

# Market panics and their effect on those who leave and those who stay

Scott Condie\*

## ABSTRACT

This paper studies the welfare and asset pricing consequences of market panics driven by ambiguity aversion. When information is perceived to be highly ambiguous, ambiguity-averse traders exit the market entirely, raising the risk premium and creating opposing effects for expected-utility traders who remain: a wealth loss from falling prices and an expanded trading opportunity from reduced risk sharing. This paper characterizes the size of these effects and their impact on welfare. In a two-period extension the premium associated with the possibility of a future panic is calculated. Even when all fundamental uncertainty is normal, the distribution of interim prices is left-skewed and fat-tailed. Price variance decomposes into a fundamental component and an ambiguity-driven component, providing structural microfoundations for well-documented non-normalities in equity returns.

Keywords: ambiguity, welfare, entry

JEL classification: D81, G11, G12, G14

---

\*Department of Economics, Brigham Young University. Email: [ssc@byu.edu](mailto:ssc@byu.edu). Telephone: +1.801.422.5306.

# 1 Introduction

Investors speak colloquially of traders getting “spooked” when there are sudden increases in market uncertainty. Spooked traders tend to sell off risky asset holdings quickly and move to assets that they perceive to be safer. For those who are contemplating staying in the market, the existence of potentially spooked traders poses risks and opportunities. One risk comes from the fact that if many traders get spooked, the sale of their risky asset holdings will drive prices down for those who remain in the market, lowering the wealth of these individuals. The opportunity comes from the reduced risk-sharing from fewer traders in the market, which increases the premium associated with holding market risk for those who remain.

This paper studies this phenomenon by building a model where some individuals in the market experience market panic, a phenomenon that is modelled as a sudden increase in the perceived ambiguity around signals about future asset payoffs combined with an aversion to this ambiguity. It has been shown in other work that ambiguity aversion can lead to individuals choosing to not participate in markets when they perceive signals about future payoffs to be ambiguous (Epstein & Wang 1994, Dow & Werlang 1992, Mukerji & Tallon 2001, Easley & O’Hara 2009). In a two-period extension of the model, we show that the prospect of a future panic depresses the price of the risky asset by a panic premium that depends on the fraction of traders susceptible to panic and the expected severity of the ambiguity shock.

By deriving the distribution of time-1 prices in the model it is shown that whenever panic is possible, the distribution of prices is left-skewed and leptokurtic relative to a Gaussian with the same mean and variance. Both features arise endogenously from the equilibrium point at which individuals panic and exit the market. A variance decomposition shows that total price variance equals fundamental uncertainty plus an ambiguity-driven component that vanishes when ambiguity is absent and grows with the severity of potential panics.

Prior to the panic, the premium embedded in the price of holding risk implies that a portion of today’s expected return compensates specifically for the risk that ambiguity-averse traders will exit in the future, and that this component disappears if a panic does not materialize. Investors seeking to decompose their expected returns into fundamental and panic-related components, or to hedge the tail risk associated with a potential panic, can use the variance decomposition derived here to quantify

the ambiguity-driven share of price variance. The finding that the interim price distribution is left-skewed and fat-tailed—even when the underlying asset payoff is symmetric and normally distributed—has direct implications for risk management: standard tools based on normality will systematically understate the probability of large leftward price movements in markets where ambiguity-driven exit is possible.

The model also clarifies one mechanism that makes a market structurally prone to panic. The two ingredients developed in this paper include the distribution of ambient uncertainty and the fraction of potentially panicking investors. In this paper, the probability that a panic occurs equals the probability that the ambiguity shock exceeds the equilibrium panic threshold, so markets where uncertainty is more dispersed—where large ambiguity realizations are more likely—are more prone to panics holding market composition fixed. Equally important is the share of ambiguity-averse traders: a larger fraction of potentially panicking investors lowers the equilibrium threshold at which a panic is triggered, amplifying panic risk from both directions at once. These two determinants have different policy implications. For market regulators, the distribution of uncertainty is the more tractable target: policies that narrow the range of plausible interpretations of public signals—clearer disclosure standards, more timely and standardized reporting, or reduced information asymmetry between firms and investors—compress the distribution of perceived ambiguity and shift the panic threshold outward, making the market less vulnerable. Because the threshold is endogenous, even modest reductions in ambient uncertainty can have outsized effects on the panic probability when the current threshold sits in a region where the ambiguity distribution places substantial probability mass. For firms seeking equity finance, the implications are more immediate: a firm whose assets are difficult to value, whose business model is opaque, or whose cash flows depend on hard-to-interpret contingencies will attract a higher share of ambiguity-averse investors and generate a larger perceived ambiguity shock, both of which raise the panic premium embedded in its cost of equity. Firms can reduce this premium by taking actions that lower perceived ambiguity—voluntary disclosure, analyst coverage, simplified capital structures, or index inclusion—since these alter the composition of the investor base and raise the panic threshold to a level where panic is less likely to occur.

## Related Literature

This paper builds on the maxmin expected utility (MMEU) model of Gilboa & Schmeidler (1989), in which agents evaluate uncertain acts by the minimum expected utility over a set of priors. A central implication of MMEU preferences, established by Dow & Werlang (1992), is portfolio inertia: an ambiguity-averse agent may optimally hold zero position in a risky asset across a range of prices, because no single prior supports deviating from this position. This paper exploits that feature to generate endogenous market exit—“panic”—by a subset of traders. A related result appears in Mukerji & Tallon (2001), who show that ambiguity aversion by any fraction of traders can cause certain asset markets to collapse entirely, since ambiguity-averse agents refuse to participate; our setting differs in that we take the presence of both trader types as given and ask how the prospect of exit by one group reshapes incentives for the other.

The papers most directly related to ours are Easley & O’Hara (2009) and Easley & O’Hara (2010). Easley & O’Hara (2009) study a setting in which ambiguity about asset payoffs causes some investors to withdraw from the market and examine how regulatory disclosure policy can restore participation. Our paper shares their central premise—that ambiguity drives nonparticipation—but studies ambiguous signals and shifts the focus from the regulator’s problem to the welfare of participants who remain after others exit, the empirical implications of this type of equilibrium including the price distribution and volume. Easley & O’Hara (2010) embed ambiguity into a standard market microstructure framework with informed and uninformed traders, showing that ambiguity about the information content of order flow generates equilibrium bid-ask spreads that pure expected utility models cannot produce. The parallel to our setting is their treatment of traders who perceive signal quality as ambiguous; our model follows a similar approach by representing ambiguity as uncertainty over the mean of signal noise, but asks how this affects the endogenous composition of the market and the welfare of those who choose to remain.

A closely related strand of the literature studies how ambiguity about the quality or precision of public information affects equilibrium asset prices. Epstein & Schneider (2008) show, in a dynamic rational expectations model, that ambiguity about signal precision amplifies the price impact of news and generates excess volatility relative to the standard Bayesian benchmark. Their focus is on asset pricing moments; ours is on participation and welfare. This paper also connects to work on markets in which

ambiguity-averse and expected utility traders coexist. Condie & Ganguli (2011) show that when a subset of traders is ambiguity averse, the resulting equilibria may fail to aggregate private information, so that asset prices become less informative than in a standard rational expectations equilibrium. Our model takes a simpler, closed-form approach—closer in spirit to a CAPM-style setting—to make welfare comparisons tractable, and focuses on the two-way feedback between market composition and the incentives of expected utility traders to remain in or exit the market.

The present paper also shares important structural features with Condie & Ganguli (2017), who study how ambiguity about private signals affects equilibrium asset prices. As in the present model, ambiguity-averse traders in Condie & Ganguli (2017) perceive the noise in their signal as drawn from a set of distributions with means in  $[-\delta, \delta]$ , which under MMEU preferences generates the same symmetric interval of posterior means and the same pattern of portfolio inertia. The two papers diverge in focus and in the question they address. Condie & Ganguli (2017) examine a rational expectations setting in which the private signal carries information about asset value, and the central question is how ambiguity about signal quality distorts the degree to which that private information is revealed in prices; the main finding is that ambiguity can cause prices to imperfectly aggregate information in ways that differ qualitatively from both the fully-revealing and no-trade benchmarks. The present paper instead uses the same signal-ambiguity structure with a public signal in a setting designed to make the exit decision and its welfare consequences tractable, and asks how the anticipation of ambiguity-driven exit shapes the current price and the distribution of future prices. The two papers thus illustrate, from complementary vantages, how the same representation of signal ambiguity generates qualitatively different pricing phenomena depending on whether the focus is on information aggregation or on market participation and welfare.

The present paper is also closely related to Illeditsch (2011) and Illeditsch, Ganguli & Condie (2021), both of which share the same MMEU signal-ambiguity structure and study its consequences for equilibrium asset markets. Illeditsch (2011) embeds ambiguous public announcements in a market with MMEU traders and shows that uncertainty over the information content of the signal—represented as an interval of possible signal means, exactly as in the present model—generates portfolio inertia around announcement dates and excess return volatility relative to the Bayesian benchmark. The mechanism is the same one that underlies panic in our model:

because no single prior in the trader’s set supports deviating from the current position, the trader optimally holds a fixed position in the asset. The key distinction is the nature of the response and its market-level consequence. In Illeditsch (2011) the response is inertia—traders do not adjust their portfolios in response to the ambiguous signal—whereas in the present paper the response is exit: ambiguity-averse traders withdraw entirely from the market when the realized ambiguity shock is sufficiently large. Exit is a more extreme form of the same inertia logic, and it generates consequences—a shift in market composition, a welfare transfer to remaining traders, and a discrete price decline—that inertia alone does not produce.

