}(\sigma_i)} \,/\, \overline{\sigma_i}$	0.41
$\mathrm{CV}(\sigma) \,/\, \mathrm{CV}(\mu)$	1.70
$\rho(\mu_i, \sigma_i)$	-0.10
Mean N per quarter	34.4
Quarters	137

individual forecasts and an average of 34 panelists per quarter.

Summary statistics. Table I.A.1 reports the key cross-sectional statistics, averaged across all 137 quarters. The average cross-sectional standard deviation of individual means is $\mathrm{SD}(\mu_i) = 0.60$ percentage points, while the average cross-sectional standard deviation of individual perceived volatilities is $\mathrm{SD}(\sigma_i) = 0.45$ pp. Because these quantities have different scales—the average μ_i is 2.45 pp while the average σ_i is 1.09 pp—we compare coefficients of variation, defined as the ratio of time-series-averaged dispersion to time-series-averaged level:

$$\mathrm{CV}(\mu) \equiv \frac{\overline{\mathrm{SD}(\mu_i)}}{\overline{\mu_i}} = \frac{0.60}{2.45} = 0.24\,, \qquad \mathrm{CV}(\sigma) \equiv \frac{\overline{\mathrm{SD}(\sigma_i)}}{\overline{\sigma_i}} = \frac{0.45}{1.09} = 0.41\,.$$

Relative to their respective levels, disagreement about the volatility of GDP growth *exceeds* disagreement about its expected level by a factor of $0.41/0.24 = 1.7$. The average within-quarter correlation between μ_i and σ_i is $\bar{\rho} = -0.10$. Knowing a forecaster's directional view provides almost no information about their perceived volatility.

Cross-sectional belief distributions examples. Figures I.A.1 and I.A.2 display scatter plots of individual forecasters' (σ_i, μ_i) pairs for representative quarters spanning the 2008 financial crisis and the COVID-19 pandemic, respectively. During the GFC (Figure I.A.1), the belief cloud shifts downward between the pre-crisis and crisis quarters as

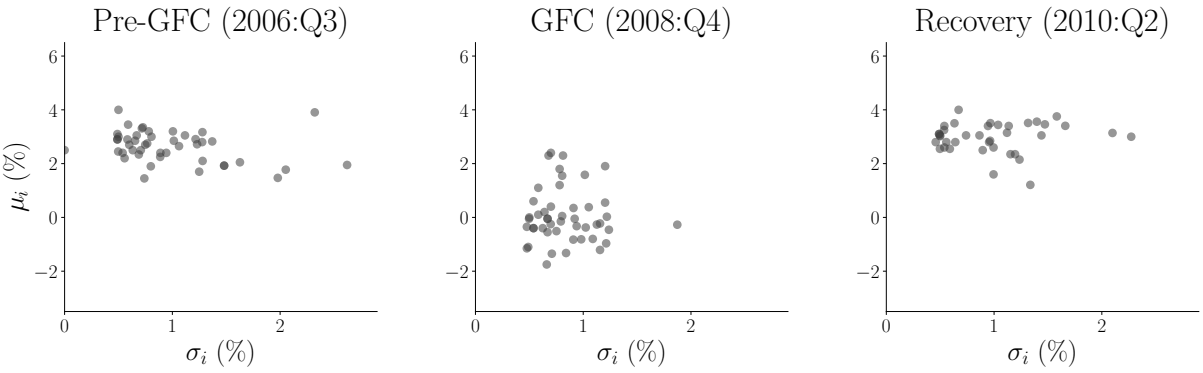


Figure I.A.1: **Individual forecasters' beliefs around the Global Financial Crisis.** Each dot represents one SPF panelist's next-year GDP growth forecast: σ_i (perceived volatility, x-axis) versus μ_i (expected growth, y-axis). Moments are extracted from individual histogram forecasts using bin midpoints. Axis scales are identical across panels.

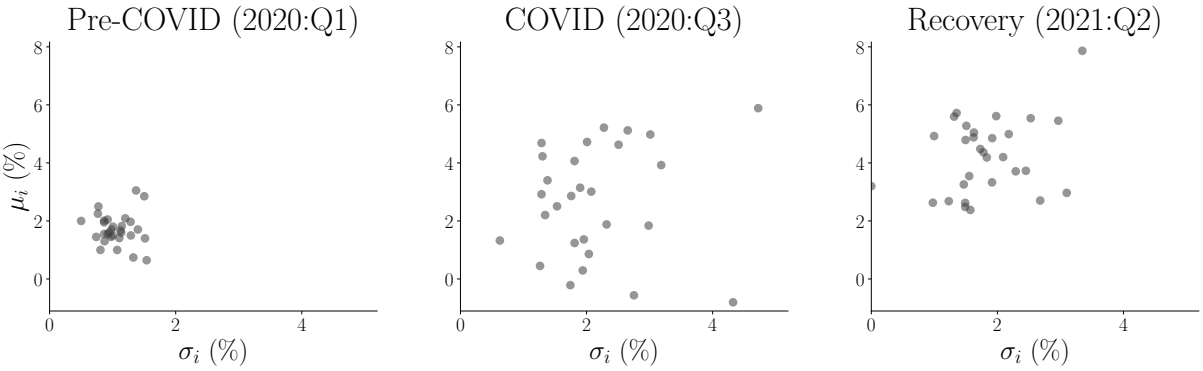


Figure I.A.2: **Individual forecasters' beliefs around the COVID-19 pandemic.** Same construction as Figure I.A.1. Axis scales are identical across panels. Note the different axis scales compared to Figure I.A.1, reflecting the wider histogram bins in the post-2020 period.

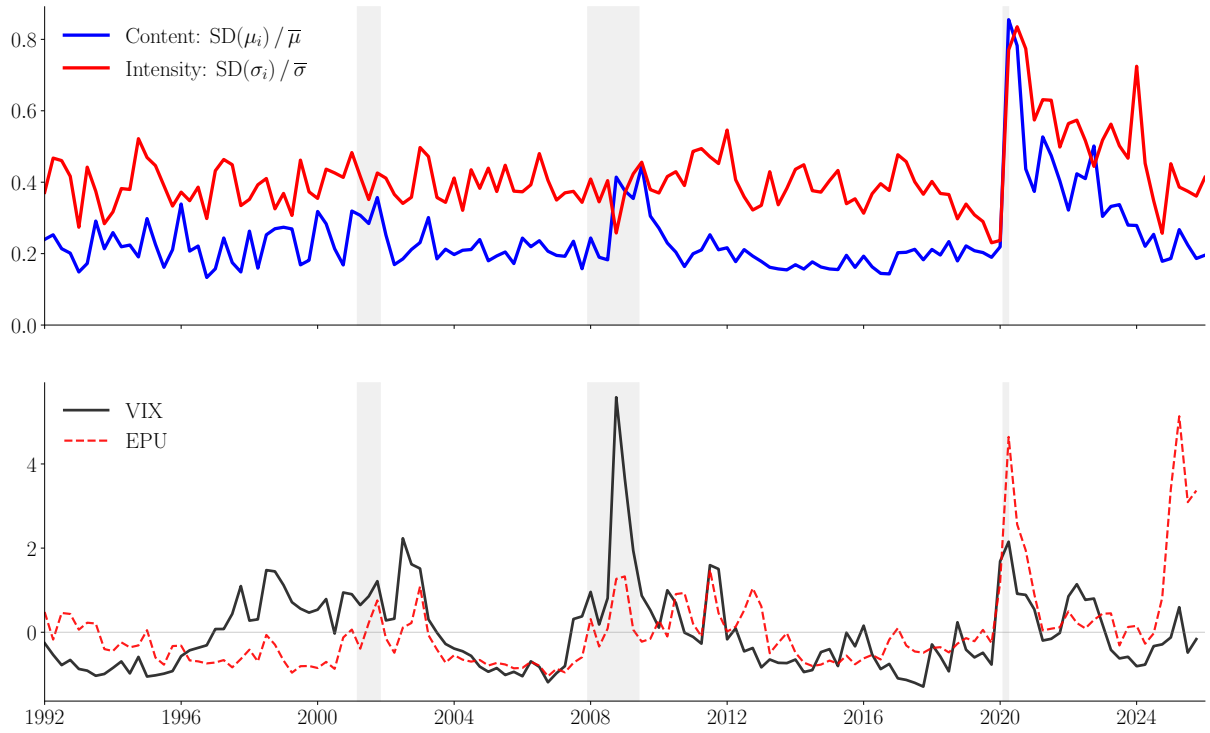


Figure I.A.3: **Cross-sectional disagreement and uncertainty indices.** Top panel: quarterly cross-sectional standard deviation of individual forecast means (blue) and perceived volatilities (red), each normalised by the full-sample average level of the corresponding moment ($\overline{\mu} = 2.45$ pp and $\overline{\sigma} = 1.09$ pp). Bottom panel: standardised quarterly averages of the VIX (solid black) and the Baker–Bloom–Davis Economic Policy Uncertainty index (dashed red). Grey bars indicate NBER recession dates.

forecasters revise expected growth, while the horizontal spread in σ_i persists throughout—reflecting sustained disagreement about the uncertainty of the outlook even as directional views converge during the recovery. During COVID (Figure I.A.2), the cloud expands markedly in both dimensions at the pandemic onset, then compresses during the recovery, with the σ_i dimension contracting more slowly than the μ_i dimension.

Time-series dynamics. Figure I.A.3 plots the quarterly time series of normalised disagreement—SD(μ_i)/$\overline{\mu}$ and SD(σ_i)/$\overline{\sigma}$, where $\overline{\mu}$ and $\overline{\sigma}$ are the full-sample average levels (top panel)—alongside standardised VIX and Economic Policy Uncertainty indices (bottom panel). Grey bars indicate NBER recession dates. Content disagreement spikes during the 2001 and 2008 recessions and the COVID-19 pandemic, tracking the VIX.

Intensity disagreement follows distinct dynamics: it remains elevated relative to content disagreement throughout the sample, including during calm periods.

I.A.2 News arrival with Dirichlet disagreements

Recall that the three states—no news m, good news g, and bad news b—are given by $m = 1-v$, $g = vu$, and $b = v(1-u)$, where (v, u) follows a two-dimensional Beta distribution. Under the special parameter restriction $\alpha_{v0} = \alpha_{u0}+\beta_{u0}$, the induced joint distribution over (m, g, b) collapses to a Dirichlet distribution with parameters $(\beta_{v0}, \alpha_{u0}, \beta_{u0})$. In this case, the increase in intensity sentiment always dominates the content-concentration channel. The result extends beyond paths with balanced sentiment in news content, $U_t = 1/2$, and holds for all possible realizations of news. Result IA1 establishes this general result.

Result IA1. *Suppose* $p_T(n_T) = e^{\sigma n_T}$ *and* $\alpha_{vt} = \alpha_{ut} + \beta_{ut}$. *Compare two realized histories at the same calendar time* t *with the same net news* n_t *but different news counts* $\nu_t > \nu'_t$. *The history with higher* ν_t *has strictly higher risk-neutral variance* $\mathrm{Var}^*_t[p_T/p_t - 1]$ *and strictly lower normalized price* $p_t/e^{\sigma n_t}$.

Proof. Under the Dirichlet specification, the remaining counts (n_m, n_g, n_b) with $n_m+n_g+n_b = T-t$ follow a Dirichlet-Multinomial with cell parameters $(\beta_{vt}, \alpha_{ut}, \beta_{ut})$ and precision $S = \beta_{vt} + \alpha_{ut} + \beta_{ut}$. The MGF of $n_g - n_b$ is

$$M_n(s) = \sum_{n_m+n_g+n_b=n} \frac{n!}{n_m!\, n_g!\, n_b!} \frac{(\beta_{vt})_{n_m}\,(\alpha_{ut})_{n_g}\,(\beta_{ut})_{n_b}}{(S)_n}\, e^{s(n_g-n_b)}\,, \tag{I.A.1}$$

where $n = T - t$ and $(a)_k = a(a+1)\cdots(a+k-1)$. Multiplying by $\frac{(S)_n}{n!}\, z^n$ and summing over $n \geqslant 0$,

$$\mathcal{G}(s, z) \;\equiv\; \sum_{n=0}^{\infty} \frac{(S)_n}{n!}\, z^n\, M_n(s) \;=\; (1-z)^{-\beta_{vt}}\,(1-ze^{s})^{-\alpha_{ut}}\,(1-ze^{-s})^{-\beta_{ut}}\,. \tag{I.A.2}$$

Set $\delta = \nu_t - \nu'_t > 0$. The parameter shift $\beta_{vt} \to \beta_{vt} - \delta$, $\alpha_{ut} \to \alpha_{ut} + \delta/2$, $\beta_{ut} \to \beta_{ut} + \delta/2$ leaves S invariant. Differentiating:

$$\frac{\partial \mathcal{G}}{\partial \delta} = \mathcal{G}(s, z; \delta)\, \underbrace{\left[\log(1-z) - \tfrac{1}{2}\log(1-ze^{s}) - \tfrac{1}{2}\log(1-ze^{-s})\right]}_{B(s,z)}\,. \tag{I.A.3}$$

