

Market panics, bank leverage, and risk capital

Scott Condie*

July 20, 2026

Abstract

Panics on the asset side of banks' balance sheets like those that occurred in the 2007–2009 financial crisis can lead to banks exiting some markets, which transfers risk to the leveraged banks that remain. This paper studies the phenomenon by modeling some banks as ambiguity averse. In the model, banks face a risk-based capital requirement capping leverage at a multiple of a downside risk measure. Slack in normal times, this requirement binds in the panic. In equilibrium, scarce equity deepens crashes. If the crash is deep enough, banks that would otherwise exit do not, building a floor for the crash. The prospect of a panic leaves the interim bank asset price distribution left-skewed and leptokurtic even with normal fundamentals. Expected shortfall loads on the possibility of a panic continuously and ahead of the event, whereas value-at-risk-based capital is flat and then jumps. For financial stability this procyclicality points to a countercyclical capital buffer indexed to ambiguity rather than to volatility or credit growth. In U.S. Call Reports, bank asset returns are shown to be left-skewed and leptokurtic. A new ambiguity measure, built from banks' provisioning disagreement, forecasts changes in the tail distribution of bank assets.

Keywords: financial stability, macroprudential policy, countercyclical capital buffer, ambiguity, bank capital, fire sales, expected shortfall, procyclicality, tail risk.

JEL classification: G01, G21, G28, D81, G12, G11.

1 Introduction

Banks and other leveraged intermediaries hold portfolios of risky assets, such as loans, securities, and trading positions, funded by equity capital and by runnable liabilities, and against which they must hold regulatory capital. Both bankers and their regulators are frequently concerned that a market for these assets may be gripped by a panic: a sudden withdrawal of investors that drives prices down and concentrates risk on those who remain. This paper provides a framework for thinking about tail risk and bank capital in these circumstances

*Department of Economics, Brigham Young University. Email: scott_condie@byu.edu.

and illustrates how the possibility of a panic shapes the distribution of asset values. It also shows how the downside risk measures on which capital is built, namely value at risk under the Basel II market-risk framework and expected shortfall under the Fundamental Review of the Trading Book (Basel Committee on Banking Supervision 2019), respond to that possibility. It also shows how the capital requirement, when binding, helps to determine the depth of the fire sale. The model provides conceptual guidance for a risk office or a supervisor who fears a panic and wants to quantify, before the panic occurs, the additional tail risk that the panic implies and the capital that should be held against it. Using that intuition, the paper builds a new measure of ambiguity from the cross-sectional dispersion of banks' loss provisioning, and shows that it forecasts both the deepening of the left tail of bank asset returns one to four quarters ahead and, in a bank-level panel, the depletion of regulatory capital.

In the financial crisis of 2007–2009 losses mounted on subprime mortgages and investors could no longer confidently value securities built on U.S. housing. Confronted with this ambiguity, holders withdrew and one market after another seized: asset-backed commercial paper, repo against mortgage collateral (Brunnermeier & Pedersen 2009), auction-rate securities, and, after the Reserve Primary Fund “broke the buck,” prime money market funds. Caballero & Krishnamurthy (2008) characterize the episode as a flight to quality driven by Knightian uncertainty, in which agents who cannot assign probabilities retreat to assets they understand. Crucially, the intermediaries that had to absorb the assets others were fleeing were thinly capitalized and highly levered, so the withdrawal became a fire sale. With little equity to bear the risk, and with risk-based capital requirements tightening as measured risk rose, the banks that remained could hold the risky portfolio only at a steep markdown. This episode is not unique. In the “dash for cash” of March 2020 even U.S. Treasury and investment-grade corporate markets dislocated until the Federal Reserve intervened, and the failure of Silicon Valley Bank and other regional lenders in March 2023 touched off a flight from an entire class of bank funding. In each of these episodes, ambiguity about the value of bank assets, interacting with the leverage and capital of those who must hold them, was a plausible mechanism that drove participants out of the market and reshaped the distribution of prices.