Illeditsch et al. (2021) study how ambiguity about private signal quality generates information inertia: prices adjust sluggishly to private information because ambiguity-averse traders, perceiving uncertainty over what their signal implies, do not immediately trade on it. As in Condie & Ganguli (2017) and the present paper, the ambiguity is represented as uncertainty over the noise mean of a private signal, yielding the same interval of posterior means under MMEU. The central question in Illeditsch et al. (2021) is dynamic information aggregation—how quickly, and how fully, private information is impounded into prices when some traders are ambiguity averse—whereas the present paper holds the information structure fixed and asks how the prospect of ambiguity-driven exit shapes the current price level and the distribution of future prices. Taken together, Illeditsch (2011), Illeditsch et al. (2021), Condie & Ganguli (2017), and the present paper form a coherent body of work showing how the same representation of signal ambiguity generates qualitatively distinct phenomena—excess volatility, information inertia, imperfect price revelation, and market panics—depending on whether ambiguity bears on public or private signals and whether the equilibrium response is inertia, gradual adjustment, or full exit.

Our finding that the prospect of a future panic depresses the current asset price is closely related to the “peso problem” identified by Rietz (1988). Rietz showed that a small probability of a rare economic disaster can generate a large equity risk premium, suggesting a potential resolution to the puzzle of Mehra & Prescott (1985), even if the disaster never materializes in the observed sample. Barro (2006) revived and extended this idea by calibrating the disaster probability and magnitude to twentieth-century macroeconomic catastrophes, showing that the model can quantitatively match the equity premium and the low risk-free rate. Subsequent work by Gabaix (2012) and Wachter (2013) showed that allowing the probability of a rare disaster to vary over

time can explain a wide range of additional asset pricing phenomena, including excess volatility and return predictability. Our panic premium is structurally analogous: the possibility that ambiguity-averse traders will exit in the future acts as a rare-disaster-like event whose anticipated occurrence is priced into the current asset value. The key difference is that in our model the “disaster” arises endogenously from the interaction between ambiguity aversion and market structure, rather than from an exogenous consumption shock, and its severity depends on the fraction of traders susceptible to panic and the distribution of the ambiguity shock.

The distributional results of this paper—left skewness and excess kurtosis of time-1 prices—connect to the empirical literature on return non-normality and to theoretical work on ambiguity-driven asset pricing moments. Empirically, Bollerslev & Todorov (2011) show that equity tail risk is separately priced and that the variance risk premium is almost entirely attributable to compensation for jump tail events; our endogenous panic event provides a structural origin for precisely such a tail. Carr & Wu (2009) and Bollerslev, Tauchen & Zhou (2009) document that the variance risk premium—the difference between risk-neutral and physical return variance—is negative and time-varying, implying that investors pay to hedge variance uncertainty above and beyond fundamental payoff risk. Our variance decomposition identifies the corresponding theoretical quantity: the ambiguity component  $x_a^2 \text{Var}(\min(\Delta, y(x_e)))$  is the portion of price variance generated by the ambiguity shock rather than by fundamental uncertainty, and it is precisely the variance over which expected-utility traders cannot insure. On the theoretical side, Miao, Wei & Zhou (2019) show that ambiguity aversion can explain the negative variance risk premium through pessimistic probability distortions; our model complements their result by providing a quantity-theoretic decomposition in which the ambiguity-driven excess variance is traced explicitly to the distribution of the censored ambiguity shock. Both the empirical and theoretical connections reinforce the rare-disasters parallel: the panic event contributes to expected returns through the time-0 panic premium and to return non-normality through the left-skewed, leptokurtic distribution of time-1 prices.

This paper begins by developing an equilibrium model of market panics due to ambiguity in section 2. The equilibrium and welfare implications of that model are derived in section 3. Section 4 then adds one pre-period to the model where the realization of the amount of ambiguity is unknown and the premium available to market participants as well as the distributional properties of market prices are derived. Sec-

tion 5 concludes.

## 2 Model

Consider a market operating at time 1, the uncertainty of which will be resolved at time 2. The assets in the market are composed of one risk-free asset and one risky asset. The risk-free asset pays a normalized net return of 0 at time 2, while the future value of the risky asset is given by  $v \sim N(\mu_0, \rho_0^{-1})$ , where  $\rho_0$  is the ex ante precision of the risky asset. There is a unit mass of traders in the market who each have initial endowments of one unit of the risky asset and no units of the risk-free asset. Of these traders, fraction  $x_e$  are expected utility traders, while fraction  $x_a = 1 - x_e$  perceive ambiguity in the signal about asset values in a way that will be discussed. Traders can trade at time 1 based on signals they receive about the future value of the risky asset. The signal received by traders is given by  $s = v + \epsilon$  where  $\epsilon \sim N(0, \rho_\epsilon^{-1})$ .

Bayesian updating of the signal for an individual who does not perceive ambiguity in the signal leads to a posterior distribution of the future value of the asset of  $v|s \sim N(m(s), \rho_I^{-1})$  where  $\rho_I = \rho_0 + \rho_\epsilon$  and  $m(s) = \mu_0 + \frac{\rho_\epsilon}{\rho_I}(s - \mu_0)$ . Some individuals perceive the information content of the signal to be ambiguous. Specifically, these individuals believe that the noise term  $\epsilon$  is drawn from a set of normal distributions with means  $\mu \in [-\delta, \delta]$  and precision  $\rho_\epsilon$ . Updating the prior given a signal with each of these possible distributions leads to a set of posterior distributions for the future value of the asset given by  $v|s \sim N(m_z(s), \rho_I^{-1})$  where  $m_z(s) = m(s) - \frac{\rho_\epsilon}{\rho_I}z$  for  $z \in [-\delta, \delta]$ . This is the set of normal distributions with means in the interval  $[\mu_0 + \frac{\rho_\epsilon}{\rho_I}(s - \mu_0 - \delta), \mu_0 + \frac{\rho_\epsilon}{\rho_I}(s - \mu_0 + \delta)]$  and precision  $\rho_I$ . We define  $\Delta = \delta \frac{\rho_\epsilon}{\rho_I}$  and write this set of updated means as  $[m(s) - \Delta, m(s) + \Delta]$ .

For all traders in the market, terminal wealth is given by  $w = b + \theta v$  where  $b$  is the number of units of the risk-free asset held and  $\theta$  is the number of units of the risky asset held. Given a trader's initial endowment of one unit of the risky asset and no units of the risk-free asset, the budget constraint is given by

$$b + p\theta = p \tag{1}$$

Plugging this into terminal wealth shows that terminal wealth is the random variable

given by

$$w = p + \theta(v - p). \quad (2)$$

Individuals who either do not perceive ambiguity in the signal or are not averse to ambiguity have mean-variance expected utility over terminal wealth given by

$$U = Ew - \frac{\gamma}{2}\sigma_w^2. \quad (3)$$

Plugging in terminal wealth and recognizing that given a signal  $s$ ,  $E(v|s) = m(s)$  and  $\sigma_I^2 = \frac{1}{\rho_I}$  are the conditional mean and variance of the future value of the asset, respectively, gives

$$U = p(s) + \theta(m(s) - p(s)) - \frac{\gamma}{2}\theta^2\sigma_I^2. \quad (4)$$

Given a price  $p(s)$ , the optimal portfolio choice for these expected utility (EU) traders is

$$\theta_e = \gamma^{-1}\rho_I[m(s) - p(s)] \quad (5)$$

and the indirect utility is

$$U(p, s) = p(s) + \frac{1}{2}\gamma^{-1}\rho_I[m(s) - p(s)]^2. \quad (6)$$

Taken together, expected utility traders have the fraction  $x_e$  of the total initial endowment of the risky asset.

As a benchmark, consider the case where all traders in the market are expected utility traders with homogeneous beliefs. In this case the market clearing condition would be

$$\theta_e(p(s)) = 1. \quad (7)$$

Given the optimal portfolio choice for expected utility traders, this implies that the equilibrium price in this case would be

$$p(s) = m(s) - \gamma\sigma_I^2. \quad (8)$$

For use throughout the rest of the paper, we define

$$\Lambda = \gamma\sigma_I^2 \quad (9)$$

to be the baseline risk premium in this market. The quantity  $\Lambda$  is the amount of risk perceived in the market, normalized by risk aversion, and expressed in terms of expected future wealth. This quantity plays an important role in both the decisions of traders in the market and in equilibrium outcomes. We will often refer to this quantity as simply *market risk*.

Put in terms of market risk, demand for EU traders is

$$\theta_e = \Lambda^{-1}[m(s) - p(s)] \quad (10)$$

and indirect utility is

$$U(p, s) = p(s) + \frac{1}{2}\Lambda^{-1}[m(s) - p(s)]^2. \quad (11)$$

For the fraction  $x_a$  of traders who are ambiguity averse (AA), utility is

$$U = \min_{z \in [-\delta, \delta]} \left\{ E[w|s, z] - \frac{\gamma}{2}\sigma_w^2 \right\} \quad (12)$$

Plugging in terminal wealth to this utility function leads to the objective function

$$U = p(s) + \min_{z \in [-\delta, \delta]} \left\{ \theta(m_z(s) - p(s)) - \frac{\gamma}{2}\theta^2\sigma_I^2 \right\} \quad (13)$$

where  $m_z(s) = m(s) - \frac{\rho_\epsilon}{\rho_I}z$  as defined above.