Expanding via $-\log(1-x) = \sum_{k\geqslant 1} x^k/k$,

$$B(s,z) = \sum_{k=1}^{\infty} \frac{\cosh(ks)-1}{k}\, z^k\,. \tag{I.A.4}$$

Every coefficient is strictly positive for $s \neq 0$. Since $\mathcal{G}$ also has all positive z^n-coefficients, $\frac{\partial \mathcal{G}}{\partial \delta} = \mathcal{G}\cdot B$ has strictly positive z^n-coefficients for $n \geqslant 1$ by the Cauchy product. Dividing by $\frac{(S)_n}{n!} > 0$,

$$\frac{\partial M_n(s)}{\partial \delta} > 0 \qquad \text{for all } n \geq 1,\ s \neq 0\,. \tag{I.A.5}$$

With $p_T = e^{\sigma n_T}$: $p_t^{-1} = e^{-\sigma n_t} M_{T-t}(-\sigma)$, $\mathbb{E}_t[p_T] = e^{\sigma n_t} M_{T-t}(\sigma)$, $\mathrm{Var}_t^*[p_T/p_t - 1] = M_{T-t}(\sigma)M_{T-t}(-\sigma) - 1$. Applying (I.A.5) at $s = \pm\sigma$: higher ν_t strictly raises $M_{T-t}(\pm\sigma)$, hence strictly lowers $p_t/e^{\sigma n_t} = 1/M_{T-t}(-\sigma)$ and strictly raises Var_t^*. □

I.A.3 Wealth distribution and derivative exposure

Result IA2 (Wealth Distribution and News Intensity Beliefs Across Excess Return Expectations)**.** *Define the variables $z(r)$, k, and γ:*

$$z(r) = \frac{(1+r)g_t^* m_t^* - g_t^* M_t}{G_t m_t^* - g_t^* M_t}\,, \quad k = \frac{G_t m_t^* - g_t^* M_t}{g_t^* M_t - \frac{B_t}{b_t^*} g_t^* m_t^*}\,, \quad \gamma = \frac{g_t^* m_t^*}{G_t m_t^* - g_t^* M_t}\,. \tag{I.A.6}$$

i) The wealth share of agents who believe the underlying's expected return is r satisfies

$$f_t(r) = \frac{\gamma k^{\beta_{ut}} z^{\alpha_{vt}-1}(1-z)^{\beta_{ut}+\beta_{vt}-1} B(\beta_{ut},\beta_{vt})}{(k+1)^{\beta_{ut}} B(\alpha_{vt},\beta_{vt})B(\alpha_{ut},\beta_{ut})} F_1\left(\beta_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\alpha_{ut}; \beta_{ut}+\beta_{vt}; 1-z, \frac{k(1-z)}{1+k}\right),$$

for $z > 0$, and

$$f_t(r) = \frac{\gamma k(-kz)^{\alpha_{vt}-1}(1+kz)^{\alpha_{ut}} B(\alpha_{ut},\beta_{vt})}{(1+k)^{\alpha_{ut}+\beta_{vt}-1} B(\alpha_{vt},\beta_{vt})B(\alpha_{ut},\beta_{ut})} F_1\left(\alpha_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}; 1+kz, \frac{1+kz}{1+k}\right),$$

for $z < 0$. $F_1()$ represents Appell hypergeometric function, and $(\alpha_{vt}, \beta_{vt}, \alpha_{ut}, \beta_{ut})$ follows (10).

ii) The wealth-weighted average news arrival rate v among investors with expectation r, which we call $\mathbb{E}[v|r]$, satisfies, when $z \geqslant 0$,

$$\mathbb{E}[v|r] = z + z(1-z)\cdot \frac{\beta_{ut}}{\beta_{ut}+\beta_{vt}} \cdot \frac{F_1\left(\beta_{ut}+1; \alpha_{vt}+\beta_{vt}, 1-\alpha_{ut}; \beta_{ut}+\beta_{vt}+1; 1-z, \frac{k(1-z)}{1+k}\right)}{F_1\left(\beta_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\alpha_{ut}; \beta_{ut}+\beta_{vt}; 1-z, \frac{k(1-z)}{1+k}\right)},$$

and when $z < 0$,

$$\mathbb{E}[v|r] = -kz + (-kz)(1+kz) \cdot \frac{\alpha_{ut}}{\alpha_{ut}+\beta_{vt}} \frac{F_1\left(\alpha_{ut}+1; \alpha_{vt}+\beta_{vt}, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}+1; 1+kz, \frac{1+kz}{1+k}\right)}{F_1\left(\alpha_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}; 1+kz, \frac{1+kz}{1+k}\right)}$$

The variables $z(r)$ *and* k *follow the same definitions as in* (I.A.6)*, and* $F_1(\dots)$ *represents the Appell hypergeometric function.*

Proof for IA2 . Consider the group of agents whose subjective expectation of the market return equals r: $\mathbb{E}_t^{(v,u)}[p_{t+1}/p_t - 1] = r$. We express u as a function of v and r:

$$u = \frac{p_{t+1,m} - p_{t+1,b}}{p_{t+1,g} - p_{t+1,b}} + \frac{(1+r)p_t - p_{t+1,m}}{(p_{t+1,g} - p_{t+1,b})v} \tag{I.A.7}$$

When $r < p_{t+1,m}/p_t - 1$, u stays in $[0,1]$ if and only if $v \in [(p_{t+1,m} - (1+r)p_t)/(p_{t+1,m} - p_{t+1,b}), 1]$. When $r > p_{t+1,m}/p_t - 1$, u stays in $[0,1]$ if and only if $v \in [((1+r)p_t - p_{t+1,m})/(p_{t+1,g} - p_{t+1,m}), 1]$. Substituting (I.A.7) into (1), the conditional distribution of v given r satisfies $f_t(v|r) \propto (v+kz)^{\alpha_{ut}-1}(v-z)^{\beta_{ut}-1} v^{\alpha_{vt}-\alpha_{ut}-\beta_{ut}}(1-v)^{\beta_{vt}-1}$, where z as a function of r represents the scaled return and k is a constant:

$$z = \frac{(1+r)p_t - p_{t+1,m}}{p_{t+1,g} - p_{t+1,m}} = \frac{1+r-M_t/m_t^*}{G_t/g_t^* - M_t/m_t^*}, \quad k = \frac{p_{t+1,g} - p_{t+1,m}}{p_{t+1,m} - p_{t+1,b}} = \frac{G_t/g_t^* - M_t/m_t^*}{M_t/m_t^* - B_t/b_t^*}.$$

Integrating the joint distribution $f_t(v,z)$ over v yields the marginal distribution of z:

$$f_t(z) = \left(\frac{k}{k+1}\right)^{\beta_{ut}} \frac{B(\beta_{ut},\beta_{vt})}{B(\alpha_{ut},\beta_{ut})B(\alpha_{vt},\beta_{vt})} z^{\alpha_{vt}-1}(1-z)^{\beta_{ut}+\beta_{vt}-1}$$
$$F_1\left(\beta_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\alpha_{ut}, \beta_{ut}+\beta_{vt}; 1-z, \frac{k(1-z)}{1+k}\right)$$

for $z > 0$ and

$$f_t(z) = \frac{k}{(k+1)^{\alpha_{ut}}} \frac{B(\alpha_{ut},\beta_{vt})}{B(\alpha_{ut},\beta_{ut})B(\alpha_{vt},\beta_{vt})} (-kz)^{\alpha_{vt}-1}(1+kz)^{\alpha_{ut}+\beta_{vt}-1}$$
$$F_1\left(\alpha_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}; 1+kz, \frac{1+kz}{1+k}\right)$$

for $z < 0$. Substituting into $f_t(r) = p_t/(p_{t+1,g} - p_{t+1,m}) f_t(z)$, we obtain the two equations in Result IA2.

From the distribution $f_t(v|r)$, the conditional expectation $\mathbb{E}[v|r]$ can be expressed in terms of ratios of Appell hypergeometric functions. This expression allows us to compute the wealth-weighted average derivative position for this group of agents. Specifically, when

$z \geqslant 0$, we obtain:

$$\mathbb{E}[v|r] = z + z(1-z) \cdot \frac{\beta_{ut}}{\beta_{ut}+\beta_{vt}} \cdot \frac{F_1\left(\beta_{ut}+1; \alpha_{vt}+\beta_{vt}, 1-\alpha_{ut}; \beta_{ut}+\beta_{vt}+1; 1-z, \frac{k(1-z)}{1+k}\right)}{F_1\left(\beta_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\alpha_{ut}; \beta_{ut}+\beta_{vt}; 1-z, \frac{k(1-z)}{1+k}\right)}.$$

When $z < 0$, we obtain:

$$\mathbb{E}[v|r] = -kz + (-kz)(1+kz) \cdot \frac{\alpha_{ut}}{\alpha_{ut}+\beta_{vt}} \frac{F_1\left(\alpha_{ut}+1; \alpha_{vt}+\beta_{vt}, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}+1; 1+kz, \frac{1+kz}{1+k}\right)}{F_1\left(\alpha_{ut}; \alpha_{vt}+\beta_{vt}-1, 1-\beta_{ut}; \alpha_{ut}+\beta_{vt}; 1+kz, \frac{1+kz}{1+k}\right)}.$$

□

I.A.4 Bayesian Individuals: Numerical Illustration

The main text introduces the Bayesian Individuals extension (Section 2.6), in which agent (v, u) holds prior beliefs $\tilde{v} \sim \text{Beta}(\zeta v, \zeta(1-v))$ and $\tilde{u} \sim \text{Beta}(\zeta u, \zeta(1-u))$ over the true arrival intensity and content parameters, with concentration $\zeta > 0$ governing prior precision. After observing the news history up to node (ν_t, n_t, t)—with $n_{gt} = (\nu_t + n_t)/2$ good news and $n_{bt} = (\nu_t - n_t)/2$ bad news events—agent (v, u)'s posterior means are

$$\hat{v} = \frac{\zeta\, v + \nu_t}{\zeta + t}, \qquad \hat{u} = \frac{\zeta\, u + n_{gt}}{\zeta + \nu_t}. \tag{I.A.8}$$

In the dogmatic limit $\zeta \to \infty$, we recover $\hat{v} = v$ and $\hat{u} = u$; the effective beliefs coincide with the agent indices and the model reduces to the baseline of Sections 2–3. For finite ζ, each agent's effective beliefs are pulled toward the empirical frequencies. Because the agent index (v, u) ranges over $(0, 1)^2$, the effective beliefs at node (ν_t, n_t, t) are confined to the rectangle

$$\hat{v} \in \left(\frac{\nu_t}{\zeta + t}, \frac{\zeta + \nu_t}{\zeta + t}\right), \qquad \hat{u} \in \left(\frac{n_{gt}}{\zeta + \nu_t}, \frac{\zeta + n_{gt}}{\zeta + \nu_t}\right). \tag{I.A.9}$$

As ζ decreases, this rectangle shrinks: individual learning compresses the support of effective beliefs, reducing the population of agents available to take extreme positions. In the limit $\zeta \to 0^+$, all agents' posteriors converge to the empirical frequencies $(\nu_t/t,\ n_{gt}/\nu_t)$, eliminating effective disagreement entirely.

Wealth distribution with Bayesian agents. Under log utility and complete markets, each agent's one-period wealth growth equals the ratio of their subjective probability to the risk-neutral probability for the realized state (equation (3) in the main text). For

a Bayesian agent indexed by (v, u), the subjective probability of each state at time τ is determined by their *current* posterior—which depends on the history observed up to τ. However, because the Beta prior is conjugate to the Bernoulli likelihood, the product of one-step-ahead predictive probabilities along any sample path telescopes into a ratio of Beta functions.

To see this, consider the intensity dimension. Agent (v, u) holds prior $\tilde{v} \sim \text{Beta}(\zeta v, \zeta(1-v))$. At time τ, having observed ν_τ news events and $\tau - \nu_\tau$ no-news periods, the agent's predictive probability that news arrives at $\tau + 1$ is $(\zeta v + \nu_\tau)/(\zeta + \tau)$. If news does arrive, the predictive probability was $(\zeta v + \nu_\tau)/(\zeta + \tau)$; if not, it was $(\zeta(1-v) + \tau - \nu_\tau)/(\zeta + \tau)$. The product of these one-step-ahead probabilities over $\tau = 0, \ldots, t-1$ yields the marginal likelihood of the observed sequence under the agent's prior. By standard properties of exchangeable sequences (the Beta-Binomial model), this product depends only on the sufficient statistics $(\nu_t, t - \nu_t)$ and equals

$$\prod_{\tau=0}^{t-1} \Pr_{\tau}^{(v,u)}(\text{news or no-news at } \tau) \;=\; \frac{B(\zeta v + \nu_t,\ \zeta(1-v) + t - \nu_t)}{B(\zeta v,\ \zeta(1-v))}\,. \tag{I.A.10}$$

The same argument applies to the content dimension, conditional on news arriving. Agent (v, u)'s prior over content is $\tilde{u} \sim \text{Beta}(\zeta u, \zeta(1-u))$. Among the ν_t news events, n_{gt} are good and n_{bt} are bad, and the marginal likelihood is

$$\prod_{\tau:\text{news at } \tau} \Pr_{\tau}^{(v,u)}(\text{good or bad at } \tau) \;=\; \frac{B(\zeta u + n_{gt},\ \zeta(1-u) + n_{bt})}{B(\zeta u,\ \zeta(1-u))}\,. \tag{I.A.11}$$

Because the risk-neutral probabilities (v^*_τ, u^*_τ) at each step do not depend on the agent index (v, u), they cancel when comparing wealth ratios across agents. The time-t wealth distribution therefore satisfies

$$f_t(v,u) \;\propto\; f_0(v,u) \cdot \frac{B(\zeta v + \nu_t,\ \zeta(1-v) + t - \nu_t)}{B(\zeta v,\ \zeta(1-v))} \cdot \frac{B(\zeta u + n_{gt},\ \zeta(1-u) + n_{bt})}{B(\zeta u,\ \zeta(1-u))}\,, \tag{I.A.12}$$

where $f_0(v, u)$ is the initial wealth (population) distribution from (1). This is the key computational object: it gives the wealth weight of agent (v, u) as a function of the agent index, the prior precision ζ, and the realized news counts (ν_t, n_{gt}, n_{bt}). In the dogmatic limit $\zeta \to \infty$, each Beta ratio converges to a power function—$B(\zeta v + \nu_t, \zeta(1-v) + t -$

$\nu_t)/B(\zeta v, \zeta(1-v)) \to v^{\nu_t}(1-v)^{t-\nu_t}$ up to constants—recovering (9) in the main text.