This paper develops an equilibrium model in which this kind of panic arises endogenously. The model characterizes the downside risk measures that the resulting price distribution implies, namely value at risk and expected shortfall, together with the capital held against them. The market is populated by two kinds of banks. Expected-utility banks form beliefs in the standard Bayesian way. Ambiguity-averse banks perceive the noise in a public signal as drawn from a set of distributions and evaluate positions under the worst case, in the

maximin tradition of Gilboa & Schmeidler (1989). A sufficiently large ambiguity shock leads the ambiguity-averse banks to hold a zero position, that is, to panic and exit. Each bank is leveraged. It funds its portfolio with equity capital and risk-free liabilities, and it is bound by a risk-based capital requirement that caps leverage at a multiple of the downside risk of its position, measured by a risk measure. The requirement is slack in normal times but binds in a panic, where the scarce capital of the banks that must absorb the portfolio of exiting intermediaries determines how far the price falls. A two-period extension makes the amount of ambiguity, and therefore the potential magnitude of the panic, itself uncertain. That extension shows how a potential panic shapes the distribution of asset values, and hence the amount of capital that should be required. In this setting prices incorporate the possibility of a panic before it occurs, and that possibility changes the properties of the distribution of interim prices from which value at risk and expected shortfall are computed.

The equilibrium analysis yields several results. Ambiguity-averse banks exit when the risk premium that would prevail in their absence falls short of the ambiguity premium they require to stay. When they exit, asset values fall discretely. How far they fall depends on capital. When the banks that remain are well capitalized, they absorb the portfolio at a premium set only by their scarcer risk-bearing capacity. When their capital is scarce, the requirement binds and the fire sale deepens further, by an amount that is larger the less equity they hold and the more conservative the risk measure. This is an intermediary-capital channel in which risk prices spike as net worth shrinks (He & Krishnamurthy 2013, Adrian & Boyarchenko 2012). Because expected shortfall is more stringent than value at risk at any common confidence level, a binding expected-shortfall requirement produces the deeper crash. The deepening is bounded, however. If the potential fire sale is severe enough that the discounted premium exceeds what the ambiguity-averse banks require to stay in the market, they remain, placing a floor under the price.

In the period prior to the possible panic, the anticipation of a possible future panic depresses the current price by a panic premium that is increasing in the share of panic-susceptible banks and in the expected severity of the ambiguity shock. When the panic is expected to be capital-constrained, the premium is larger when the banking sector is thinly capitalized and the binding requirement conservative. Part of the expected return on the asset is therefore compensation specifically for bearing the risk of a future exit. The unconditional distribution of interim prices is left-skewed and leptokurtic even when the fundamental payoff is symmetric and normally distributed, and its variance decomposes into a fundamental component and an ambiguity-driven component. The ambiguity component vanishes when ambiguity is absent, and the remaining banks cannot insure against it, because it originates in the exit of their counterparties. A thinner banking sector, or a stricter risk

measure, moves the panic further out and thereby deepens the left skew and fattens the left tail of the interim price distribution. This comparative static is absent from the frictionless model, and it ties the shape of the loss distribution directly to bank capital.

From this distribution the paper derives expected shortfall, the average loss in the worst outcomes, in closed form as a function of the panic probability and the amount of ambiguity. It responds continuously and ahead of the event, rising as a panic becomes more likely. It recedes only in the limit where the panic is nearly certain and already embedded in the price. Value at risk behaves very differently: it is flat in the panic probability and then jumps once the panic is likely enough to reach the chosen quantile. Variance is moreover not a sufficient statistic for this risk. A normal model, matched to the correct total variance, understates the magnitude of tail risk, because it ignores the skewness the panic imparts, and in the panic-dominated region the understatement is large. In the model's calibration the true expected shortfall exceeds its variance-matched normal benchmark by about a quarter, widening toward 40% as the banking sector thins. Because this loss distribution is an equilibrium object, so is the capital built on it. The conservatism of the risk measure then carries a two-sided effect that the statistical comparison of value at risk and expected shortfall does not reach: a more conservative measure accumulates capital earlier and more gradually, but a binding one deepens the very fire sale it guards against.