Demand for AA traders is

$$\theta_a = \begin{cases} \Lambda^{-1}(m(s) - \Delta - p(s)) & m(s) - \Delta - p(s) > 0 \\ 0 & m(s) - \Delta - p(s) \leq 0 \leq m(s) + \Delta - p(s) \\ \Lambda^{-1}(m(s) + \Delta - p(s)) & m(s) + \Delta - p(s) \leq 0 \end{cases} \quad (14)$$

This demand function reflects the fact that if AA traders are going to hold a long position in the market, the expected payoff to holding the risky asset must be positive under the worst-case scenario belief. The ambiguity inherent in the market signal for these traders is larger enough that they are only willing to hold the risky asset with the additional *ambiguity premium*  $\Delta$ .

Indirect utility for these AA investors is

$$U(p, s) = \begin{cases} p(s) + \frac{1}{2}\Lambda^{-1}(m(s) - \Delta - p(s))^2 & m(s) - \Delta - p(s) > 0 \\ p(s) & m(s) - \Delta - p(s) \leq 0 \leq m(s) + \Delta - p(s) \\ p(s) + \frac{1}{2}\Lambda^{-1}(m(s) + \Delta - p(s))^2 & m(s) + \Delta - p(s) > 0 \end{cases} . \quad (15)$$

Ambiguity averse traders hold the fraction  $x_a = 1 - x_e$  of the total initial endowment of the risky asset.

The market clearing condition in this market is that total demand by expected utility traders and ambiguity averse traders must equal the total supply of the asset.

$$x_e\theta_e(p(s)) + x_a\theta_a(p(s)) = 1. \quad (16)$$

### 3 Equilibrium

There are two classes of equilibria to consider in this model. In the first type, ambiguity averse traders feel anxious enough about the market signal that has been received that they choose to exit the market. That is, it is optimal for them to hold a zero position. In the second type, ambiguity averse traders do not choose to exit the market and sell off their initial endowment of the risky asset to expected utility traders. We will derive equilibrium under each of these two possibilities and use these derivations to understand conditions under which each type of equilibrium occurs. Then we will study the welfare implications of each type of equilibrium for both types of traders.

Under an equilibrium where ambiguity averse investors are spooked and exit the market, the market clearing condition becomes

$$x_e\theta_e(p(s)) = 1. \quad (17)$$

Given the optimal portfolio choice for expected utility traders, this implies that the equilibrium price in this case is

$$p(s) = m(s) - \frac{\Lambda}{x_e}. \quad (18)$$

This equilibrium price has the implicit risk premium

$$y(x_e) = \frac{\Lambda}{x_e}. \quad (19)$$

The traditional risk premium  $\Lambda = \gamma\sigma_I^2$  is scaled up by the inverse of the fraction of expected utility traders in the market. If this fraction is small, the risk premium is large because there is less risk sharing occurring in the market. To derive conditions under which this is an equilibrium, we need to check that ambiguity averse traders choose to hold a zero position in the risky asset at this price. This requires that

$$\begin{aligned} m(s) - \Delta &\leq p(s) \leq m(s) + \Delta \\ m(s) - \Delta &\leq m(s) - y(x_e) \leq m(s) + \Delta \\ -\Delta &\leq -y(x_e) \leq \Delta. \end{aligned} \quad (20)$$

Since  $y(x_e) > 0$ , this condition reduces to  $y(x_e) \leq \Delta$ . That is, the equilibrium where ambiguity averse traders exit the market occurs when the risk premium in the market is less than the ambiguity premium required by ambiguity averse traders to stay in the market.

For an equilibrium where both types of traders remain in the market holding non-zero positions, the market clearing condition is

$$x_e\theta_e(p(s)) + x_a\theta_a(p(s)) = 1. \quad (21)$$

Since AA traders' range of potential beliefs is centered around  $m(s)$ , the beliefs of EU traders, it can be verified that there will never be an equilibrium where AA traders take a short position in the risky asset while EU traders take a long position, or where EU traders take a short position and AA traders take a long position. Therefore, if this equilibrium exists, both types of traders take long positions in the risky asset and the market clearing condition becomes

$$\Lambda^{-1}(m(s) - p(s))x_e + \Lambda^{-1}(m(s) - \Delta - p(s))x_a = 1 \quad (22)$$

This implies that the equilibrium price in this case is

$$\begin{aligned} p(s) &= m(s) - \Delta x_a - \frac{\Lambda}{x_e} x_e \\ &= m(s) - \Delta x_a - y(x_e) x_e. \end{aligned} \tag{23}$$

This equilibrium price is a wealth-weighted average of the ambiguity premium  $\Delta$  and the risk premium  $y(x_e)$ . This will be an equilibrium if AA traders choose to hold a positive position in the risky asset at this price. This requires that

$$\begin{aligned} m(s) - \Delta &> p(s) \\ m(s) - \Delta &> m(s) - \Delta x_a - y(x_e) x_e \\ y(x_e) x_e &> \Delta(1 - x_a) \\ y(x_e) &> \Delta. \end{aligned} \tag{24}$$

Therefore, an equilibrium where both types of traders remain in the market occurs when the risk premium in the market is greater than the ambiguity premium required by ambiguity averse traders to stay in the market.

**Result 1.** *In equilibrium, ambiguity averse traders will exit the market if and only if the risk premium when they are absent is less than the ambiguity premium they require to stay in the market. This will happen when*

- *The level of perceived ambiguity  $\Delta$  is high,*
- *Market risk  $\Lambda$  is low, and/or*
- *The fraction of expected utility traders in the market  $x_e$  is high.*

Intuitively, when a small number of traders perceive great ambiguity in a precise signal, these traders will be spooked and exit the market, selling off their risky asset holdings to those who remain in the market.

### 3.1 Welfare for those who stay and those who leave

Given the equilibrium conditions derived above, we can compute welfare for both types of traders under the equilibrium where some investors become spooked and

exit the market and under the equilibrium where all investors remain in the market. When AA traders become spooked, utility for EU traders is

$$\begin{aligned}
U_e^*(s, x_e) &= m(s) - y(x_e) + \frac{1}{2}\Lambda^{-1}y(x_e)^2 \\
&= m(s) - \frac{\Lambda}{x_e} + \frac{1}{2}\Lambda^{-1}\left(\frac{\Lambda}{x_e}\right)^2 \\
&= m(s) - \frac{\Lambda}{x_e} + \frac{1}{2}\frac{\Lambda}{x_e^2}
\end{aligned} \tag{25}$$

Throughout the paper, variables with a  $\star$  superscript will refer to quantities calculated in the panic equilibrium and variables with a  $\circ$  superscript refer to quantities calculated in the non-panic equilibrium.

We see that the risk premium  $y(x_e)$  has two effects on welfare. First, increases in  $y(x_e)$  decrease the value of the trader's endowment (or previous position), which lowers utility. That is, it has a negative wealth effect on trader welfare. Additionally, deviations of the premium away from zero increase the trading opportunities available in the market in a quadratic way. That is, as the premium  $y(x_e)$  increases, the trading effect is positive. In other words, the market offers greater opportunity to the trader the more  $y(x_e)$  deviates from zero.

Utility for AA traders in the panic equilibrium is

$$U_a^*(s, x_e) = p(s) = m(s) - y(x_e). \tag{26}$$

Their utility is the market value of their initial endowment of the risky asset. They do not receive any trading benefit because they have exited the market.

In the equilibrium where both types of traders remain in the market, utility for EU traders is

$$\begin{aligned}
U_e^\circ(s, x_a, x_e) &= m(s) - \Delta x_a - y(x_e)x_e + \frac{1}{2}\Lambda^{-1}(\Delta x_a + y(x_e)x_e)^2 \\
&= m(s) - \Delta x_a - y(x_e)x_e + \frac{1}{2}\Lambda^{-1}(\Delta x_a + \Lambda)^2
\end{aligned} \tag{27}$$

and for AA traders is

$$\begin{aligned}
U_a^\circ(s, x_a, x_e) &= m(s) - \Delta x_a - y(x_e)x_e + \frac{1}{2}\Lambda^{-1}(y(x_e)x_e - \Delta x_e)^2 \\
&= m(s) - \frac{\Lambda}{2} - \Delta + \frac{1}{2\Lambda}\Delta^2 x_e^2
\end{aligned} \tag{28}$$

### 3.1.1 Welfare comparison for EU traders under differing ambiguity levels

Here we compare the welfare of EU traders across two equilibria that differ only in the realized ambiguity level. In both equilibria EU traders are in the market — they must be, since they are the residual holders of the risky asset regardless of whether AA traders exit. The comparison therefore asks not whether EU traders prefer to stay, but rather which equilibrium outcome is more favorable to them. Consider two levels of perceived ambiguity, parameterized by  $\Delta_1$  and  $\Delta_2$  where  $\Delta_2 > \Delta_1$ . In these two worlds, all other parameters are the same; the level of ambiguity under  $\Delta_2$  is sufficiently high that AA traders exit the market, while under  $\Delta_1$  they remain. The difference in welfare for EU traders between these two equilibria is

$$\begin{aligned}
&U_e(s, x_e; \Delta_2) - U_e(s, x_a, x_e; \Delta_1) \\
&= \left[ m(s) - y(x_e) + \frac{1}{2}\Lambda^{-1}y(x_e)^2 \right] - \left[ m(s) - \Delta_1 x_a - y(x_e)x_e + \frac{1}{2}\Lambda^{-1}(\Delta_1 x_a + y(x_e)x_e)^2 \right] \\
&= (\Delta_1 - y(x_e))x_a + \frac{1}{2}\Lambda^{-1} [y(x_e)^2 - (\Delta_1 x_a + y(x_e)x_e)^2]
\end{aligned} \tag{29}$$

Recall that the condition for AA traders to exit the market under  $\Delta_2$  and remain under  $\Delta_1$  is that  $\Delta_1 \leq y(x_e) \leq \Delta_2$ . Given that, we can decompose this difference in welfare between the two regimes into two components. The first component is the wealth effect given by

$$dW = (\Delta_1 - y(x_e))x_a < 0. \tag{30}$$

This effect is negative because when AA traders exit the market, the market value of their asset holdings drops by  $y(x_e)$  which is greater than the value  $\Delta_1 x_a + y(x_e)x_e$  that it would have dropped if AA traders had stayed in the market. Therefore, EU traders lose wealth when AA traders exit the market relative to when they stay in the market.