Equilibrium computation. Given the wealth distribution (I.A.12), equilibrium prices are computed by backward induction on the trinomial tree, as in the dogmatic case. The wealth-weighted aggregate beliefs are

$$V_t = \int_0^1 \int_0^1 \hat{v}(v,u)\, f_t(v,u)\, \mathrm{d}v\, \mathrm{d}u\,, \qquad U_t = \frac{\int_0^1 \int_0^1 \hat{v}(v,u)\, \hat{u}(v,u)\, f_t(v,u)\, \mathrm{d}v\, \mathrm{d}u}{V_t}\,, \quad \text{(I.A.13)}$$

where $\hat{v}(v,u)$ and $\hat{u}(v,u)$ are the effective beliefs from (I.A.8) and f_t is normalized to integrate to one. The risk-neutral measure (v_t^*, u_t^*) follow the same expressions as (4), with each agent's effective beliefs substituted for the dogmatic beliefs.

Two forces now shape the evolution of the wealth distribution. First, the *wealth channel* (market selection) operates exactly as in the dogmatic case: agents whose beliefs are vindicated by realized news gain wealth and pricing influence. Second, *individual learning* shifts each agent's posterior toward the data. These two forces compound rather than offset: after a quiet period, the wealth channel transfers wealth toward low-v agents, while individual learning simultaneously pulls every agent's $\hat{v}$ downward. The dogmatic case—where only the wealth channel operates—is therefore the conservative benchmark.

Prices and sentiment. Figure I.A.4 reports prices and Mr. Market's arrival-intensity belief V_t along three sample paths from Figure 2—the no-news path (red, $\nu_2 = 0$), the good-news-then-no-news path (black, $\nu_2 = 1$, $n_2 = 1$), and the volatile path (blue, $\nu_2 = 2$, $n_2 = 0$)—for $\zeta \in \{\infty, 20, 10, 2\}$.

The fan of lines across ζ values reveals that individual learning amplifies the pricing effects of heterogeneous beliefs. Along the no-news path (red), lower ζ depresses V_t further—agents individually update toward $\hat{v} = 0$, compounding the wealth transfer toward calm-world believers—and prices rise more sharply. The "no news is good news" effect documented in Result 3 is thus stronger with Bayesian agents. Along the volatile path (blue), the effect reverses: lower ζ elevates V_t further, as agents update toward $\hat{v} = 1$, and prices are depressed. The spread in V_t between the quiet and volatile paths widens monotonically as ζ falls.

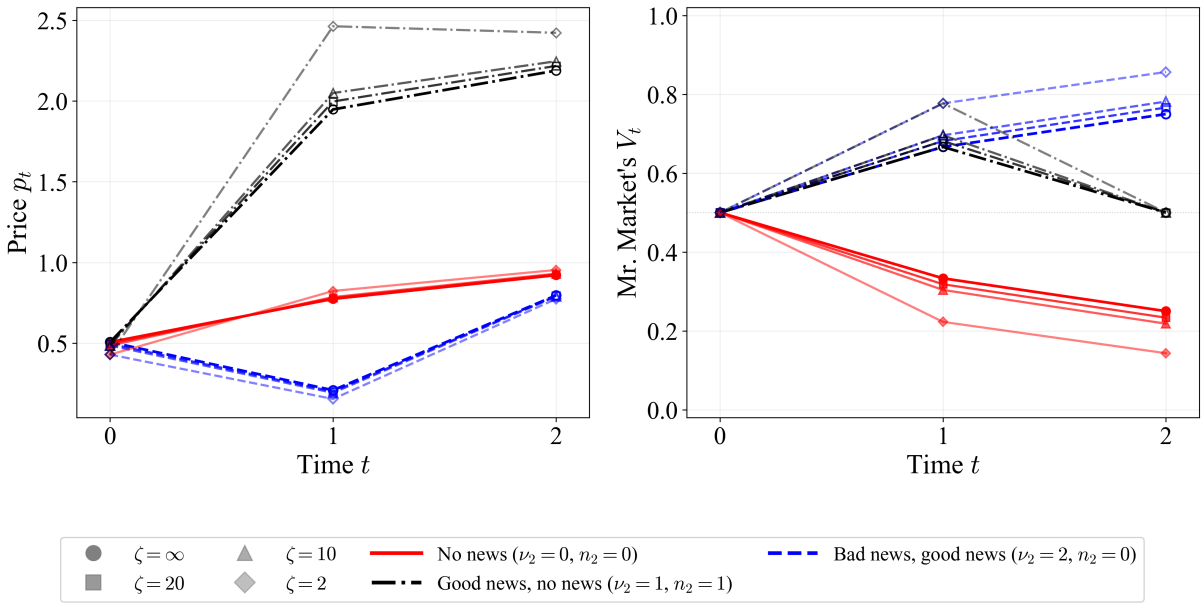


Figure I.A.4: **Bayesian Individuals: prices and arrival-intensity beliefs along three sample paths.** The left panel reports the market price p_t and the right panel reports Mr. Market's arrival-intensity belief V_t for $\zeta \in \{\infty, 20, 10, 2\}$. Red solid lines trace the no-news path ($\nu_2 = 0$, $n_2 = 0$), black dash-dot lines trace good news followed by no news ($\nu_2 = 1$, $n_2 = 1$), and blue dashed lines trace the volatile path ($\nu_2 = 2$, $n_2 = 0$). Marker shapes distinguish ζ values. The terminal payoff is $p_T(n_T) = e^{\sigma n_T}$ with $T = 3$ and $\sigma = \log\sqrt{5}$; the initial wealth distribution is uniform in both v and u.

Portfolio boundaries. Figures I.A.5–I.A.6 extend Figure 3 to the Bayesian setting, plotting portfolio allocation boundaries on the *effective belief* domain $(\hat{v}, \hat{u})$ for $\zeta \in \{\infty, 20, 10, 2\}$. Panel (a) in each figure reproduces the dogmatic case from Figure 3 exactly, since $\hat{v} = v$ and $\hat{u} = u$ when $\zeta = \infty$. As ζ decreases, the coloured contours progressively compress into the effective belief rectangle (I.A.9), shown as a light shaded region. The gray reference lines (corresponding to the $t = 0$ boundaries) always span the full $[0, 1]^2$ domain, since no data has arrived at $t = 0$.

Figure I.A.5 corresponds to the no-news path ($\nu_2 = 0$, $n_2 = 0$). The effective belief domain compresses from the right: $\hat{v}_{\max} = \zeta/(\zeta+2)$ shrinks from 1 to 0.5 as ζ falls from ∞ to 2. Agents who would have believed in high arrival intensity have all been pulled toward "calm world" by observing no news. The risk-neutral measure (marked $\times$) remains inside the effective domain for all ζ values shown.

Figure **??** corresponds to the black path ($\nu_2 = 1$, $n_2 = 1$). The domain compresses

symmetrically in $\hat{v}$ (one news event is the neutral case for intensity learning), but the $\hat{u}$ floor rises because all agents update toward the 100% good-news rate among arrived content. Mr. Market stays at $V_t = 0.5$ across all ζ but U_t rises from 0.667 to 0.777 as individual learning amplifies the content signal.

Figure I.A.6 corresponds to the volatile path ($\nu_2 = 2$, $n_2 = 0$). This case yields the most dramatic compression, with $\hat{v} \in (0.5, 1)$ and $\hat{u} \in (0.25, 0.75)$ at $\zeta = 2$: after observing a 100% news arrival rate (one good, one bad), every agent believes news is at least 50% likely. A notable feature appears in panel (d): the risk-neutral measure at $(\hat{v}^*, \hat{u}^*) = (0.889, 0.167)$ falls *below* the effective domain floor at $\hat{u} = 0.25$. This occurs because the stochastic discount factor over-weights bad states, pushing u^* below any agent's posterior belief. Individual learning has eliminated the population of extreme pessimists who would normally "inhabit" the risk-neutral measure. This phenomenon cannot occur in the dogmatic case, where the full $(0, 1)^2$ is always populated.

Discussion. The Bayesian extension reveals three features that reinforce the main text's conclusions. First, the dogmatic case is the conservative benchmark: individual learning compounds with market selection, amplifying rather than dampening the pricing effects of heterogeneous beliefs. Second, the effective belief domain (I.A.9) provides a visual diagnostic of learning speed—its compression rate is controlled by the ratio of the sample size (t or ν_t) to the prior precision ζ. When ζ is large relative to the number of trading periods, individual learning has negligible impact and the dogmatic model is an accurate approximation. Third, when ζ is small enough that individual learning materially compresses the belief domain, the risk-neutral measure can fall outside the support of any agent's effective beliefs—a qualitative departure from the dogmatic case that highlights the role of extreme agents in spanning the pricing kernel.

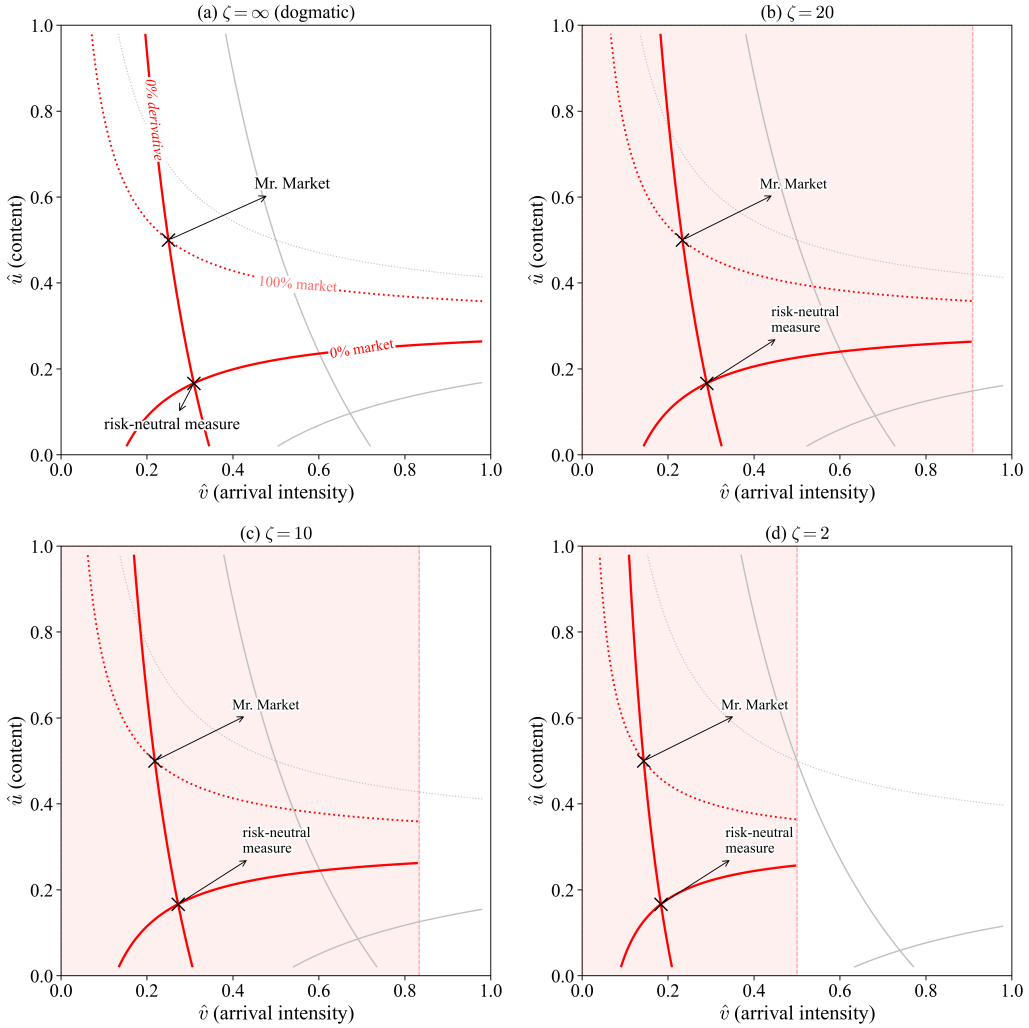


Figure I.A.5: **Portfolio boundaries with Bayesian agents: no-news path** ($\nu_2 = 0$, $n_2 = 0$). Each panel plots the boundaries of agents' asset allocation on the effective belief domain $(\hat{v}, \hat{u})$ at $t = 2$, for $\zeta \in \{\infty, 20, 10, 2\}$. Panel (a) is the dogmatic case, identical to Figure 3(a). Red curves are portfolio boundaries (solid: 0% market and 0% derivative; dotted: 100% market); gray curves show the $t = 0$ reference. The shaded rectangle marks the effective belief domain (I.A.9). Crosses ($\times$) mark Mr. Market and the risk-neutral measure (all panels). Labels and curve annotations appear in panel (a) only. Terminal payoff $p_T(n_T) = e^{\sigma n_T}$ with $T = 3$, $\sigma = \log\sqrt{5}$, and uniform initial wealth distribution.