These predictions are borne out in the data, most pointedly in the assets banks actually hold. In the quarterly U.S. Call Reports, covering 1.24 million bank-quarters from 1984 to 2026, the return on bank assets is left-skewed and leptokurtic, an asymmetry masked in bank *equity* by leverage and the government guarantee. Its left tail deepens sharply in crises, turning strongly negative in 2008–2009. This is the realized counterpart of the model's panic, in the very portfolio that bank capital is held against. A proxy for ambiguity is then built from the banks themselves. The magnitude of differences in bank views on the distribution of asset values is estimated with the cross-sectional dispersion of their loss provisioning, purged of what realized credit quality (their nonperforming loans and net charge-offs) explains. What remains is disagreement about losses not yet evident in the loan book. This bank-data measure behaves as the model suggests. It spikes in the banking crises of the 1980s and of 2008–2009, and stays muted in market-stress episodes such as 2020 that did not become banking crises. It forecasts the deepening of the left tail of bank asset returns one to four quarters ahead. At the bank level, with bank and quarter fixed effects, it also forecasts the depletion of the regulatory capital the model concerns: a bank whose provisioning is harder to explain from its own realized losses is subsequently more likely to cut its dividend and to see its Tier 1 risk-based ratio fall, relative to other banks in the same quarter. This appears to be the first measure of ambiguity built from bank provisioning. It lets the model's central

object, ambiguity about asset values, be observed in the data on which bank capital is assessed, where an equity-market index, tracking market rather than banking stress, cannot see it.

The paper’s contribution is to derive expected shortfall, and the loss distribution that underlies it, as an equilibrium object in a market of leveraged, capital-constrained, ambiguity-averse banks, so that the capital requirement both measures the tail risk of a possible panic and, once it binds, shapes the panic’s depth. The paper connects this object to the regulation of bank risk capital and, empirically, to a new measure of ambiguity read directly from banks’ own loss provisioning, observed in the reports on which their capital is assessed. For financial stability, the analysis points to a macroprudential rule: a countercyclical capital buffer indexed to ambiguity rather than to volatility or the credit-to-GDP gap, built on a coherent tail measure so that capital accumulates ahead of a panic rather than procyclically at its onset.

1.1 Related literature

This paper draws on four literatures: the measurement of ambiguity from the dispersion of beliefs, the duality between coherent risk measures and ambiguity aversion, the equilibrium asset pricing of ambiguity-driven exit and limited participation, and the intermediary-capital and financial-stability work on fire sales and the procyclicality of risk-based capital. Its contribution is threefold. The first is to derive expected shortfall and its relationship to value at risk as equilibrium objects in a market of leveraged, capital-constrained, ambiguity-averse banks. The second is to show that a binding capital requirement sets the depth of the fire sale while re-entering ambiguity-averse capital places a floor under it. The third is to measure ambiguity directly from the cross-sectional dispersion of banks’ loss provisioning.

Measuring ambiguity from the dispersion of beliefs. Ambiguity concerns the *distribution* of returns rather than their realizations, so it cannot be read off a return series the way volatility can. A productive tradition instead measures it as a dispersion of *beliefs*. Anderson, Ghysels & Juergens (2009) use the disagreement among professional forecasters and Brenner & Izhakian (2018) the volatility of the probabilities of returns, and both find it priced in equity markets. The bank-data measure here carries this approach into the assets intermediaries hold by using the dispersion of banks’ loss forecasts, their loan-loss provisions, purged of what realized credit quality explains. Ambiguity is measured in bank assets rather than in equity claims on them, since leverage and beliefs about a government backstop may obscure the tail risk in question (Section 3.2). The measure is, if anything, negatively correlated with the Brenner & Izhakian (2018) equity-market index and rises in banking crises

that leave it muted, so it captures a banking-sector ambiguity an equity-based measure does not. This appears to be the first application of the dispersion-of-beliefs approach to bank provisioning.

Risk measures and the duality with ambiguity. Artzner, Delbaen, Eber & Heath (1999) show that every coherent risk measure is a worst-case expected loss over a set of priors, the same object as the maxmin expected utility of Gilboa & Schmeidler (1989) (see also Maccheroni, Marinacci & Rustichini (2006)). The regulatory risk measure and the ambiguity aversion that drives the panic are thus two faces of one construct, and this duality is the conceptual hinge the paper builds on. Expected shortfall is such a coherent measure even for the skewed, atom-bearing loss the panic here generates (Acerbi & Tasche 2002), the backdrop to its adoption under the Fundamental Review of the Trading Book (Basel Committee on Banking Supervision 2019). The current paper builds on these representations of preferences under uncertainty and derives their implications for banks' portfolios and capital holdings. A related strand computes worst-case value at risk for a single optimizing agent facing an exogenous ambiguity set (El Ghaoui, Oks & Oustry 2003, Kang, Li & Li 2020), or a model-uncertain market-maker whose spreads trigger liquidity dry-ups (Routledge & Zin 2009). The measures here instead come from the equilibrium of competitive, capital-constrained banks, so the loss distribution is generated jointly with prices.