The second component is the trading effect given by

$$dT = \frac{1}{2}\Lambda^{-1} [y(x_e)^2 - (\Delta_1 x_a + y(x_e)x_e)^2]. \quad (31)$$

This effect is positive because the risk premium  $y(x_e)$  is further from zero when AA traders exit the market than when they stay in the market. Therefore, EU traders earn a larger premium for staying in the market when AA traders exit. To determine which effect dominates, factor (29) using  $y(x_e)x_e = \Lambda$  and  $y(x_e) - \Lambda = y(x_e)x_a$ :

$$U_e(s, x_e; \Delta_2) - U_e(s, x_a, x_e; \Delta_1) = x_a(y(x_e) - \Delta_1) \cdot \frac{y(x_e) + \Delta_1 x_a - \Lambda}{2\Lambda}. \quad (32)$$

Each factor is non-negative in the no panic regime:  $x_a \geq 0$ ;  $y(x_e) \geq \Delta_1$  (the condition for AA traders to remain); and  $y(x_e) + \Delta_1 x_a \geq y(x_e) = \Lambda/x_e > \Lambda$  since  $x_e < 1$ . Therefore the difference is always non-negative: the panic equilibrium is always weakly better for EU traders than the no panic equilibrium, for any  $\Delta_1 \leq y(x_e)$ , with equality only at the boundary  $\Delta_1 = y(x_e)$ . Since EU traders are in the market in both equilibria, this is a statement about which equilibrium outcome they benefit from more, not about an exit decision on their part. The trading effect always dominates the wealth effect: EU traders gain more from the higher risk premium that arises when AA traders leave than they lose from the larger price drop on their initial endowment. Put another way, expected utility traders who can withstand the the initial wealth effect of the panic will find themselves happier on average if they stay in the market.

### 3.1.2 Welfare comparison under differing levels of wealth at risk of flight

In this section we study how changes in the level of wealth that is at risk of getting spooked affects the welfare of traders who stay and those who go. We analyze the impact of varying initial endowments and their implications for market participation and utility outcomes.

Ambiguity averse traders choose to leave the market when the risk premium in the market is less than their ambiguity premium, i.e., when  $y(x_e) \leq \Delta$ . Given  $\Delta = \frac{\rho_\epsilon}{\rho_I} \delta$  and  $y(x_e) = \frac{\gamma \sigma_I^2}{x_e}$ , the wealth-level of ambiguity averse investors at which they choose to exit the market is

$$\hat{x}_a = 1 - \frac{\gamma}{\delta \rho_\epsilon}. \quad (33)$$

**Result 2.** *The amount of spooked wealth that will lead to ambiguity averse traders being willing to stay in the market is*

1. *Increasing in the amount of perceived ambiguity  $\delta$ ,*
2. *Increasing in the precision of the signal  $\rho_\epsilon$ , and*
3. *Decreasing in the risk aversion of traders  $\gamma$ .*

The first result follows from the fact that as ambiguity increases *ceteris paribus*, AA traders have an increased desire to leave the market. In order for them to stay, their departure needs to have such a large impact on the risk premium that the risk-premium with them gone is attractive enough for them to stay. So, as ambiguity increases, the amount of wealth that is at risk of flight that will lead to AA traders staying in the market increases.

The second result follows similarly, since the ambiguity premium  $\Delta$  is increasing in the precision of the signal that is received. A more precise signal with a given ambiguity level leads to a larger ambiguity premium because the signal has a larger weight in the posterior beliefs of the traders. So low-variance/high-ambiguity signals look worse to AA traders. One might interpret these kinds of signals as having high correlation with the true, underlying state of the market, but where the model between the signal and the market is difficult to understand. In other words, where there is little question that the signal is being observed correctly, but lots of uncertainty about what that observation means.

Finally, the third result follows from the fact that as risk aversion increases, the risk premium required by remaining EU traders increases, making it more likely that AA traders will want to stay in the market. Therefore, the amount of wealth at risk of flight that will lead to AA traders staying in the market decreases as risk aversion increases.

It is noteworthy that in this model the condition for AA traders to drop out of the market does not depend on the signal that is conveyed. Whether in good news or bad news, the market panics that happen here happen because of the inherent aversion to the ambiguity in the signal. This model shows that it is sufficient for ambiguity averse investors to believe that the signal is unreliable, rather than needing the signal to necessarily be bad. To understand this distinction, the next section discusses the differences between this model and models where some traders interpret the signal to be bad news, but without ambiguity.

### 3.2 A comparison with belief dispersion models

Consider a world where instead of ambiguity aversion, the fraction  $x_p$  of wealth interprets the signal  $s$  as being more negative than the fraction  $x_e$  of wealth. Specifically, suppose that the  $x_p$  fraction of wealth believes that the signal is  $s - \zeta$  while the  $x_e$  fraction of wealth interpret the signal as before. We will assume that all traders receive the same signal, they just interpret it differently. Put differently, we assume that each set of traders believes the other set to be incorrectly interpreting the signal.

In this case, updated beliefs for the two sets of traders are given by

$$m_e(s) = \mu_0 + \frac{\rho_e}{\rho_I}(s - \mu_0) \quad (34)$$

and

$$m_p(s) = \mu_0 + \frac{\rho_e}{\rho_I}(s - \zeta - \mu_0). \quad (35)$$

Here the equilibrium price satisfies the market clearing condition given in equation (16) continues to hold here, as does the demand for traders of type  $e$  given in equation (5). Demand for traders of type  $p$  is given by

$$\theta = \Lambda^{-1}[m_p(s) - p(s)]. \quad (36)$$

Together, the market clearing condition implies that the equilibrium price is

$$p(s) = m(s) - x_p\zeta - \Lambda. \quad (37)$$

which implies an equilibrium risk premium of  $\Lambda + x_p\zeta$  and holdings for the two types of traders that satisfy

$$\theta_e = \Lambda^{-1}[x_p\zeta + \Lambda] \quad (38)$$

and

$$\theta_p = \Lambda^{-1}[(x_p - 1)\zeta + \gamma\sigma_I^2]. \quad (39)$$

This setup has equity premium  $\Lambda + x_p\zeta$  and trade volume of

$$V(\zeta) = \Lambda^{-1}x_ex_p\zeta. \quad (40)$$

Compare this to the case of AA traders with ambiguity level  $\Delta$  and EU traders

Table 1: Comparison of equilibrium premium and volume under ambiguity aversion and belief dispersion

	Common beliefs	Ambiguity w/ exit	Ambiguity w/o exit	Belief dispersion
Equity premium	$\Lambda$	$y(x_e) = \frac{\Lambda}{x_e}$	$\Lambda + \Delta x_a$	$\Lambda + x_p \zeta$
Trade Volume	0	$x_a$	$\Lambda^{-1} x_e x_a \Delta$	$\Lambda^{-1} x_e x_p \zeta$
Wealth Effect	0	$-y(x_e)$	$-\Delta x_a - \Lambda$	$-x_p \zeta - \Lambda$
Trading Effect	0	$\frac{1}{2} \Lambda^{-1} y(x_e)^2$	$\frac{1}{2} \Lambda^{-1} (\Delta x_a + \Lambda)^2$	$\frac{1}{2} \Lambda^{-1} (x_p \zeta + \Lambda)^2$

described previously. If AA traders drop out of the market after receiving the signal, i.e. if  $\Delta \geq y(x_e)$ , volume is  $x_a$  since AA traders exit the market. On the other hand, if AA traders remain in the market, volume is  $\Lambda^{-1} x_e x_a \Delta$ . This implies that volume in the ambiguity aversion case is

$$V(\Delta) = \begin{cases} x_a & \Delta \geq y(x_e) \\ \Lambda^{-1} x_e x_a \Delta & \Delta < y(x_e) \end{cases} \quad (41)$$

We notice here that the difference for EU traders when AA traders leave vs. when they stay is that the equity premium is higher when they leave, trade volume is higher, the wealth effect is larger (i.e. the immediate crash is larger), and the trading effect is also larger. The belief dispersion model looks qualitatively similar to the model with ambiguity aversion where AA, especially if the pessimism  $\zeta$  is equal to the ambiguity premium  $\Delta$ .

## 4 Valuing the possibility of panics before they happen

In order to better understand the nature of the possibility of market panics and their impact even before they happen, this section asks how traders who are subject to potential panic price that possibility prior to it happening.