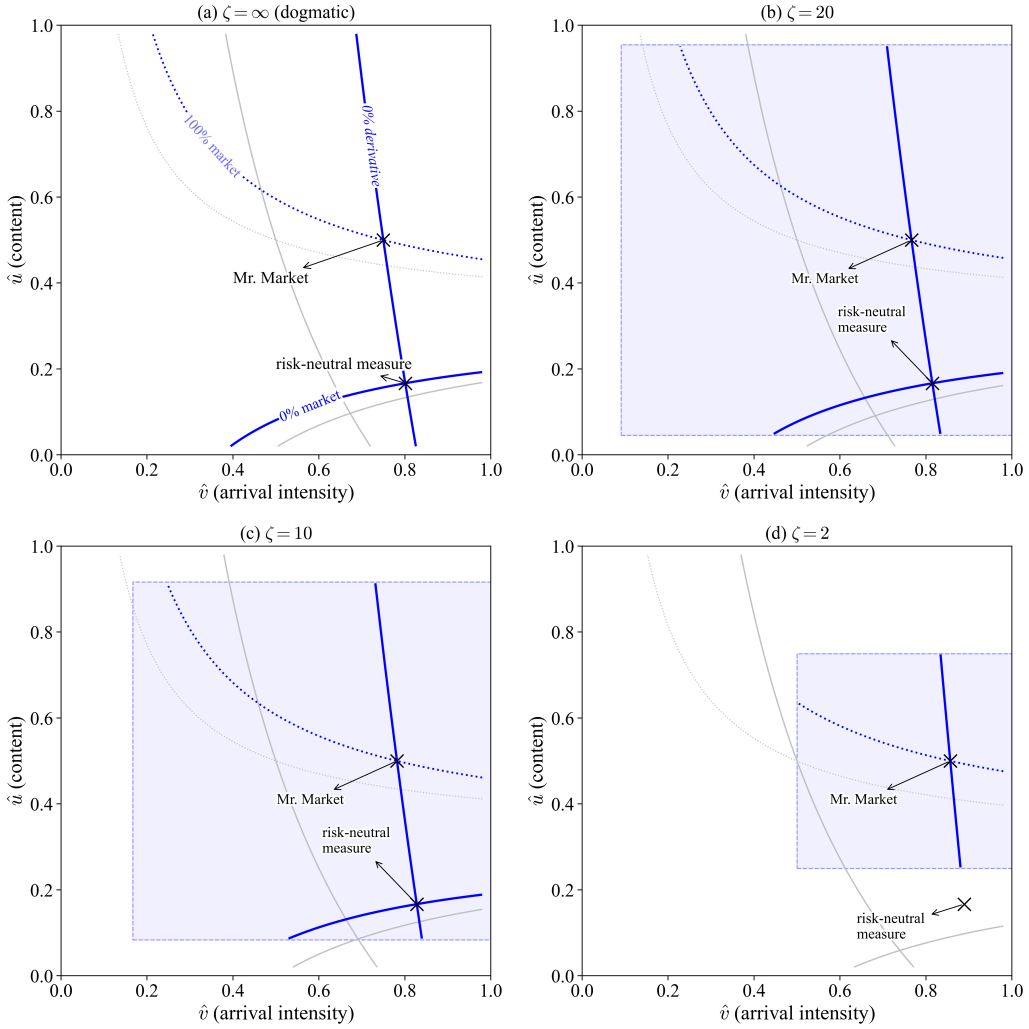


Figure I.A.6: **Portfolio boundaries with Bayesian agents: high realised intensity** ($\nu_2 = 2$, $n_2 = 0$). Same layout as Figure I.A.5. At $\zeta = 2$ (panel d), the effective domain shrinks to $\hat{v} \in (0.5, 1)$, $\hat{u} \in (0.25, 0.75)$, and the risk-neutral measure ($\times$) falls outside the shaded rectangle: $\hat{u}^* = 0.167 < 0.25$. No agent in the Bayesian population holds posterior beliefs as pessimistic as risk-neutral measure.

I.A.5 Extension to CRRA Preferences

Assume all agents have CRRA preferences with risk aversion parameter γ over terminal wealth. In this case, the representative agent (Mr. Market) differs from the wealth-weighted average of agents' beliefs. Since markets are complete, the competitive allocation is Pareto optimal. We therefore characterize the equilibrium by constructing a social

welfare function with Negishi weights $\psi_t(v,u)$ (Negishi (1960)) at time t:

$$\max_{w_T^{(v,u)}} \int \psi_t(v,u)\, \mathbb{E}_t^{(u,v)} \left[\frac{\left(w_T^{(v,u)}\right)^{1-\gamma}}{1-\gamma} \right] dvdu \quad \text{subject to} \quad \int w_T^{(v,u)}(n_T) dvdu = p_T(n_T)\,,$$

where the Negishi weight is inversely proportional to subjective expected marginal utility of terminal wealth $\mathbb{E}_t^{(u,v)} \left[\left(w_T^{(v,u)}\right)^{-\gamma} \right]$.

Up to a normalization constant, the Lagrange multiplier associated with the state-by-state feasibility constraint corresponds to the Arrow–Debreu state price. Applying the first-order condition and substituting $w_T^{(v,u)}$ into the feasibility constraint implies that the Arrow–Debreu price of state n_T at time t, $q_t(n_T)$, is proportional to the marginal utility of the terminal payoff multiplied by a consensus probability of that state.

$$q_t(n_T) \propto p_T^{-\gamma}(n_T) \left[\int \left(\psi_t(v,u) \Pr^{(v,u)}(n_T) \right)^{\frac{1}{\gamma}} dudv \right]^{\gamma}. \tag{I.A.14}$$

Using the zero interest rate condition $\sum_{n_T} q_t(n_T) = 1$, we obtain the market price

$$p_t = \sum_{n_T} q_t(n_T) p_T(n_T) = \frac{\sum_{n_T} p_T^{1-\gamma}(n_T) \left[\int \left(\psi_t(v,u) \Pr^{(v,u)}(n_T) \right)^{\frac{1}{\gamma}} dudv \right]^{\gamma}}{\sum_{n_T} p_T^{-\gamma}(n_T) \left[\int \left(\psi_t(v,u) \Pr^{(v,u)}(n_T) \right)^{\frac{1}{\gamma}} dudv \right]^{\gamma}}. \tag{I.A.15}$$

$\psi_t(v,u)$ plays a role analogous to the wealth distribution $f_t(v,u)$ under logarithmic preferences in the main text, aggregating agents' heterogeneous beliefs into the equilibrium price. Under logarithmic preference, $\psi_t(v,u)$ coincides exactly with $f_t(v,u)$. Setting $\gamma = 1$ in (I.A.15) therefore recovers Result 1.

Because the allocation $w_T^{(v,u)}$ implied by the Negishi weights is identical across periods, the Negishi weights satisfy the following evolution equation, under which the weight increases for agents with more ex-post accurate beliefs. This yields the following result.

Result IA3 (Preference-Invariant Selection)**.** *In any complete-market economy with CRRA preferences, the Negishi weight of agent (v,u) relative to agent (v',u') satisfies*

$$\frac{\psi_t(v,u)}{\psi_t(v',u')} = \frac{\psi_0(v,u)}{\psi_0(v',u')} \prod_{\tau=1}^{t} \frac{\Pr^{(v,u)}(s_\tau)}{\Pr^{(v',u')}(s_\tau)}\,, \tag{I.A.16}$$

which is independent of γ, the risk-neutral measure, and all equilibrium prices. In particular, if $\psi_0(v,u)$ is specified by the same functional form as $f_0(v,u)$ in (1), then $\psi_t(v,u)$ has the same functional form as $f_t(v,u)$ in (9) for all t, γ, and sample paths.

Proof. In a complete market, the first-order condition for agent (v,u) in state n_T is:

$$\psi_t(v,u)\Pr^{(v,u)}(n_T)\big(w_T^{(v,u)}(n_T)\big)^{-\gamma} = \ell_t(n_T)$$

where $\ell_t(n_T)$ is the Lagrange multiplier on the feasibility constraint. This gives $w_T^{(v,u)}(n_T) = \big(\psi_t(v,u)\Pr^{(v,u)}(n_T)/\ell(n_T)\big)^{1/\gamma}$. For this allocation to be consistent across periods (i.e., $w_T^{(v,u)}(n_T)$ does not depend on t for the same equilibrium), the Negishi weight evolution must satisfy $\psi_{t+1}/\psi_t = \Pr^{(v,u)}(s_{t+1})\ell_{t+1}(n_T)/\ell_t(n_T)$. But $\ell_{t+1}(n_T)/\ell_t(n_T)$ cancels in the cross-sectional ratio (I.A.16), leaving a pure likelihood ratio. □

The market selection mechanism and the survival dynamics of heterogeneous beliefs are therefore invariant to risk aversion γ: the central planner's welfare weights shift toward the most accurate forecasters at the same rate as log-utility economy wealth shares.

Under CRRA preferences with any $\gamma > 0$, the Negishi weight of an agent with beliefs (v,u) decays exponentially at rate $KL(v_{true},v) + v_{true}\,KL(u_{true},u)$ per trading period, as in Result 4 . Consequently, the Poisson-limit survival of intensity disagreement and the Brownian-limit elimination of intensity disagreement hold exactly as stated in the main text. Every step in the proof of Result 4 carries over verbatim after replacing the wealth share $f(v,u)$ with the Negishi weight $\psi(v,u)$. The Poisson and Brownian limit results follow from KL-divergence properties and therefore do not depend on the pricing model.

The CRRA equilibrium admits a representation in terms of an effective consensus probability measure.

Result IA4 (CRRA Consensus Representation)**.** *Define the CRRA consensus probability of terminal state n_T as*

$$\widetilde{\pi}_t(n_T) = \frac{\left[\int \left(\psi_t(v,u)\Pr^{(v,u)}(n_T)\right)^{\frac{1}{\gamma}} dvdu\right]^{\gamma}}{\sum_n \left[\int \left(\psi_t(v,u)\Pr^{(v,u)}(n)\right)^{\frac{1}{\gamma}} dvdu\right]^{\gamma}}. \tag{I.A.17}$$

Let $\widetilde{\mathbb{P}}_t$ denote the corresponding consensus probability measure. Then the equilibrium price and the risk–neutral variance satisfy

$$p_t = \frac{\widetilde{\mathbb{E}}_t\big[p_T^{1-\gamma}\big]}{\widetilde{\mathbb{E}}_t\big[p_T^{-\gamma}\big]}, \qquad \mathrm{Var}_t^*[p_T/p_t] = \frac{\widetilde{\mathbb{E}}_t\big[p_T^{2-\gamma}\big]\,\widetilde{\mathbb{E}}_t\big[p_T^{-\gamma}\big]}{\Big(\widetilde{\mathbb{E}}_t\big[p_T^{1-\gamma}\big]\Big)^2} - 1. \tag{I.A.18}$$

Proof. Equation (I.A.15) establishes the equilibrium price in a form equivalent to the proposition. For the risk-neutral variance formula,

$$\mathbb{E}_t^*\left[\left(\frac{p_T}{p_t}\right)^2\right] = \frac{\sum_{n_T} q_t(n_T)\, p_T^2(n_T)/p_t^2}{\sum_{n_T} q_t(n_T)} = \frac{\sum_{n_T} p_T^{2-\gamma}\, \widetilde{\pi}_t(n_T)}{\sum_{n_T} p_T^{-\gamma}\, \widetilde{\pi}_t(n_T)} \cdot \frac{\left(\sum_{n_T} p_T^{-\gamma}\, \widetilde{\pi}_t(n_T)\right)^2}{\left(\sum_{n_T} p_T^{1-\gamma}\, \widetilde{\pi}_t(n_T)\right)^2}$$
$$= \frac{\widetilde{\mathbb{E}}_t[p_T^{2-\gamma}]\, \widetilde{\mathbb{E}}_t[p_T^{-\gamma}]}{(\widetilde{\mathbb{E}}_t[p_T^{1-\gamma}])^2}.$$

Subtracting 1 gives $\mathrm{Var}_t^*[p_T/p_t]$. □

The first expression shows that the market price equals the ratio of two consensus expectations of different moments of the terminal payoff, while the second shows that the risk-neutral variance depends on other moments of p_T under the same measure.