Equilibrium ambiguity, exit, and participation. The exit at the center of the model has been studied in equilibrium settings that stop short of computing a downside risk measure. Ambiguity-averse investors decline to participate in markets whose signals they cannot confidently interpret (Easley & O'Hara 2009), model uncertainty produces limited participation through a threshold at which uncertain investors withdraw (Cao, Wang & Zhang 2005), and maxmin preferences generate portfolio inertia (Dow & Werlang 1992) and ambiguity-driven pricing (Condie & Ganguli 2017, Anthropelos & Schneider 2024). Closest to the signal structure here, Illeditsch, Ganguli & Condie (2021) show that risk and ambiguity aversion together leave optimal portfolios unresponsive to public news that is worse than expected, so the equilibrium price does not reflect it, an inertia that vanishes when the news is idiosyncratic. The banks here go further than non-response. A large enough ambiguity shock drives them out of the market altogether, and the capital constraint on those who remain turns that exit into a fire sale. On the pricing-moment side, uncertainty aversion generates a sizable variance premium (Drechsler 2013, Miao, Wei & Zhou 2019). These papers supply the ingredients of the analysis here, namely endogenous exit, an ambiguity premium, and excess variance, but none derives value at risk, expected shortfall, or the shape of the loss

distribution as done here. Because the tail risk this exit creates is forward-looking, quantifying it ahead of time requires a model of the exit itself and of the balance sheets that must absorb it.

Ambiguity-driven panics, fire sales, and procyclical capital. Finally, the paper connects to work in which ambiguity is the source of discrete, panic-like events and to the capital consequences of the fire sales they produce. The mechanics of the model in this paper have similarities with each of these, although the specifics of the models and the questions each paper addresses differ. Dicks & Fulghieri (2019) make uncertainty aversion turn idiosyncratic shocks into systemic runs, Caballero & Krishnamurthy (2008) take Knightian uncertainty as the trigger for flight-to-quality episodes, and Caballero & Simsek (2013) generate fire sales from complexity that spikes perceived counterparty risk. Closest to the mechanism here, Anderson (2019) obtains ambiguity-driven freezes and fire sales in securitization markets, reproducing the timing and duration of the crisis freeze. That model explains the freeze as an outcome among the investors who hold a securitized asset: banks there are risk-neutral, and it is the purchasers of securitized assets who face differing degrees of ambiguity. The present paper places the same ambiguity-driven exit on a leveraged bank balance sheet governed by a risk-based capital requirement, so the depth of the fire sale is no longer a property of the freeze alone but is set by the scarcity of the remaining banks' equity and by the conservatism of the risk measure. Because of this balance-sheet structure the downside risk measures and the regulatory capital built on them become equilibrium objects, and pricing the possibility of a panic before it occurs yields the interim tail risk and its procyclicality. Where Anderson (2019) accounts for why a securitization market freezes, the object here is the bank capital a possible freeze requires, how the requirement feeds back on the freeze's depth, and how the underlying ambiguity can be read from banks' own loss provisioning.

Closest to the empirics here, Boyarchenko (2012) attributes the 2007–2008 widening of credit spreads to shifts in ambiguity about financial firms' payoffs. Where Boyarchenko (2012) measures ambiguity from the price of credit risk, here it is measured from banks' own provisioning. That a binding requirement deepens the fire sale is the discrete counterpart of the intermediary-capital channel, in which risk prices spike as net worth shrinks (He & Krishnamurthy 2013, Adrian & Boyarchenko 2012), and of the procyclicality of value-at-risk-based capital (Adrian & Shin 2014, Brunnermeier & Pedersen 2009, Danielsson, Shin & Zigrand 2012). The contribution here is to trace that procyclicality to the discontinuity of value at risk in the panic probability, which expected shortfall removes.