In this extension of the model, we assume that there are now two decision periods  $t \in \{0, 1\}$  and everyone begins the model at time  $t = 0$  with an initial endowment of one unit of the risky asset and no units of the risk-free asset, as before. At time 0 the ambiguity  $\Delta$  that will be in the signal at time 1 is not known. This framing allows us to explore how the possibility of panics characterized by ambiguity aversion

are priced prior to their potential arrival. We assume that the distribution of  $\Delta$  has CDF  $F(\Delta)$ , which is commonly known. After trading at time 0, the realization of ambiguity  $\Delta$  occurs and the proceeds as in the previous sections at time  $t = 1$ .

At time  $t = 0$ , given beliefs about ambiguity  $F(\Delta)$ , all traders can compute the expected equilibrium price at time  $t = 1$  and the expected utility at time  $t = 1$  for both types of traders.

For each trader type at time  $t = 0$ , there are two possible scenarios that could occur at time  $t = 1$ :

- The ambiguity level  $\Delta$  is low enough that AA traders remain in the market. This occurs with probability  $F(y(x_e)) = 1 - \pi_e^*$ . We call this regime the *no panic* regime.
- The ambiguity level  $\Delta$  is high enough that AA traders exit the market. This occurs with probability  $1 - F(y(x_e)) = \pi_e^*$ . We call this regime the *panic* regime.

We will assume that AA traders do not perceive ambiguity in the distribution of  $\Delta$  at time  $t = 0$  and therefore their time 0 expected utility is the expected value of their utility at time  $t = 1$  across these two scenarios.

#### 4.1 Time-1 indirect utility as a function of time-0 holdings

Since there is the possibility that individuals drop out at time 0 and thus enter time 1 with no holdings, we recalculate holdings and indirect utility here with the holdings at time 1 that were carried forward from time 0 as a parameter.

Consider a trader who begins time  $t = 1$  holding  $\theta_0$  units of the risky asset and  $(1 - \theta_0)p_0$  in risk-free wealth, where  $p_0$  is the time-0 price at which they traded. At time 1 the signal  $s$  and ambiguity level  $\Delta$  are realized, and the trader optimally rebalances at the time-1 equilibrium price  $p_1$ . A trader of type  $i$  with belief  $E_i[v]$  about the risky asset's payoff chooses time-1 holdings  $\theta_1$  to solve

$$\max_{\theta_1} \theta_1 E_i[v] + (\theta_0 - \theta_1)p_1 + (1 - \theta_0)p_0 - \frac{\Lambda}{2} \theta_1^2.$$

The first-order condition gives  $\theta_1^* = \Lambda^{-1}(E_i[v] - p_1)$ , and substituting yields the indirect utility

$$V_i(\theta_0, s, \Delta) = (1 - \theta_0)p_0 + \theta_0 p_1 + \frac{(E_i[v] - p_1)^2}{2\Lambda}. \quad (42)$$

The key feature of (42) is that the trading benefit  $\frac{1}{2\Lambda}(E_i[v] - p_1)^2$  does not depend on  $\theta_0$ . Holdings at time 0 affect utility only through the endowment-value term  $\theta_0 p_1$ .

We now specialize to each regime and trader type. In every case,  $p_1$  is the time-1 equilibrium price from the previous section and  $E_i[v]$  is the trader's belief about the asset payoff.

**No panic regime** ( $\Delta \leq y(x_e)$ ). The time-1 price is  $p_1 = m(s) - \Delta x_a - \Lambda$ .

For an EU trader ( $E_e[v] = m(s)$ ), indirect utility is

$$V_e^\circ(\theta_0) = (1 - \theta_0) p_0 + \theta_0(m(s) - \Delta x_a - \Lambda) + \frac{(\Delta x_a + \Lambda)^2}{2\Lambda}. \quad (43)$$

For an AA trader ( $E_a[v] = m(s) - \Delta$ ), the optimal holding is  $\theta_1^a = \Lambda^{-1}(\Lambda - \Delta x_e) \geq 0$ , and indirect utility is

$$V_a^\circ(\theta_0) = (1 - \theta_0) p_0 + \theta_0(m(s) - \Delta x_a - \Lambda) + \frac{(\Lambda - \Delta x_e)^2}{2\Lambda}. \quad (44)$$

**Panic regime** ( $\Delta > y(x_e)$ ). The time-1 price is  $p_1 = m(s) - y(x_e)$ .

For an EU trader, indirect utility is

$$V_e^*(\theta_0) = (1 - \theta_0) p_0 + \theta_0(m(s) - y(x_e)) + \frac{y(x_e)^2}{2\Lambda}. \quad (45)$$

For an AA trader, optimal holdings are 0 so they earn no trading benefit:

$$V_a^*(\theta_0) = (1 - \theta_0) p_0 + \theta_0(m(s) - y(x_e)). \quad (46)$$

## 4.2 Expected time-0 utility

Taking expectations over  $(s, \Delta)$  using the probability that there is a panic  $\pi_e^*$ , define the time-0 indirect utility of trader type  $i$  as  $W_i(\theta_0) \equiv E[V_i(\theta_0, s, \Delta)]$ . As a function of time-0 holdings and price,

$$W_i(\theta_0) = p_0 + \theta_0(\bar{p}_1 - p_0) + \bar{T}_i, \quad (47)$$

where  $\bar{p}_1 \equiv E[p_1]$  is the expected time-1 price,

$$\bar{p}_1 = (1 - \pi_e^*)(\mu_0 - E[\Delta \mid \Delta \leq y(x_e)] x_a - \Lambda) + \pi_e^*(\mu_0 - y(x_e)), \quad (48)$$

and  $\bar{T}_i$  is the expected trading benefit, which differs across types but does not depend on  $\theta_0$ .

For EU traders,

$$\bar{T}_e = (1 - \pi_e^*) \frac{E[(\Delta x_a + \Lambda)^2 \mid \Delta \leq y(x_e)]}{2\Lambda} + \pi_e^* \frac{y(x_e)^2}{2\Lambda}. \quad (49)$$

For AA traders,

$$\bar{T}_a = (1 - \pi_e^*) \frac{E[(\Lambda - \Delta x_e)^2 \mid \Delta \leq y(x_e)]}{2\Lambda}. \quad (50)$$

AA traders earn no trading benefit in the panic regime because they exit the market entirely.

The structure of (47) makes clear that the time-0 portfolio choice  $\theta_0$  affects expected utility only through the expected price change  $\bar{p}_1 - p_0$ , while the trading benefit  $\bar{T}_i$  enters as a type-specific lump sum. The difference in trading benefits,  $\bar{T}_e - \bar{T}_a$ , captures the welfare advantage that EU traders enjoy from being able to trade profitably in both regimes at time 1.<sup>1</sup>

### 4.3 Time-0 equilibrium price

With this utility specification time-1 equilibrium prices are determined entirely by traders' demands—which depend on beliefs and risk aversion—and not on initial holdings. Whether a trader enters time 1 holding  $\theta_0 = 1$  or  $\theta_0 = 0$ , the time-1 price  $p_1(s, \Delta)$  is the same, because each trader rebalances to the same optimal position  $\theta_1^*$  regardless of starting point. It follows that  $\bar{p}_1$  and the trading benefits  $\bar{T}_e$ ,  $\bar{T}_a$  are independent of the time-0 allocation.

Since  $W_i(\theta_0)$  is affine in  $\theta_0$  (equation (47)), any trader facing  $\bar{p}_1 > p_0$  would demand an unbounded long position, while any trader facing  $\bar{p}_1 < p_0$  would sell entirely. The unique market-clearing price is therefore

$$p_0 = \bar{p}_1, \quad (51)$$

where  $\bar{p}_1$  is given by (48).

---

<sup>1</sup>This is a function of the assumption that traders care only about terminal wealth and nothing the only thing that happens between time 0 and 1 is the realization of the amount of ambiguity.

## 4.4 The panic premium

If there were no possibility of ambiguity ( $\Delta = 0$  with certainty), the time-1 price would always be  $p_1 = m(s) - \Lambda$  and the time-0 price would be

$$p_0^{\text{no amb}} = \mu_0 - \Lambda.$$

To see how the possibility of a future panic depresses the time-0 price, we can rewrite (51) using (48) as

$$p_0 = \mu_0 - \Lambda - x_a E[\min(\Delta, y(x_e))]. \quad (52)$$

To verify, note that the no-panic regime contributes  $x_a E[\Delta \mid \Delta \leq y(x_e)]$  weighted by  $(1 - \pi_e^*)$ , while the panic regime contributes  $x_a y(x_e)$  weighted by  $\pi_e^*$ ; together these give  $x_a E[\min(\Delta, y(x_e))]$ . The panic premium relative to the no-ambiguity benchmark is therefore

$$\Omega = p_0^{\text{no amb}} - p_0 = x_a E[\min(\Delta, y(x_e))] \geq 0. \quad (53)$$

The panic premium  $\Omega$  is increasing in  $x_a$ : a larger share of traders who are susceptible to panic amplifies the premium. It is bounded above by  $x_a y(x_e) = \Lambda x_a/x_e$ , which obtains in the limit where panic is certain ( $\pi_e^* \rightarrow 1$ ).