Because the Negishi weights $\psi_t(v,u)$ follow exactly the same Beta-distribution updating rule as the log-utility wealth shares, the dispersion of the CRRA consensus measure $\widetilde{\pi}_t(n_T)$ responds to realized news in a similar manner. Using the expression for $\psi_t(v,u)$ in (9) and substituting it into (I.A.17) gives

$$\widetilde{\pi}_t(n_T) \propto \sum_{\substack{\forall n_g - n_b = n_T - n_t, \\ n_g + n_m + n_b = T-t}} \frac{(T-t)!}{n_m! n_g! n_b!} \frac{\left[B(\frac{\alpha_{vt}+T-t-n_m+\gamma-1}{\gamma}, \frac{\beta_{vt}+n_m+\gamma-1}{\gamma})\right]^\gamma}{B(\alpha_{vt},\beta_{vt})} \frac{\left[B(\frac{\alpha_{ut}+n_g+\gamma-1}{\gamma}, \frac{\beta_{ut}+n_b+\gamma-1}{\gamma})\right]^\gamma}{B(\alpha_{ut},\beta_{ut})}. \tag{I.A.19}$$

When disagreement over the news arrival rate (v) is high, periods without news concentrate Negishi weights among low-v agents. This shift increases the likelihood of paths with high n_m and consequently compresses the variance of n_T under $\widetilde{\mathbb{P}}_t$. In contrast, when there is no disagreement about the news arrival rate, periods without news leave wealth broadly distributed across agents with heterogeneous u beliefs. This increases the likelihood of paths with large net news $n_g - n_b$ and thus raises the variance of $n_T = n_t + n_g - n_b$. These forces carry over to the risk-neutral variance in (I.A.18) for exponential terminal payoffs $p_T = e^{\sigma n_T}$. Consequently, the qualitative distinction of volatility persistence remains a robust qualitative feature of the equilibrium.

The impact of realized news frequency on asset valuations may instead depend on the risk aversion. For typical equity investors with relative risk aversion $\gamma > 1$, a less dispersed distribution $\widetilde{\pi}_t(n_T)$ implies a higher price because $p_T^{-\gamma}$ exhibits greater curvature

than $p_T^{1-\gamma}$. Consequently, when disagreement about the news arrival rate is high, the "no news is good news" finding remains robust. This finding need not hold when $\gamma < 1$. In particular, as $\gamma \to 0$, $p_t \approx \widetilde{\mathbb{E}}_t[p_T]$ can increase when $\widetilde{\pi}_t(n_T)$ becomes more dispersed, since investors who assign greater weight to high-volatility outcomes bid up prices.

I.A.6 Two Risky Assets

The M-node tree accommodates multiple assets, but rather than pursuing full generality we focus on a single question: how does concentrated speculation in one asset affect the broader market risk premium and individual market betas?

Consider a setting with $M = 4$ and $K = 2$. In each period, binary news—either 'good news' or 'bad news'—arrives for both assets A and B. We construct the 'market-wide' news by combining these individual shocks, defining the state space as $(good, good)$, $(good, bad)$, $(bad, good)$, and (bad, bad).[2] To ensure dynamic completeness, agents trade assets A and B (unit net supply), a one-period risk-free bond, and a one-period derivative with zero net supply.

When good news arrives for asset A, agents who are bullish on asset A accumulate wealth. Their views regarding asset B then influence asset B's pricing. If news frequently arrives in opposite directions for the two assets—the $(good, bad)$ and $(bad, good)$ combinations—the aggregate market generally appears less volatile since the joint news shocks are largely idiosyncratic. Conversely, when news frequently arrives in the same direction, alternating between extremely good (both assets have good news) and extremely bad (both assets have bad news), the aggregate market exhibits higher volatility. Building on this observation, we specify the one-period derivative contract as a market volatility swap: it yields positive payoffs in states $(good, good)$ and (bad, bad), and zero otherwise.

[2]Because all four states involve news for at least one asset, the no-news channel is absent in the diffusion limit of this two-asset extension. The intensity-disagreement mechanism operates here only through the Poisson (jump) component of the two-asset economy (Internet Appendix I.A.6.2), where co-jump and idiosyncratic jump intensities play the role of arrival-rate disagreement. A richer $M > 4$ specification incorporating asset-specific no-news states would preserve the full news/no-news distinction in the multi-asset diffusion limit.

I.A.6.1 A Diffusion Limit

We consider the continuous-time limit to obtain tractable formulas for prices and risk-neutral return variances. The interval $[0, T]$ is partitioned into $N = T/\delta$ periods, each of length δ. The market consists of assets A and B, with terminal payoffs $p_{i,T}(n_{i,T}) = e^{\sigma_i \sqrt{\frac{T}{N}} n_{i,T}}$ for $i \in \{A, B\}$, where the composite market index is $p_T = p_{A,T} + p_{B,T}$.

To allow for distinct belief heterogeneity regarding the news of A and B (i.e., $\theta_A \neq \theta_B$), we model the population mass using two separate Beta distributions, $B(\alpha_i, \beta_i)$, joined by a *bivariate Bernoulli copula*. The parameters for asset $i \in \{A, B\}$ are independently scaled as $\alpha_{i,0} = \frac{1}{2}\theta_i N + \frac{1}{2}\eta_i \sqrt{N}$ and $\beta_{i,0} = \frac{1}{2}\theta_i N - \frac{1}{2}\eta_i \sqrt{N}$, where $\alpha_{i0} + \beta_{i0} = \theta_i N$.[3] In the large-$N$ limit, the wealth-weighted beliefs over log payoffs converge to a two-dimensional Gaussian distribution with mean $\mu_{i,t} = \frac{1-\phi}{\theta_i+\phi}\eta_i \sigma_i \sqrt{T} + \frac{\theta_i+1}{\theta_i+\phi}\eta_{i,t}\sigma_i\sqrt{\phi T}$, variance $\sigma^2_{i,t} = \frac{\theta_i+1}{\theta_i+\phi}(1-\phi)\sigma_i^2 T$, and covariance $\sigma_{AB,t} = \sqrt{\frac{\theta_A+1}{\theta_A+\phi}\frac{\theta_B+1}{\theta_B+\phi}}(1-\phi)\rho\sigma_A\sigma_B T$, where $\phi = t/T$ and $\eta_{i,t} = n_{i,t}/\sqrt{\phi N}$.

The following result gives closed-form prices and risk-neutral variances for individual assets and the aggregate market.

Result IA5 (Brownian Limit of Two Risky Assets)**.** *Let $\mu_{Z,t} = \mu_{A,t} - \mu_{B,t}$, and $\sigma^2_{Z,t} = \sigma^2_{A,t} + \sigma^2_{B,t} - 2\sigma_{AB,t}$. In the Brownian limit, the market price $p_t = p_{A,t} + p_{B,t}$ is given by:*

$$p_t = \exp\left(\mu_{A,t} - \frac{1}{2}\sigma^2_{A,t}\right) \mathscr{W}\left(\kappa_t\left(\mu_{Z,t} - \sigma^2_{A,t} + \sigma_{AB,t}\right)\right), \tag{I.A.20}$$

where $\mathscr{W}(x) = 1 + e^{-x}$ and the scaling factor is $\kappa_t = (1 + \frac{\pi}{8}\sigma^2_{Z,t})^{-1/2}$. The market share of asset A is $\frac{p_{A,t}}{p_t} = \mathscr{W}\left(\kappa_t\,\mu_{Z,t}\right)^{-1}$.

The risk-neutral variance of the market return is

$$\mathrm{Var}^*_t[p_T/p_t] = \frac{\exp(\sigma^2_{A,t})\mathscr{W}\left(\mu_{Z,t} + \frac{1}{2}(\sigma^2_{A,t} - \sigma^2_{B,t})\right)}{\mathscr{W}\left(\kappa_t\left(\mu_{Z,t} - \sigma^2_{A,t} + \sigma_{AB,t}\right)\right)} - 1\,. \tag{I.A.21}$$

[3] Specifically, at time t (or step ϕN), let $q_{A,t} = \alpha_{A,t}/(\alpha_{A,t} + \beta_{A,t})$ and $q_{B,t} = \alpha_{B,t}/(\alpha_{B,t} + \beta_{B,t})$ denote the marginal posterior probabilities of 'good news' for each asset. The simultaneous draws of $s_i \in \{\text{'good'}, \text{'bad'}\} \times \{\text{'good'}, \text{'bad'}\}$ have the joint distribution: $h_{1,t} = q_{A,t}q_{B,t} + \rho\varepsilon_t$, $h_{2,t} = q_{A,t}\bar{q}_{B,t} - \rho\varepsilon_t$, $h_{3,t} = \bar{q}_{A,t}q_{B,t} - \rho\varepsilon_t$, $h_{4,t} = \bar{q}_{A,t}\bar{q}_{B,t} + \rho\varepsilon_t$, with $q_{i,t} + \bar{q}_{i,t} = 1$ and $\varepsilon_t^2 = q_{A,t}\bar{q}_{A,t}q_{B,t}\bar{q}_{B,t}$. The parameter $\rho \in (-1, 1)$ represents the fundamental correlation between individual net news shocks. Summing over either asset immediately recovers its exact marginal probability.

The risk-neutral variance of $p_{A,T}/p_{A,t}$ and its covariance with the market return are

$$\mathrm{Var}_t^*[p_{A,T}/p_{A,t}] = \frac{\exp(\sigma_{A,t}^2)\mathscr{W}\left(\kappa_t\,\mu_{Z,t}\right)^2}{\mathscr{W}\left(\kappa_t\left(\mu_{Z,t}-\sigma_{A,t}^2+\sigma_{AB,t}\right)\right)\mathscr{W}\left(\kappa_t\left(\mu_{Z,t}+\sigma_{A,t}^2-\sigma_{AB,t}\right)\right)}-1\,, \quad \text{(I.A.22)}$$

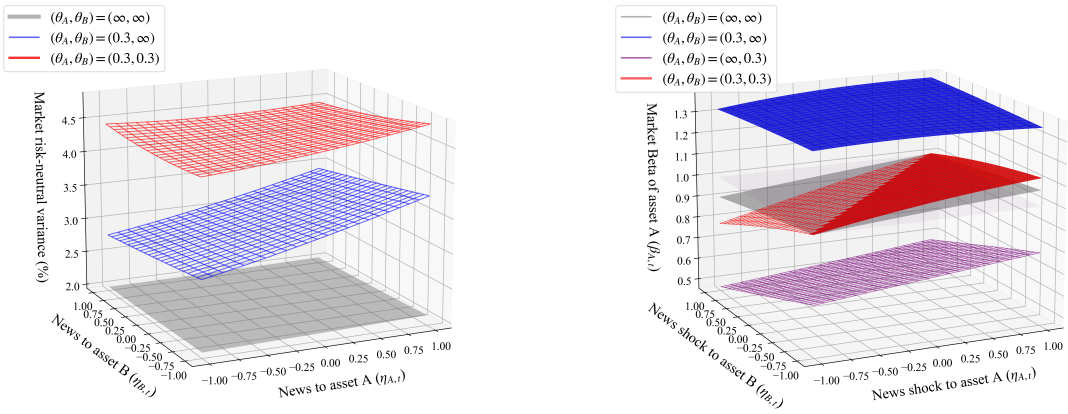
$$\mathrm{Cov}_t^*[p_{A,T}/p_{A,t},\, p_T/p_t] = \frac{\exp(\sigma_{A,t}^2)\mathscr{W}\left(\kappa_t\,\mu_{Z,t}\right)}{\mathscr{W}\left(\kappa_t\left(\mu_{Z,t}-\sigma_{A,t}^2+\sigma_{AB,t}\right)\right)}-1\,. \quad \text{(I.A.23)}$$

When the two assets are identical—$\sigma_A = \sigma_B$, $\rho = 1$, $\theta_A = \theta_B$, and identical news histories ($\eta_{A,t} = \eta_{B,t}$)—every $\mathscr{W}(\cdot)$ term reduces to a constant and Result IA5 collapses to Result 6.[4] Under logarithmic utility, the no-arbitrage condition dictates that any expected risky excess return satisfies $\mathbb{E}_t[p_{i,T}/p_{i,t}] - 1 = \mathrm{Cov}_t^*[p_{i,T}/p_{i,t},\, p_T/p_t]$, as demonstrated by Martin and Wagner (2019). Consequently, we can compute the market beta of asset A by taking the ratio of (I.A.23) to (I.A.21).

To isolate these cross-sectional dynamics, we adopt a symmetric parameterization for Figure I.A.7. We set fundamental volatilities equal ($\sigma_A = \sigma_B = 20\%$) and, importantly, fix the fundamental correlation to zero ($\rho = 0$). This zero-correlation benchmark ensures that any observed cross-asset spillovers, volatility amplifications, or beta compressions are entirely endogenous—driven purely by the structure of heterogeneous beliefs rather than mechanical fundamental linkages. We evaluate these dynamics at an intermediate time $t/T = 0.3$, contrasting the homogeneous benchmark ($\theta_A = \theta_B = \infty$) with cases of localized and generalized disagreement where at least one $\theta_i = 0.3$ for $i \in \{A, B\}$.