We can also compare  $p_0$  to the expected time-1 price in each regime. Let  $\bar{p}_1^\circ \equiv E[p_1 \mid \Delta \leq y(x_e)]$  and  $\bar{p}_1^* \equiv E[p_1 \mid \Delta > y(x_e)]$  denote the conditional expected prices in the no-panic and panic regimes respectively:

$$\begin{aligned} \bar{p}_1^\circ &= \mu_0 - E[\Delta \mid \Delta \leq y(x_e)] x_a - \Lambda, \\ \bar{p}_1^* &= \mu_0 - y(x_e). \end{aligned}$$

The no-panic expected price exceeds the panic expected price by

$$\bar{p}_1^\circ - \bar{p}_1^* = x_a [y(x_e) - E[\Delta \mid \Delta \leq y(x_e)]] \geq 0,$$

since  $E[\Delta \mid \Delta \leq y(x_e)] \leq y(x_e)$ . Because  $p_0$  is the probability-weighted average of these two conditional prices,

$$p_0 = (1 - \pi_e^*) \bar{p}_1^\circ + \pi_e^* \bar{p}_1^*,$$

the time-0 price lies between them:

$$\bar{p}_1^* \leq p_0 \leq \bar{p}_1^\circ. \quad (54)$$

The time-0 price is below what will prevail if no panic occurs and above what will prevail if panic does occur. In this sense the time-0 price partially anticipates the panic: it incorporates the expected damage but not the full realized damage.

## 4.5 The role of panic probability

The dependence of the time-0 price on the panic probability  $\pi_e^*$  follows directly from the convex combination

$$p_0 = (1 - \pi_e^*) \bar{p}_1^\circ + \pi_e^* \bar{p}_1^*.$$

Since  $\bar{p}_1^* \leq \bar{p}_1^\circ$ , the time-0 price is decreasing in  $\pi_e^*$ : as the probability of a panic rises, more weight shifts to the lower panic-regime price, and the time-0 price falls. At the extremes, if panic is impossible ( $\pi_e^* = 0$ ) the time-0 price equals the no-panic expected price  $\bar{p}_1^\circ$ , and if panic is certain ( $\pi_e^* = 1$ ) it equals the panic expected price  $\bar{p}_1^* = \mu_0 - y(x_e)$ .

A key implication is that most of the time—whenever panic does not occur—the realized time-1 price exceeds the time-0 price:

$$p_1^\circ - p_0 = \pi_e^* (\bar{p}_1^\circ - \bar{p}_1^*) > 0 \quad \text{when } \pi_e^* > 0.$$

An investor who holds the risky asset from time 0 to time 1 therefore earns a positive expected capital gain conditional on no panic, and this gain is increasing in  $\pi_e^*$ . Conversely, when panic does occur, the realized price is below  $p_0$ :

$$p_1^* - p_0 = -(1 - \pi_e^*) (\bar{p}_1^\circ - \bar{p}_1^*) < 0 \quad \text{when } \pi_e^* < 1.$$

This pattern is reminiscent of the pre-announcement return premium documented in equity and fixed-income markets. Empirically, average stock returns are substantially higher on days immediately preceding scheduled macroeconomic announcements—such as FOMC meetings or employment reports—than on non-announcement days (Lucca & Moench 2015). In the model, the time before an announcement corresponds to time 0, when the probability of a panic-inducing information shock is elevated. The

time-0 price is depressed by the panic premium, creating a positive expected return that is earned when the announcement resolves favorably. Higher  $\pi_e^*$ —corresponding to announcements perceived as more likely to trigger a panic—generates a larger pre-announcement discount and therefore a larger expected return conditional on no panic. The model thus provides a mechanism through which heterogeneous responses to ambiguity can generate pre-announcement risk premia without requiring any difference in information across traders.

The same mechanism also generates a pattern analogous to the variance risk premium—the well-documented finding that option-implied volatility systematically exceeds realized volatility (Carr & Wu 2009, Bollerslev et al. 2009). In the model, the time-0 price embeds the full distribution of time-1 outcomes, including the low-probability panic regime. In most realizations ( $\Delta \leq y(x_e)$ ), the panic does not occur and the realized price movement from  $p_0$  to  $p_1^o$  is smaller in magnitude than what the time-0 price “expected.” The price had anticipated a possible crash that did not materialize, just as option-implied volatility reflects tail scenarios that typically do not occur. Only in the panic regime does the realized price drop exceed the average expected movement. Bollerslev & Todorov (2011) decompose the variance risk premium empirically and show that compensation for rare tail events—crashes that are priced in but seldom observed—accounts for a large share of the total equity and variance risk premia, which is precisely the structure of the panic premium in our model. Moreover, Miao et al. (2019) show that ambiguity aversion, rather than risk aversion alone, is the primary driver of the variance risk premium, attributing approximately 96% of the mean variance premium to ambiguity aversion in a calibrated model. This provides further support for the connection between our ambiguity-based panic mechanism and the implied-versus-realized volatility gap observed in options markets.

## 4.6 The distribution of time-1 prices: general ambiguity

Since  $s = v + \epsilon$  is the sum of two independent normals— $v \sim N(\mu_0, \rho_0^{-1})$  and  $\epsilon \sim N(0, \rho_\epsilon^{-1})$ —the posterior mean  $m(s) = \mu_0 + \frac{\rho_\epsilon}{\rho_I}(s - \mu_0)$  is also normal. By the law of total variance,

$$\sigma_m^2 \equiv \text{Var}(m(s)) = \text{Var}(v) - \text{Var}(v|s) = \frac{1}{\rho_0} - \frac{1}{\rho_I} = \frac{\rho_\epsilon}{\rho_0 \rho_I}, \quad (55)$$

so  $m(s) \sim N(\mu_0, \sigma_m^2)$ . The variance  $\sigma_m^2$  captures how much of the prior uncertainty about  $v$  is resolved by observing  $s$ .

The time-1 prices in both regimes admit the unified representation

$$p_1 = m(s) - \Lambda - x_a \min(\Delta, y(x_e)). \quad (56)$$

To verify: in the no-panic regime  $\min(\Delta, y(x_e)) = \Delta$ , giving  $m(s) - x_a \Delta - \Lambda$ ; in the panic regime  $\min(\Delta, y(x_e)) = y(x_e)$ , and  $\Lambda + x_a y(x_e) = \Lambda(1 + x_a/x_e) = \Lambda/x_e = y(x_e)$ , giving  $m(s) - y(x_e)$ . Because  $s$  is informative about  $v$  but independent of the ambiguity shock  $\Delta$ , the random variables  $m(s)$  and  $\Delta$  are independent, so (56) decomposes  $p_1$  into a normal component  $m(s) - \Lambda$  and an independent censored component  $-x_a \min(\Delta, y(x_e))$ .

Conditioning on  $\Delta = \Delta_0$ , the price  $p_1 | \Delta = \Delta_0$  is normal with mean  $\mu_0 - \Lambda - x_a \min(\Delta_0, y(x_e))$  and variance  $\sigma_m^2$ . Integrating over  $\Delta$  gives the unconditional density:

$$f_{p_1}(p) = \int_0^{y(x_e)} \phi_N(p; \mu_0 - x_a z - \Lambda, \sigma_m^2) f_\Delta(z) dz + \pi_e^* \phi_N(p; \mu_0 - y(x_e), \sigma_m^2), \quad (57)$$

where  $f_\Delta$  is the density of  $\Delta$  and  $\pi_e^* = 1 - F(y(x_e))$ . The first term is the no-panic contribution—a location mixture of normals whose mean is randomized by the draw  $\Delta | \Delta \leq y(x_e)$ —and the second term is the panic contribution, a normal centered at  $\mu_0 - y(x_e)$ . Equation (57) holds for any  $F$  with support on  $[0, \infty)$ ; no closed-form simplification is available without further specifying  $F$ .

**Mean and variance.** Using the law of iterated expectations and the independence of  $m(s)$  and  $\Delta$ ,

$$E[p_1] = \mu_0 - \Lambda - x_a E[\min(\Delta, y(x_e))] = p_0, \quad (58)$$

confirming (52) for any  $F$ . For the variance, the law of total variance applied to (56) gives

$$\text{Var}(p_1) = \sigma_m^2 + x_a^2 \text{Var}(\min(\Delta, y(x_e))). \quad (59)$$

The variance of the censored variable  $Z = \min(\Delta, y(x_e))$  can be expressed using the survival function  $\bar{F}(z) = 1 - F(z)$ :

$$\text{Var}(Z) = 2 \int_0^{y(x_e)} z \bar{F}(z) dz - \left( \int_0^{y(x_e)} \bar{F}(z) dz \right)^2, \quad (60)$$

where the first integral uses  $E[Z^2] = 2 \int_0^{y(x_e)} z P(Z > z) dz$  and  $P(Z > z) = \bar{F}(z)$  for  $z < y(x_e)$ . Both formulas hold for any  $F$  with support on  $[0, \infty)$ .

**Non-normality and skewness.** The distribution (57) is non-normal whenever  $F$  is non-degenerate. To characterize the direction of skewness, write  $p_1 - p_0 = A - x_a B$ , where  $A = m(s) - \mu_0 \sim N(0, \sigma_m^2)$  and  $B = Z - E[Z]$  are independent. Since  $A$  is symmetric, all odd central moments of  $A$  vanish, so

$$E[(p_1 - p_0)^3] = -x_a^3 \mu_3(Z), \quad (61)$$

where  $\mu_3(Z) = E[(Z - E[Z])^3]$  is the third central moment of  $Z$ . The time-1 price distribution is left-skewed if and only if  $\mu_3(Z) > 0$ , i.e. if  $Z = \min(\Delta, y(x_e))$  is right-skewed. This holds whenever panic is possible ( $\pi_e^* > 0$ ), because the point mass  $\pi_e^*$  at the upper boundary  $y(x_e)$  pushes the right tail of  $Z$  above what the continuous part of  $F$  alone would generate. Intuitively, occasional large ambiguity shocks that trigger a panic translate into occasional large leftward price movements, producing a left-skewed price distribution even when the fundamental payoff uncertainty is symmetric.

**Excess kurtosis.** By an analogous calculation using  $E[(p_1 - p_0)^4] = 3\sigma_m^4 + 6\sigma_m^2 x_a^2 \sigma_Z^2 + x_a^4 \mu_4(Z)$ , where  $\sigma_Z^2 = \text{Var}(Z)$  and  $\mu_4(Z) = E[(Z - E[Z])^4]$ , the excess kurtosis of  $p_1$  relative to a normal with the same variance is positive whenever  $\mu_4(Z) > 3\sigma_Z^4$ , i.e. whenever  $Z$  itself has excess kurtosis. The censoring of  $\Delta$  at  $y(x_e)$  and the resulting point mass generically produce excess kurtosis in  $Z$ , so the time-1 price distribution has heavier tails than a normal with matching mean and variance.

**Connection to empirical evidence.** The left skewness and excess kurtosis derived above are not artifacts of the exponential specification—they hold for any ambiguity distribution  $F$  that makes panic possible—and align closely with well-documented properties of equity returns. Empirically, equity return distributions are left-skewed

and exhibit fat tails relative to a Gaussian benchmark. Bollerslev & Todorov (2011) find that the variance risk premium is almost entirely attributable to the compensation for jump tail risk, and that this tail-risk component is separately priced in the cross-section of returns; our panic event constitutes precisely such an endogenous left-tail event. Carr & Wu (2009) and Bollerslev et al. (2009) document that the variance risk premium—the gap between risk-neutral and physical variance—is negative and time-varying, implying that investors pay a premium to hedge variance uncertainty beyond fundamental payoff risk. This corresponds directly to the  $x_a^2 \text{Var}(Z)$  term in (59): it is the component of price variance attributable to ambiguity, on top of the fundamental uncertainty  $\sigma_m^2$ , for which ambiguity-averse investors demand compensation. The density (57) also formalizes the observation that equity returns are well-approximated by a mixture of normals (Wachter 2013): the no-panic term is a location mixture of Gaussians whose mean is randomized by the draw of  $\Delta$ , while the panic term contributes an additional normal component shifted to the left.

**Connection to the ambiguity and rare-disasters literature.** On the theoretical side, the results connect to two strands of the asset pricing literature. Epstein & Schneider (2008) show that ambiguity about signal quality generates asymmetric price responses to information and a non-normal return distribution: because ambiguity-averse investors act on worst-case beliefs about signal precision, bad news lowers prices by more than good news raises them, producing left skewness. Our mechanism is distinct—ambiguity here shifts beliefs about asset value rather than signal quality—but generates the same directional prediction: left skewness and excess kurtosis arise endogenously from the equilibrium panic threshold  $y(x_e)$ , without any asymmetry in the information structure. Miao et al. (2019) show that ambiguity aversion explains the negative variance risk premium through the pessimistic probability distortions it induces; the excess-variance term  $x_a^2 \text{Var}(Z)$  in (59) provides a complementary quantity-theoretic decomposition in which the contribution of ambiguity to price variance is transparent and traceable to the distribution of the censored ambiguity shock  $Z$ .

The panic event also plays a role analogous to a rare disaster in the tradition of Barro (2006) and Gabaix (2012): a low-probability, large-magnitude event that generates fat tails, excess expected returns, and a non-normal price distribution. The critical difference is that the disaster threshold is here endogenous. The panic occurs

if and only if  $\Delta > y(x_e) = \Lambda/x_e$ , so the probability  $\pi_e^* = 1 - F(y(x_e))$  and the severity  $y(x_e)$  of the tail event both depend on equilibrium market composition. An increase in  $x_e$  (more expected-utility traders) lowers the threshold  $y(x_e)$ , raises  $\pi_e^*$ , and thus increases the probability and severity of the left-tail event—a feedback between market composition and distributional risk that has no analogue in the exogenous rare-disasters framework.

## 4.7 Return distribution with exponential ambiguity

The general density (57) does not simplify to a closed form without further structure on  $F$ . This subsection demonstrates the properties of the previous subsection for a specific example.

Specializing to  $\Delta \sim \text{Exp}(\psi)$ , with density  $f_\Delta(\delta) = \psi e^{-\psi\delta}$  for  $\delta \geq 0$ , yields closed-form expressions for all key quantities. The exponential distribution provides a parsimonious one-parameter family in which higher  $\psi$  corresponds to lower expected ambiguity ( $E[\Delta] = 1/\psi$ ).

The panic probability is

$$\pi_e^* = e^{-\psi y(x_e)}. \quad (62)$$

The panic premium (53) becomes

$$p_0^{\text{no amb}} - p_0 = \frac{x_a}{\psi} (1 - e^{-\psi y(x_e)}), \quad (63)$$

and the time-0 price is

$$p_0 = \mu_0 - \Lambda - \frac{x_a}{\psi} (1 - e^{-\psi y(x_e)}). \quad (64)$$

The total return on the risky asset, measured from purchase at time 0 to the terminal payoff, is  $r = v - p_0$ . Since the terminal value  $v \sim N(\mu_0, \rho_0^{-1})$  and  $p_0$  is a known constant, the total return is normally distributed:

$$r = v - p_0 \sim N\left(\Lambda + \frac{x_a}{\psi} (1 - e^{-\psi y(x_e)}), \rho_0^{-1}\right). \quad (65)$$

The expected total return decomposes into two components:

$$\begin{aligned}
E[r] &= \underbrace{\Lambda}_{\text{baseline risk premium}} + \underbrace{\frac{x_a}{\psi} (1 - e^{-\psi y(x_e)})}_{\text{panic premium}} \\
&= \Lambda + \Omega
\end{aligned} \tag{66}$$

where  $\Omega$  denotes the panic premium.

The baseline risk premium  $\Lambda = \gamma\sigma_I^2$  is the compensation for bearing fundamental payoff risk and would prevail even without the possibility of ambiguity. The panic premium is the additional compensation investors earn for holding an asset whose interim price may crash due to a panic. It is increasing in  $x_a$  (the fraction of panic-susceptible traders) and decreasing in  $\psi$  (since lower  $\psi$  means higher expected ambiguity and therefore a greater chance of panic).

In the limit  $\psi \rightarrow \infty$  (no ambiguity), the panic premium vanishes and  $E[r] \rightarrow \Lambda$ . In the limit  $\psi \rightarrow 0$  (very high expected ambiguity, panic nearly certain), the panic premium approaches its upper bound  $x_a y(x_e) = \Lambda x_a/x_e$ , and the expected return becomes  $\Lambda(1 + x_a/x_e) = \Lambda/x_e$ .

Figure 1 plots the panic premium as a fraction of the total premium  $E[r]$  as a function of  $\psi$  and  $x_a$ . The panic premium constitutes a larger share of the total expected return when  $\psi$  is low (high expected ambiguity) and when  $x_a$  is high (more panic-susceptible traders). In an environment with high expected ambiguity and a large fraction of panic-prone investors, the panic premium can account for a substantial portion of the total expected return.

The variance of the total return is  $\rho_0^{-1}$ , which does not depend on the ambiguity parameters  $\psi$  or  $x_a$ . The possibility of a panic raises the expected return but does not affect return volatility. This is because the total return depends only on the terminal payoff  $v$ , which is independent of the ambiguity realization  $\Delta$ . The non-normal features—left skewness and excess kurtosis—derived above for the interim price return  $p_1 - p_0$  are absent from the total return, since they are generated by the intermediate trading dynamics rather than by the terminal payoff.

**Comparative statics in  $\psi$ .** The exponential rate parameter  $\psi$  governs the entire distribution of ambiguity, and its comparative statics illuminate how the severity of the ambient uncertainty environment affects the panic premium, expected returns,

The Panic Premium as Fraction of the Total Equity Premium

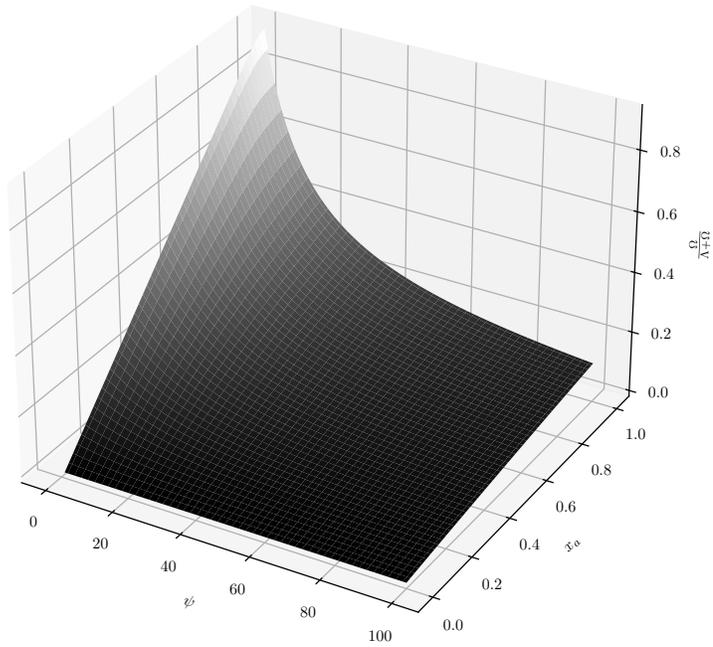


Figure 1: Panic premium as a fraction of total expected return,  $\Omega/E[r]$ , as a function of the ambiguity parameter  $\psi$  and the fraction of panic-susceptible traders  $x_a$ . The panic premium constitutes a larger share of the total expected return when expected ambiguity is higher (lower  $\psi$ ) and when more traders are susceptible to panic (higher  $x_a$ ).

and the shape of the interim price distribution.