Localized speculation generates cross-sectional spillovers even absent fundamental linkages. Figure I.A.7 illustrates: when disagreement concentrates on asset A alone, its inflated risk-neutral variance dominates the market index, mechanically elevating A's beta above one while compressing B's beta below one (right panel, purple surface). Asset B endogenously becomes a safe haven despite unchanged fundamentals. When disagreement is high for both assets (left panel), market variance peaks; when both are homogeneous, it is minimized. The betas align with the homogeneous benchmark only when both assets' news histories are identical ($\eta_{A,t} = \eta_{B,t}$).

[4] Equations (I.A.20)–(I.A.23) are *MacKay probit approximations* with very small error of order of $O(\sigma_{Z,t}^4)$. The exact two-asset pricing solution is available via Fourier inversion (Internet Appendix, Result IA6); the closed forms here are adopted for their transparent decomposition into market shares and (co)variances.



(a) Market risk premium

(b) Market beta of asset A

Figure I.A.7: **The 'Illusion' of Safety.** We fix $\phi = t/T = 0.3$, $\sigma_A = \sigma_B = 20\%$, and $\rho = 0$ in both figures. The parameter θ_i indicates the level of agreement regarding asset i, with $\theta_i = \infty$ representing the homogeneous economy limit. Because our parameter choices for Assets A and B are symmetric, the purple surface is mathematically identical to the market beta of B when $(\theta_A, \theta_B) = (0.3, \infty)$.

I.A.6.2 A Poisson Limit

We consider the continuous-time Poisson limit to obtain tractable formulas for prices and risk-neutral jump intensity in a market with 2 risky assets.

The interval $[0, T]$ is partitioned into $N = T/\delta$ periods, each of length δ. The market consists of assets A and B. In any interval δ, four mutually exclusive states can occur: no jumps, an idiosyncratic jump in A only, an idiosyncratic jump in B only, or a simultaneous co-jump in both. When a jump occurs in asset i, its fundamental value drops by a fractional amount $\Delta_i = 1 - e^{-\omega_i}$. The terminal payoffs are $p_{i,T}(s_{i,T}) = e^{-\omega_i s_{i,T}}$, where $s_{i,T}$ is the total count of jumps experienced by asset i.

To model heterogeneous beliefs, we scale the Dirichlet parameters of the 4-node tree. Let $\theta > 0$ represent the disagreement among agents. At $t = 0$, the three jump states are initialized as $\beta_{j,0} = \lambda_{j,0}\theta$ for $j \in \{A, B, AB\}$, while the no-jump state holds the remainder: $\beta_{0,0} = \frac{\theta}{\delta} - (\beta_{A,0} + \beta_{B,0} + \beta_{AB,0})$.

As the discrete step size $\delta \to 0$, the probability of the no-jump state converges to 1. This completely decouples the jump branches, so that the aggregate wealth-weighted

physical probabilities converge from a joint Dirichlet distribution to three independent continuous-time Gamma-Poisson processes.

If $n_{j,t}$ events of type j accumulate over time t, the updated Dirichlet parameters (convictions) evolve exactly as $\beta_{j,t} = \beta_{j,0} + n_{j,t}$. The wealth-weighted physical jump intensities are therefore continuously updated based on the accumulated events:

$$\lambda_{j,t} = \frac{\beta_{j,t}}{\theta + t}\,, \quad \text{for } j \in \{A, B, AB\}\,. \tag{I.A.24}$$

At time T, $Y_A = \ln p_{A,T}$ and $Y_B = \ln p_{B,T}$ denote the terminal log-prices. At time t, we have $Y_A = y_{A,t} - \omega_A n_{A,\tau} - \omega_A n_{AB,\tau}$ and $Y_B = y_{B,t} - \omega_B n_{B,\tau} - \omega_B n_{AB,\tau}$, where $y_{i,t}$ is the known log-fundamental at time t, $\tau = T - t$ is the remaining time to maturity, and $n_{j,\tau}$ are the unknown future jump counts. We define $\kappa_t = \frac{\tau}{\theta+t}$ for convenience.

The joint characteristic function $\Phi_t(z_1, z_2) = \mathbb{E}_t[\exp(iz_1 Y_A + iz_2 Y_B)]$ factors into the product of the respective moment generating functions of the three independent jumps. For complex arguments $z_1, z_2 \in \mathbb{C}$, the exact analytical closed form at time t is:

$$\begin{aligned}\Phi_t(z_1, z_2) = \exp\big(iz_1 y_{A,t} + iz_2 y_{B,t}\big) \times \big[1 - \kappa_t\big(e^{-iz_1\omega_A} - 1\big)\big]^{-\beta_{A,t}} \\ \times \big[1 - \kappa_t\big(e^{-iz_2\omega_B} - 1\big)\big]^{-\beta_{B,t}} \times \big[1 - \kappa_t\big(e^{-iz_1\omega_A - iz_2\omega_B} - 1\big)\big]^{-\beta_{AB,t}}\,,\end{aligned} \tag{I.A.25}$$

where $\beta_{j,t}$ are updated Dirichlet parameters tracking the accumulation of realized jumps.

The market weight of asset i at time t is $w_{i,t} = p_{i,t}/p_t$. The market price p_t and individual asset prices $p_{i,t}$ can be solved via Fourier inversion in the complex plane.

Result IA6 (Market price and market share)**.** *The inverse market price $p_t^{-1} = \mathbb{E}_t[(e^{Y_A} + e^{Y_B})^{-1}]$ is given by a single one-dimensional Fourier integral:*

$$p_t^{-1} = \frac{1}{2}\int_{-\infty}^{\infty} \frac{\Phi_t\big(-u + i\alpha,\, u - i(\alpha - 1)\big)}{\sin\big(\pi(iu + \alpha)\big)} du\,, \tag{I.A.26}$$

where $u \in \mathbb{R}$, and $\alpha \in (0, 1)$ is an arbitrary number. The market share of asset A at time t satisfies:

$$w_{A,t} = \frac{p_{A,t}}{p_t} = \frac{1}{2}\int_{-\infty}^{\infty} \frac{\Phi_t\big(u - i\alpha,\, -u + i\alpha\big)}{\sin\big(\pi(iu + \alpha)\big)} du\,. \tag{I.A.27}$$

The risk-neutral variance of the aggregate market return p_T/p_t is:

$$\mathrm{Var}_t^*(p_T/p_t) = \frac{e^{y_{A,t}}}{p_t}\mathscr{J}_A(\kappa_t)^{-(\beta_{A,t}+\beta_{AB,t})} + \frac{e^{y_{B,t}}}{p_t}\mathscr{J}_B(\kappa_t)^{-(\beta_{B,t}+\beta_{AB,t})} - 1\,, \tag{I.A.28}$$

where the function $\mathscr{J}_i(x) = (1 - x(e^{-\omega_i} - 1))$.

The risk-neutral variance of Asset A and its covariance with market satisfy:

$$\mathrm{Var}_t^*(p_{A,T}/p_{A,t}) = \frac{p_t}{p_{A,t}^2}\left[\frac{1}{2}\int_{-\infty}^{\infty}\frac{\Phi_t\big(u - i(1+\alpha),\, -u + i\alpha\big)}{\sin\big(\pi(iu+\alpha)\big)}du\right] - 1\,, \tag{I.A.29}$$

$$\mathrm{Cov}_t^*(p_{A,T}/p_{A,t}, p_T/p_t) = \frac{e^{y_{A,t}}}{p_{A,t}}\mathscr{J}_A(\kappa_t)^{-(\beta_{A,t}+\beta_{AB,t})} - 1\,. \tag{I.A.30}$$

Result IA7 (Jump intensity and correlation)**.** *Let the physical intensities be* $\lambda_{A,t}$, $\lambda_{B,t}$, *and* $\lambda_{AB,t}$. *The corresponding risk-neutral jump intensities are:*

$$\lambda_{A,t}^* = \lambda_{A,t}\mathscr{D}_{A,t}^{-1}\,,\quad \lambda_{B,t}^* = \lambda_{B,t}\mathscr{D}_{B,t}^{-1}\,,\quad \lambda_{AB,t}^* = \lambda_{AB,t}\left(\mathscr{D}_{A,t} + \mathscr{D}_{B,t} - 1\right)^{-1}\,, \tag{I.A.31}$$

where $\mathscr{D}_{i,t} = 1 - w_{i,t}\Delta_i$.

The instantaneous risk-neutral variance of Asset i's return, and its risk-neutral covariance with Asset j's return are given by:

$$\mathrm{Var}_t^*(dR_i)/dt = (\lambda_{i,t}^* + \lambda_{ij,t}^*)\Delta_i^2\,,\quad \mathrm{Cov}_t^*(dR_i, dR_j)/dt = \lambda_{ij,t}^*\Delta_i\Delta_j\,. \tag{I.A.32}$$

Let ρ_t *denote the physical correlation of jump arrivals, and* ρ_t^* *denote the risk-neutral correlation. We have* $|\rho_t^*| > |\rho_t|$, *and specifically,*

$$\rho_t^* = \rho_t \times \frac{\sqrt{(\lambda_{A,t} + \lambda_{AB,t})(\lambda_{B,t} + \lambda_{AB,t})}}{\sqrt{\big(\lambda_{A,t}(1 - \mathscr{D}_{A,t}^{-1}(1 - \mathscr{D}_{B,t})) + \lambda_{AB,t}\big)\big(\lambda_{B,t}(1 - \mathscr{D}_{B,t}^{-1}(1 - \mathscr{D}_{A,t})) + \lambda_{AB,t}\big)}}\,. \tag{I.A.33}$$

I.A.6.3 Proofs for results

Proof of Result IA5. Let $Y_{i,T} = \log p_{i,T}$ be the terminal log-payoff of asset i. In the diffusion limit, the joint distribution is conditionally Gaussian under the measure $\mathbb{P}$, with $Y_{i,T} \sim \mathcal{N}(\mu_{i,t}, \sigma_{i,t}^2)$. This follows a similar argument as in the proof for Result 6 using Paul and Plackett (1978). We define the cross-sectional state variable $Z_T = Y_{A,T} - Y_{B,T}$, which is normally distributed as $Z_T \sim \mathcal{N}(\mu_{Z,t}, \sigma_{Z,t}^2)$.

The terminal market share of Asset A can be rewritten exactly using the $\mathscr{W}(x) = 1 + e^{-x}$ notation: $\frac{p_{A,T}}{p_T} = \frac{e^{Y_{A,T}}}{e^{Y_{A,T}} + e^{Y_{B,T}}} = \frac{1}{1+\exp(-Z_T)} = \mathscr{W}(Z_T)^{-1}$. Taking the expectation using the MacKay Probit mapping, $\mathbb{E}[\mathscr{W}(X)^{-1}] \approx \mathscr{W}(\mu_X\kappa(\sigma_X^2))^{-1}$, immediately proves the market share identity: $\frac{p_{A,t}}{p_t} = \mathbb{E}_t[\mathscr{W}(Z_T)^{-1}] = \mathscr{W}\big(\kappa_t\mu_{Z,t}\big)^{-1}$.

The inverse market price is the expectation of the terminal inverse market price $p_t^{-1} = \mathbb{E}_t[p_T^{-1}] = \mathbb{E}_t\left[e^{-Y_{A,T}}\mathscr{W}(Z_T)^{-1}\right]$. We apply a change of measure to $\tilde{\mathbb{P}}$, defined by the Radon-Nikodym derivative $\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} = \frac{e^{-Y_{A,T}}}{\mathbb{E}_t[e^{-Y_{A,T}}]}$. The expectation in the denominator is the standard log-normal moment: $\mathbb{E}_t[e^{-Y_{A,T}}] = \exp(-\mu_{A,t} + \frac{1}{2}\sigma_{A,t}^2)$. Under $\tilde{\mathbb{P}}$, Z_T's variance is unchanged, but its mean shifts by its covariance with tilting variable $-Y_{A,T}$: $\mu_{Z,t}^{\tilde{\mathbb{P}}} = \mu_{Z,t} + \mathrm{Cov}_t(-Y_{A,T}, Z_T) = \mu_{Z,t} - \sigma_{A,t}^2 + \sigma_{AB,t}$. Applying the MacKay approximation under $\tilde{\mathbb{P}}$ yields: $\mathbb{E}_t^{\tilde{\mathbb{P}}}[\mathscr{W}(Z_T)^{-1}] = \mathscr{W}\left(\kappa_t(\mu_{Z,t} - \sigma_{A,t}^2 + \sigma_{AB,t})\right)^{-1}$. Combining these components, we obtain $p_t^{-1} = \exp\left(-\mu_{A,t} + \frac{1}{2}\sigma_{A,t}^2\right)\mathscr{W}\left(\kappa_t(\mu_{Z,t} - \sigma_{A,t}^2 + \sigma_{AB,t})\right)^{-1}$. Inverting yields (I.A.20).

Next, the risk-neutral variance of the market return is defined under the pricing measure $\mathbb{Q}$: $\mathrm{Var}_t^*[p_T/p_t] = \mathbb{E}_t^*\left[(p_T/p_t)^2\right] - 1 = \frac{1}{p_t}\mathbb{E}_t[p_T] - 1$. The expectation of the market is the sum: $\mathbb{E}_t[p_T] = F_{A,t} + F_{B,t} = F_{A,t}(1 + F_{B,t}/F_{A,t})$. We evaluate the forward ratio as $\frac{F_{B,t}}{F_{A,t}} = \exp\left(\mu_{B,t} + \frac{1}{2}\sigma_{B,t}^2 - \mu_{A,t} - \frac{1}{2}\sigma_{A,t}^2\right) = \exp\left(-\left[\mu_{Z,t} + \frac{1}{2}(\sigma_{A,t}^2 - \sigma_{B,t}^2)\right]\right)$. Thus, $(1 + F_{B,t}/F_{A,t}) = \mathscr{W}\left(\mu_{Z,t} + \frac{1}{2}(\sigma_{A,t}^2 - \sigma_{B,t}^2)\right)$. From our previous derivation of p_t^{-1}, multiplying by $F_{A,t} = \exp(\mu_{A,t} + \frac{1}{2}\sigma_{A,t}^2)$ yields:

$$\frac{F_{A,t}}{p_t} = \frac{\exp(\sigma_{A,t}^2)}{\mathscr{W}\left(\kappa_t(\mu_{Z,t} - \sigma_{A,t}^2 + \sigma_{AB,t})\right)}\,. \tag{I.A.34}$$

Multiplying (I.A.34) by $(1 + F_{B,t}/F_{A,t})$ and subtracting 1 recovers equation (I.A.21).

Next, for the individual asset risk-neutral variance, we evaluate the second non-central moment: $\mathrm{Var}_t^*[p_{A,T}/p_{A,t}] = \frac{p_t}{p_{A,t}^2}\mathbb{E}_t\left[\frac{p_{A,T}^2}{p_T}\right] - 1 = \frac{p_t}{p_{A,t}^2}\mathbb{E}_t\left[e^{Y_{A,T}}\mathscr{W}(Z_T)^{-1}\right] - 1$. We apply a forward change of measure $\mathbb{P}'$ defined by $\frac{d\mathbb{P}'}{d\mathbb{P}} = \frac{e^{Y_{A,T}}}{F_{A,t}}$. Under $\mathbb{P}'$, the mean of Z_T shifts by its positive covariance with $Y_{A,T}$: $\mu_{Z,t}^{\mathbb{P}'} = \mu_{Z,t} + \mathrm{Cov}_t(Y_{A,T}, Z_T) = \mu_{Z,t} + \sigma_{A,t}^2 - \sigma_{AB,t}$. Applying the MacKay mapping under this forward measure gives: $\mathbb{E}_t\left[e^{Y_{A,T}}\mathscr{W}(Z_T)^{-1}\right] = F_{A,t}\mathscr{W}\left(\kappa_t(\mu_{Z,t} + \sigma_{A,t}^2 - \sigma_{AB,t})\right)^{-1}$. Next, we evaluate the scalar multiplier. Since $p_{A,t} = p_t\mathscr{W}(\kappa_t\mu_{Z,t})^{-1}$, we have $\frac{p_t}{p_{A,t}^2} = \frac{1}{p_t}\mathscr{W}(\kappa_t\mu_{Z,t})^2$. Combining the multiplier with the evaluated expectation yields $\frac{F_{A,t}}{p_t}\frac{\mathscr{W}(\kappa_t\mu_{Z,t})^2}{\mathscr{W}\left(\kappa_t(\mu_{Z,t} + \sigma_{A,t}^2 - \sigma_{AB,t})\right)} - 1$. Substituting equation (I.A.34) for $F_{A,t}/p_t$ perfectly places the remaining $\mathscr{W}$ term in the denominator and attaches the $\exp(\sigma_{A,t}^2)$ scalar to the numerator, completing the proof for equation (I.A.22).

Finally, for the covariance, we expand the risk-neutral expectation of the gross returns:

$$\mathrm{Cov}_t^*[p_{A,T}/p_{A,t}, p_T/p_t] = \frac{1}{p_{A,t}}\mathbb{E}_t\left[\left(p_{A,T}\frac{p_T}{p_t}\right)\frac{p_t}{p_T}\right] - 1 = \frac{1}{p_{A,t}}\mathbb{E}_t[p_{A,T}] - 1 = \frac{F_{A,t}}{p_{A,t}} - 1\,.$$

Using the identity $p_{A,t} = p_t \mathscr{W}(\kappa_t \mu_{Z,t})^{-1}$, we rewrite this as: $\frac{F_{A,t}}{p_{A,t}} = \frac{F_{A,t}}{p_t}\mathscr{W}(\kappa_t \mu_{Z,t})$. Substituting equation (I.A.34) into this expression yields equation (I.A.23). □

Proof for Result IA6. The inverse market price is the expectation of the inverse terminal market payoff:

$$p_t^{-1} = \mathbb{E}_t\left[\frac{1}{p_{A,T} + p_{B,T}}\right] = \mathbb{E}_t\left[e^{-Y_B}\frac{1}{1+e^{Y_A - Y_B}}\right] \tag{I.A.35}$$

Let $Z = Y_A - Y_B$. We define the complex variable $v = u - i\alpha$, where $u \in \mathbb{R}$ and $\alpha \in \mathbb{R}$ is a free parameter. The Fourier transform of $g(Z)$ along this contour is: $\hat{g}(v) = \int_{-\infty}^{\infty} e^{ivZ}\frac{1}{1+e^Z}dZ$. We apply the substitution $x = e^Z$, which implies $dZ = dx/x$. The limits of integration map from $(-\infty, \infty)$ to $(0, \infty)$. The integral becomes:

$$\hat{g}(v) = \int_0^{\infty} \frac{x^{iv}}{1+x}\frac{dx}{x} = \int_0^{\infty}\frac{x^{iv-1}}{1+x}dx \tag{I.A.36}$$

This is the standard integral representation of the Euler Beta function, $B(s, 1-s) = \int_0^{\infty}\frac{x^{s-1}}{1+x}dx = \Gamma(s)\Gamma(1-s) = \frac{\pi}{\sin(\pi s)}$. For the integral to converge, the real part of the exponent $s = iv$ must strictly lie in the interval $(0, 1)$. Substituting $v = u - i\alpha$, we have $iv = i(u - i\alpha) = iu + \alpha$. Therefore, the strip of analyticity is defined by $0 < \alpha < 1$. Within this strip, the Fourier transform is exact $\hat{g}(v) = \frac{\pi}{\sin(\pi iv)} = \frac{\pi}{\sin\left(\pi(iu+\alpha)\right)}$.

By the Fourier inversion theorem, we can express the original function as:

$$g(Z) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-ivZ}\hat{g}(v)\,du\,.$$

Substituting this back into the expectation for the market price in equation (I.A.35):

$$p_t^{-1} = \mathbb{E}_t\left[e^{-Y_B}\frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-iv(Y_A - Y_B)}\frac{\pi}{\sin(\pi iv)}du\right] = \frac{1}{2}\int_{-\infty}^{\infty}\frac{\mathbb{E}_t\left[\exp\Big(-ivY_A + (iv-1)Y_B\Big)\right]}{\sin(\pi iv)}du$$

The interchange of expectation and integral is justified by Fubini's theorem, given the absolute convergence guaranteed by the damping parameter α.

We now map the terms inside the expectation to the joint characteristic function, $\Phi_t(z_1, z_2) = \mathbb{E}_t[\exp(iz_1 Y_A + iz_2 Y_B)]$. For Asset A, the exponent is $-iv = -iu - \alpha$. Setting $iz_1 = -iu - \alpha$ yields $z_1 = -u + i\alpha$. For Asset B, the exponent is $iv - 1 = iu + \alpha - 1$. Setting $iz_2 = iu + \alpha - 1$ yields $z_2 = u + i(1-\alpha)$. Substituting these arguments into the characteristic function yields the exact formula in equation (I.A.26).

Next, we derive the individual price of Asset A. Under the logarithmic pricing kernel $M_T = p_t/p_T$, the market share of Asset A satisfies $\frac{p_{A,t}}{p_t} = \mathbb{E}_t\left[\frac{1}{1+e^{Y_B-Y_A}}\right]$. Let $W = Y_B - Y_A$. This is identically $g(W) = \frac{1}{1+e^W}$. Applying the same Fourier inversion:

$$\frac{p_{A,t}}{p_t} = \mathbb{E}_t\left[\frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-iv(Y_B-Y_A)}\frac{\pi}{\sin(\pi iv)}du\right] = \frac{1}{2}\int_{-\infty}^{\infty}\frac{\mathbb{E}_t\left[\exp\left(ivY_A - ivY_B\right)\right]}{\sin(\pi iv)}du$$

Mapping the exponents to the joint characteristic function: For Asset A, the exponent is $iv = iu + \alpha$. Setting $iz_1 = iu + \alpha$ yields $z_1 = u - i\alpha$. For Asset B, the exponent is $-iv = -iu - \alpha$. Setting $iz_2 = -iu - \alpha$ yields $z_2 = -u + i\alpha$. Substituting these arguments gives $\Phi_t(u - i\alpha, -u + i\alpha)$, establishing equation (I.A.27) and completing the proof.

For the risk-neutral variance of market, we need to evaluate $\mathbb{E}_t[p_T] = \mathbb{E}_t[e^{Y_A}] + \mathbb{E}_t[e^{Y_B}]$, which corresponds to evaluating the joint characteristic function at the purely imaginary arguments $(z_1 = -i, z_2 = 0)$ and $(z_1 = 0, z_2 = -i)$. This gives the closed-form:

$$\mathbb{E}_t[p_T] = e^{y_{A,t}}\Big[1 - \kappa_t\big(e^{-\omega_A} - 1\big)\Big]^{-(\beta_{A,t}+\beta_{AB,t})} + e^{y_{B,t}}\Big[1 - \kappa_t\big(e^{-\omega_B} - 1\big)\Big]^{-(\beta_{B,t}+\beta_{AB,t})}.$$

For the risk-neutral variance of Asset A:

$$\mathrm{Var}_t^*(R_A) = \mathbb{E}_t^*[R_A^2] - (\mathbb{E}_t^*[R_A])^2 = \mathbb{E}_t\left[\left(\frac{p_t}{p_T}\right)\left(\frac{p_{A,T}}{p_{A,t}}\right)^2\right] - 1 = \frac{p_t}{p_{A,t}^2}\mathbb{E}_t\left[\frac{p_{A,T}^2}{p_{A,T}+p_{B,T}}\right] - 1\,.$$

We rewrite the term inside the expectation using terminal log-prices Y_A and Y_B:

$$\frac{e^{2Y_A}}{e^{Y_A}+e^{Y_B}} = e^{Y_A}\frac{1}{1+e^{Y_B-Y_A}} = e^{Y_A}g(W)\,,$$

where $W = Y_B - Y_A$ and $g(W) = (1+e^W)^{-1}$. Using the exact Fourier inversion derived in Theorem IA6, we substitute the integral representation for $g(W)$ with the complex variable $v = u - i\alpha$:

$$\begin{aligned}\mathbb{E}_t\left[e^{Y_A}g(W)\right] &= \mathbb{E}_t\left[e^{Y_A}\frac{1}{2\pi}\int_{-\infty}^{\infty}e^{-iv(Y_B-Y_A)}\frac{\pi}{\sin(\pi iv)}du\right]\\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty}\mathbb{E}_t\left[\exp\left((1+iv)Y_A - ivY_B\right)\right]\frac{\pi}{\sin(\pi iv)}du\,.\end{aligned}$$

We now map the exponents to the joint characteristic function $\Phi_t(z_1, z_2) = \mathbb{E}_t[\exp(iz_1Y_A + iz_2Y_B)]$: For Asset A, the exponent is $1 + iv = 1 + i(u - i\alpha) = 1 + iu + \alpha$. Setting $iz_1 = 1 + iu + \alpha$ yields $z_1 = u - i(1+\alpha)$. For Asset B, the exponent is $-iv = -i(u - i\alpha) = -iu - \alpha$. Setting $iz_2 = -iu - \alpha$ yields $z_2 = -u + i\alpha$. Substituting these arguments gives $\Phi_t(u - i(1+\alpha), -u + i\alpha)$, establishing equation (I.A.29) and completing the proof.