*Panic premium and expected return.* The panic premium (63) is strictly decreasing in  $\psi$ . To see this, note that  $\frac{\partial}{\partial \psi} \left[ \frac{1-e^{-\psi y}}{\psi} \right] = \frac{e^{-\psi y}(1+\psi y)-1}{\psi^2}$ , which is negative for all  $\psi > 0$  since  $e^t(1-t) < 1$  for  $t > 0$ .<sup>2</sup> A decrease in  $\psi$ —representing a more ambiguous environment—raises both the panic probability  $\pi_e^* = e^{-\psi y(x_e)}$  and the expected magnitude of the ambiguity shock, so investors demand a larger premium to hold the risky asset. The two limiting cases bound the premium: as  $\psi \rightarrow \infty$  the premium vanishes and  $E[r] \rightarrow \Lambda$ , recovering the frictionless risk premium; as  $\psi \rightarrow 0$  the premium approaches its maximum  $x_a y(x_e) = \Lambda x_a/x_e$  and  $E[r] \rightarrow \Lambda/x_e$ .

*Left skewness.* From (61), the left skewness of  $p_1$  is  $-x_a^3 \mu_3(Z)$  where  $Z = \min(\Delta, y(x_e))$  and  $\Delta \sim \text{Exp}(\psi)$ . The signed skewness is negative (left-skewed) for all  $\psi < \infty$ , and its magnitude is strictly decreasing in  $\psi$ . Intuitively, as  $\psi$  increases the panic probability  $\pi_e^*$  falls, the point mass that  $Z$  places at the upper boundary  $y(x_e)$  shrinks, and the right tail of  $Z$  becomes less pronounced—so the implied left tail in  $p_1$  becomes thinner. In the limit  $\psi \rightarrow \infty$  the panic probability vanishes, the point mass disappears entirely, and  $\mu_3(Z) \rightarrow 0$ : the distribution of  $p_1$  converges to a normal. Conversely, as  $\psi \rightarrow 0$  the panic becomes nearly certain, almost all of the mass of  $Z$  is concentrated at  $y(x_e)$ , and left skewness is maximized. Thus the same parameter that governs the panic premium also governs the direction and magnitude of return asymmetry: markets with higher expected ambiguity simultaneously offer higher expected returns *and* exhibit more left-skewed interim price distributions.

*Excess kurtosis.* Excess kurtosis of  $p_1$  follows the same pattern. As  $\psi$  increases and the point mass at  $y(x_e)$  shrinks, the fourth central moment of  $Z$  falls toward that of a normal distribution, and excess kurtosis of  $p_1$  diminishes toward zero. As  $\psi \rightarrow 0$ , the heavy concentration of  $Z$  near  $y(x_e)$  generates large fourth central moments and pronounced fat tails. A key implication is that the panic premium, left skewness, and excess kurtosis all move together as  $\psi$  varies: they share a common driver in the panic probability  $\pi_e^* = e^{-\psi y(x_e)}$ . An environment with a larger ambiguity premium is also one with a more non-normal, left-skewed, and leptokurtic return distribution—a joint prediction that distinguishes the ambiguity mechanism from risk-based models in which higher expected returns need not be accompanied by greater return asymmetry.

---

<sup>2</sup>Equivalently,  $E[\Delta] = 1/\psi$  decreases in  $\psi$ , and the panic premium is approximately  $x_a E[\min(\Delta, y(x_e))]$  which shares this monotonicity.

## 5 Conclusion

This paper develops an equilibrium model of market panics driven by ambiguity aversion and derives their consequences for welfare, asset prices, and return distributions. When the realized ambiguity shock exceeds the endogenously determined panic threshold, ambiguity-averse traders exit entirely, raising the risk premium and generating a discrete price decline. Despite the capital loss, expected-utility traders who remain always weakly benefit: the trading gain from the elevated risk premium dominates the wealth loss for any ambiguity level below the threshold. In the two-period extension, the time-zero price is depressed by a panic premium that is increasing in both the severity of the potential ambiguity shock and the share of panic-susceptible traders, and lies strictly between the expected prices in the panic and no-panic regimes. The full unconditional distribution of interim prices is generically non-normal even when all fundamentals are normal: it is left-skewed and leptokurtic, with the non-normality tracing entirely to the censoring of the ambiguity shock at the panic threshold. Price variance decomposes into a fundamental component and an ambiguity-driven component that vanishes when ambiguity is absent, providing structural microfoundations for the fat tails, left skewness, and excess variance premium documented in equity markets.

Several directions for future research follow naturally from these results. The closed-form expressions under exponential ambiguity provide a foundation for structural estimation: the panic premium, panic probability, and return distribution shape are all functions of a small number of parameters that could be identified from data on equity premia, variance risk premia, and return moments, with connections to observable uncertainty proxies such as the VIX. Endogenizing the fraction of ambiguity-averse traders and firms' disclosure choices would allow market composition and the panic threshold to respond to information policy, enriching both the welfare and regulatory implications of the model. A dynamic extension with time-varying ambiguity would generate time-varying panic risk connecting to the disaster-risk literature of Gabaix (2012) and Wachter (2013), while a multi-asset extension would allow for contagion cascades in which a panic in one market elevates ambiguity in related markets. Finally, the left-skewed, leptokurtic return distribution implies a distinctive implied volatility surface shape that differs from what stochastic-volatility or jump-diffusion models generate, since the non-normality here arises from an endogenous

threshold rather than an assumed process, suggesting a direct empirical test of the panic mechanism using options data.

## References

- Barro, R. J. (2006), ‘Rare disasters and asset markets in the twentieth century’, *Quarterly Journal of Economics* **121**(3), 823–866.
- Bollerslev, T., Tauchen, G. & Zhou, H. (2009), ‘Expected stock returns and variance risk premia’, *Review of Financial Studies* **22**(11), 4463–4492.
- Bollerslev, T. & Todorov, V. (2011), ‘Tails, fears, and risk premia’, *Journal of Finance* **66**(6), 2165–2211.
- Carr, P. & Wu, L. (2009), ‘Variance risk premiums’, *Review of Financial Studies* **22**(3), 1311–1341.
- Condie, S. & Ganguli, J. V. (2011), ‘Ambiguity and rational expectations equilibria’, *Review of Economic Studies* **78**(3), 821–845.
- Condie, S. & Ganguli, J. V. (2017), ‘The pricing effects of ambiguous private information’, *Economic Theory* **64**(3), 535–567.
- Dow, J. & Werlang, S. R. d. C. (1992), ‘Uncertainty aversion, risk aversion, and the optimal choice of portfolio’, *Econometrica* **60**(1), 197–204.
- Easley, D. & O’Hara, M. (2009), ‘Ambiguity and nonparticipation: The role of regulation’, *Review of Financial Studies* **22**(5), 1817–1843.
- Easley, D. & O’Hara, M. (2010), ‘Microstructure and ambiguity’, *Journal of Finance* **65**(5), 1827–1849.
- Epstein, L. G. & Schneider, M. (2008), ‘Ambiguity, information quality, and asset pricing’, *Journal of Finance* **63**(1), 197–228.
- Epstein, L. G. & Wang, T. (1994), ‘Intertemporal asset pricing under Knightian uncertainty’, *Econometrica* **62**(2), 283–322.

- Gabaix, X. (2012), ‘Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance’, *Quarterly Journal of Economics* **127**(2), 645–700.
- Gilboa, I. & Schmeidler, D. (1989), ‘Maxmin expected utility with non-unique prior’, *Journal of Mathematical Economics* **18**(2), 141–153.
- Illeditsch, P. K. (2011), ‘Ambiguous information, portfolio inertia, and excess volatility’, *Journal of Finance* **66**(6), 2213–2247.
- Illeditsch, P. K., Ganguli, J. V. & Condie, S. (2021), ‘Information inertia’, *Journal of Finance* **76**(1), 443–479.
- Lucca, D. O. & Moench, E. (2015), ‘The pre-FOMC announcement drift’, *Journal of Finance* **70**(1), 329–371.
- Mehra, R. & Prescott, E. C. (1985), ‘The equity premium: A puzzle’, *Journal of Monetary Economics* **15**(2), 145–161.
- Miao, J., Wei, B. & Zhou, H. (2019), ‘Ambiguity aversion and the variance premium’, *Quarterly Journal of Finance* **9**(2), 1–36.
- Mukerji, S. & Tallon, J.-M. (2001), ‘Ambiguity aversion and incompleteness of financial markets’, *Review of Economic Studies* **68**(4), 883–904.
- Rietz, T. A. (1988), ‘The equity risk premium: A solution’, *Journal of Monetary Economics* **22**(1), 117–131.
- Wachter, J. A. (2013), ‘Can time-varying risk of rare disasters explain aggregate stock market volatility?’, *Journal of Finance* **68**(3), 987–1035.