Last, we evaluate the risk-neutral covariance $\text{Cov}_t^*(R_A, R_M) = \mathbb{E}_t[M_T R_A R_M] - 1$. Substituting the logarithmic pricing kernel $M_T = p_t/p_T$ and the gross returns:

$$\mathbb{E}_t\left[\left(\frac{p_t}{p_T}\right)\left(\frac{p_{A,T}}{p_{A,t}}\right)\left(\frac{p_T}{p_t}\right)\right] - 1 = \mathbb{E}_t\left[\frac{p_{A,T}}{p_{A,t}}\right] - 1 = \frac{\mathbb{E}_t[p_{A,T}]}{p_{A,t}} - 1\,.$$

The term $\mathbb{E}_t[p_{A,T}]$ is the physical expectation of Asset A's terminal fundamental value, which isolates the idiosyncratic and co-jump components affecting A. Evaluating $\mathbb{E}_t[\exp(Y_A)]$ yields $e^{y_{A,t}}\mathscr{J}_A(\kappa_t)^{-(\beta_{A,t}+\beta_{AB,t})}$, completing all the proofs. □

Proof for Result IA7. Let the continuous-time interval dt be represented by a single period of length $\delta \to 0$ in the discrete $M = 4$ node tree. The four states are $j \in \{0, A, B, AB\}$, representing no jump, an idiosyncratic jump in A, an idiosyncratic jump in B, and a co-jump. The wealth-weighted physical probabilities of these states are $h_{00} = 1 - (\lambda_{A,t} + \lambda_{B,t} + \lambda_{AB,t})\delta$, $h_A = \lambda_{A,t}\delta$, $h_B = \lambda_{B,t}\delta$, and $h_{AB} = \lambda_{AB,t}\delta$. Let $p_t = p_{A,t} + p_{B,t}$ be the aggregate market value, and $w_{i,t} = p_{i,t}/p_t$. A jump in asset i reduces its value by the fraction $\Delta_i = 1 - e^{-\omega_i}$. The value of the aggregate market portfolio in the next period, $p_{t+1}(j)$, depends on the realized state: $p_{t+1}(0) = p_t$, $p_{t+1}(i) = p_t - p_{i,t}\Delta_i = p_t\mathscr{D}_{i,t}$, and $p_{t+1}(AB) = p_t - p_{A,t}\Delta_A - p_{B,t}\Delta_B = p_t(\mathscr{D}_{A,t} + \mathscr{D}_{B,t} - 1)$.

Using the generalized change of measure derived in equation (23) under logarithmic utility, the risk-neutral probability of state j is $h_j^* = h_j\frac{p_t}{p_{t+1}(j)}$. Substituting the state-dependent market prices yields (I.A.31).

The instantaneous risk-neutral variance of asset A's return is the sum of the squared jump sizes multiplied by their risk-neutral intensities: $\text{Var}_t^*(dR_A)/dt = \lambda_{A,t}^*\Delta_A^2 + \lambda_{AB,t}^*\Delta_A^2 = (\lambda_{A,t}^* + \lambda_{AB,t}^*)\Delta_A^2$. Covariance comes solely from co-jumps: $\text{Cov}_t^*(dR_A, dR_B)/dt = \lambda_{AB,t}^*\Delta_A\Delta_B$.

To prove the correlation amplification, define the physical correlation of the jump processes as $\rho_t = \frac{\lambda_{AB,t}}{\sqrt{(\lambda_{A,t}+\lambda_{AB,t})(\lambda_{B,t}+\lambda_{AB,t})}}$. Set $\mathscr{D}_{A,t} = 1 - w_{A,t}\Delta_A$, $\mathscr{D}_{B,t} = 1 - w_{B,t}\Delta_B$, and $\mathscr{D}_{AB,t} = 1 - w_{A,t}\Delta_A - w_{B,t}\Delta_B$. The risk-neutral correlation is $\rho_t^* = \frac{\lambda_{AB,t}^*}{\sqrt{(\lambda_{A,t}^*+\lambda_{AB,t}^*)(\lambda_{B,t}^*+\lambda_{AB,t}^*)}}$. Substituting the risk-neutral intensities gives:

$$\rho_t^* = \frac{\lambda_{AB,t}/\mathscr{D}_{AB,t}}{\sqrt{(\frac{\lambda_{A,t}}{\mathscr{D}_{A,t}} + \frac{\lambda_{AB,t}}{\mathscr{D}_{AB,t}})(\frac{\lambda_{B,t}}{\mathscr{D}_{B,t}} + \frac{\lambda_{AB,t}}{\mathscr{D}_{AB,t}})}} = \frac{\lambda_{AB,t}}{\sqrt{(\lambda_{A,t}\frac{\mathscr{D}_{AB,t}}{\mathscr{D}_{A,t}} + \lambda_{AB,t})(\lambda_{B,t}\frac{\mathscr{D}_{AB,t}}{\mathscr{D}_{B,t}} + \lambda_{AB,t})}}\,.$$

Multiplying and dividing by the physical denominator $\sqrt{(\lambda_{A,t} + \lambda_{AB,t})(\lambda_{B,t} + \lambda_{AB,t})}$ im-

mediately yields the multiplier structure in equation (I.A.33). Since $\mathscr{D}_{AB,t} < \mathscr{D}_{A,t}, \mathscr{D}_{B,t}$, the ratios $\frac{\mathscr{D}_{AB,t}}{\mathscr{D}_{A,t}}$ and $\frac{\mathscr{D}_{AB,t}}{\mathscr{D}_{B,t}}$ are strictly less than 1. This strictly reduces the denominator of ρ_t^* relative to ρ_t, proving that $|\rho_t^*| > |\rho_t|$. □

I.A.7 Speculation on autocorrelation

An equivalent formulation of our trinomial tree, using a periodic recombining binary tree in which agents' beliefs about the binary distribution of news switch periodically, translates heterogeneous beliefs about volatility into heterogeneous views on return autocorrelation. To be precise, we consider the following parameterization of subjective beliefs on the trinomial tree (g, m, b) using (h, ε):

$$m = 2h(1-h) - \varepsilon\,, \quad g = h(h+\varepsilon)\,, \quad b = (1-h)(1-h+\varepsilon)\,, \tag{I.A.37}$$

where $h \in (0,1)$ and ε is a small number around zero. The technical condition for (m, g, b) to be a probability measure is $\max(-h, h-1) \leq \varepsilon \leq \min(h, 1-h, 2h(1-h))$. The more extreme h is, the narrower the range of ε is around 0.

Although equivalent to some of the trinomial tree settings, the recombining binary tree model has a rather different interpretation: h captures agents' overall bullishness about the market going up; ε captures agents' beliefs about the inter-temporal correlation of news between two sub-periods ($t \to t+0.5$ and $t+0.5 \to t+1$).

The expression of m from (I.A.37) indicates two channels that can drive volatility disagreement. The first channel is through disagreement in ε: agents who believe in a positive correlation between sub-periods perceive a higher volatility of fundamentals between t and $t+1$. Figure I.A.8 shows a numerical example with three periods where $h = 0.5$ across all agents and $\varepsilon \in [-1/2, 1/2]$ follows a uniform distribution. Under this setting, individual agents' beliefs in (I.A.37) become

$$m = 1/2 - \varepsilon, \quad g = 1/4 + 1/2\,\varepsilon, \quad b = 1/4 + 1/2\,\varepsilon\,.$$

Let $g = v/2$, $m = 1-v$, and $b = v/2$ by replacing $v = 1/2 + \varepsilon$. This is equivalent to a trinomial tree model when news content disagreement is absent and $u = 0.5$ for all agents.

The red path's sub-period history on the binary tree is (up, down, up, down) and

the blue path history is (down, down, up, up). The disagreement regarding ε breaks the binary tree model's permutation symmetry and results in different pricing and sentiment regarding volatility. On the red path, those who believe in a negative ε become hindsight winners, and Mr. Market's belief of ε is negative at $t = 2$. By contrast, the blue path yields positive sentiment $\mathcal{E}$ because of positive news correlations between periods 0–0.5 and 0.5–1, and between 1–1.5 and 1.5–2. The lower $\mathcal{E}$ (recall that $\mathcal{E} = V - 0.5$) means a lower sentiment of news intensity, which pushes prices higher: 'no news is good news'.

The second channel is through disagreement in h; agents with extreme bullishness (in terms of h) or bearishness (in terms of $1 - h$) think the market is more volatile across multiple sub-periods. In Figure I.A.9, we show another example with $\varepsilon = 0$ and h varying across agents with a uniform distribution. Under this setting, (I.A.37) becomes

$$m = 2h(1-h), \quad g = h^2, \quad b = (1-h)^2,$$

and this is essentially a recombining binary tree where agents disagree about the probability h of an up-move. The pricing of the market asset in this special case is equivalent to the model in Martin and Papadimitriou (2022) upon recombining odd and even periods. The permutation symmetry of the binary tree is preserved, and pricing is not different after experiencing the red or blue news sample paths.

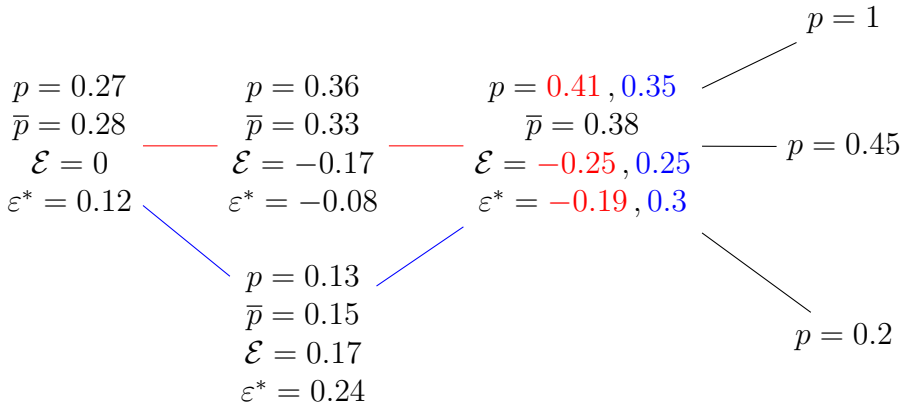


Figure I.A.8: **A numerical example of ε-disagreement.** Agents' beliefs ε are uniformly distributed and they agree on $h = 1/2$. We only report pricing at $t = 0, 1, 2, 3$.

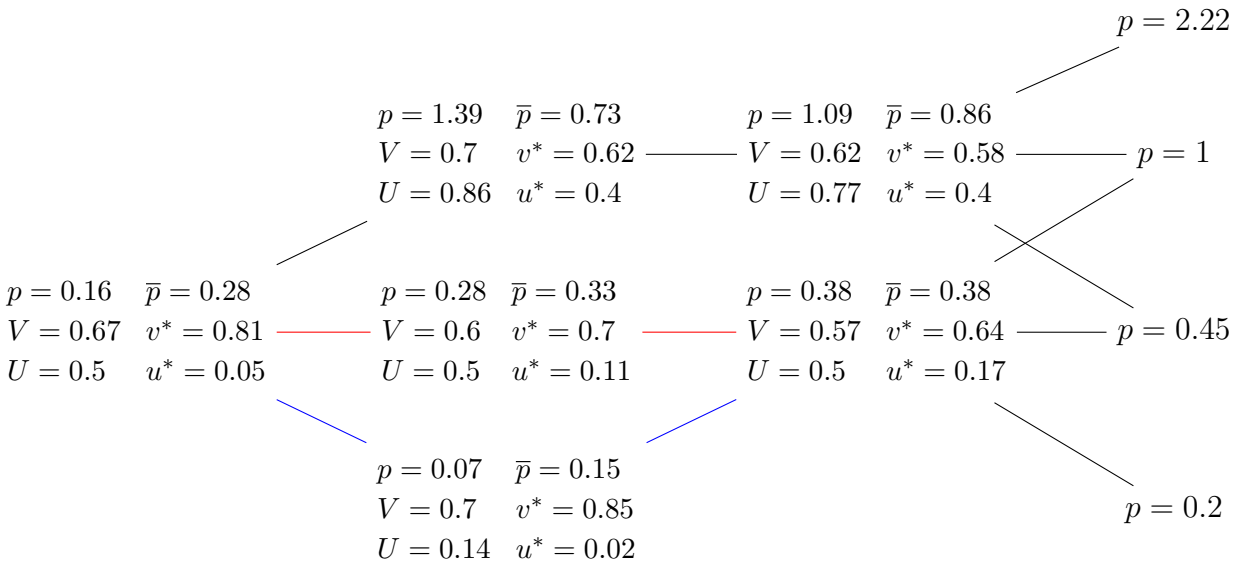


Figure I.A.9: **Recombining binary tree: h-disagreement.** This example uses the binary tree from Martin and Papadimitriou (2022) where h is uniformly distributed. We only report pricing at $t = 0, 1, 2, 3$ (while binary tree also have nodes at $t = 0.5, 1.5, 2.5$). p is the market price, $\overline{p}$ is the price in the homogeneous economy where everyone agrees with the median agent $h = 1/2$. (v, u) represents Mr. Market's beliefs and (v^*, u^*) is the risk-neutral measure.

References

Martin, Ian, and Dimitris Papadimitriou. 2022. "Sentiment and speculation in a market with heterogeneous beliefs." *American Economic Review*, 112(8): 2465–2517.

Martin, Ian W. R., and Christian Wagner. 2019. "What is the Expected Return on a Stock?" *Journal of Finance*, 74(4): 1887–1929.

Negishi, Takashi. 1960. "Welfare economics and existence of an equilibrium for a competitive economy." *Metroeconomica*, 12(2-3): 92–97.

Paul, S. R., and R. L. Plackett. 1978. "Inference Sensitivity for Poisson Mixtures." *Biometrika*, 65(3): 591–602.