The remainder of the paper proceeds as follows. Section 2 develops the model in four parts. Section 2.1 lays out the setup with the bank balance sheet and a risk-based capital

requirement. Section 2.2 derives the equilibrium and the panic it produces, including the capital-constrained fire sale and the re-entry of ambiguity-averse capital. Section 2.3 develops the two-period extension that prices the possibility of a panic and yields the interim price distribution, value at risk, and expected shortfall. Section 2.4 draws out the implications for leverage, capital, and the procyclicality of the requirement. Section 3 takes the model to data. Its centerpiece is the distribution of bank asset returns in the U.S. Call Reports and a new bank-level measure of ambiguity, built from the cross-sectional dispersion of banks' loss provisioning, that forecasts both the deepening of the left tail one to four quarters ahead and, in a bank-level panel, the depletion of regulatory capital. Section 4 concludes. Appendix A collects the proofs, and Appendix B documents the construction of the bank-data ambiguity measure.

2 The model

2.1 Setup

Consider a market for a portfolio of risky bank assets (loans and credit securities), referred to as the bank's portfolio. The market operates at times 0 and 1, and the assets in the bank's portfolio mature at time 2. The bank's assets are marked to market at time 1. Two assets are available to hold: a risk-free asset paying a normalized net return of 0 at maturity, and the risky portfolio, whose value at maturity is $v \sim N(\mu_0, \rho_0^{-1})$, where ρ_0 is its ex ante precision. The market is populated by a unit mass of banks. A fraction x_e of these banks form beliefs in the standard Bayesian way, while a fraction $x_a = 1 - x_e$ are at risk of panic when some signals about the value of banks' portfolios are received. These banks are modeled as perceiving the public signal about banks' portfolio value as ambiguous, in a manner made precise below. At the interim date (period 1) every bank observes a common public signal about the value of the portfolio¹ given by $s = v + \epsilon$, with $\epsilon \sim N(0, \rho_\epsilon^{-1})$, and rebalances accordingly.²

Bayesian updating for a bank that does not perceive ambiguity leads to a posterior $v|s \sim N(m(s), \rho_I^{-1})$ where $\rho_I = \rho_0 + \rho_\epsilon$ and $m(s) = \mu_0 + \frac{\rho_\epsilon}{\rho_I}(s - \mu_0)$. Ambiguity-averse banks instead believe that the noise term ϵ is drawn from a set of normal distributions with means $\mu \in [-\delta, \delta]$ and precision ρ_ϵ . Updating each of these priors leads to a set of posteriors

¹Examples of potential signals that could match the qualitative properties studied in the paper include a new reading of a house-price index, a delinquency print, or a ratings action, *inter alia*.

²The term *bank* covers any leveraged holder of these assets: commercial banks, but also the broker-dealers, structured vehicles, and other intermediaries that carry credit on balance sheet and fund it with runnable liabilities.

$v|s \sim N(m_z(s), \rho_I^{-1})$ with $m_z(s) = m(s) - \frac{\rho_\epsilon}{\rho_I} z$ for $z \in [-\delta, \delta]$: the set of normal distributions with means in $[m(s) - \Delta, m(s) + \Delta]$ and precision ρ_I , with Δ and σ_I^2 defined by

$$\Delta = \delta \frac{\rho_\epsilon}{\rho_I}, \quad \sigma_I^2 = \frac{1}{\rho_I}. \quad (1)$$

2.1.1 Balance sheet, capital, and leverage

Each bank of type $i \in \{e, a\}$ enters the trading date with a legacy portfolio and with *equity capital* $W_i > 0$, a primitive of the model. Holdings are normalized so that the supply of the risky portfolio is one unit. Let $\mathcal{W}_e \equiv x_e W_e$ denote the aggregate equity of expected-utility banks. A bank that chooses to hold θ units of the portfolio, worth $p\theta$, funds them with its equity W_i and with liabilities D , where

$$D = p\theta - W_i, \quad \lambda_i \equiv \frac{p\theta}{W_i}, \quad (2)$$

so that assets equal equity plus liabilities, $p\theta = W_i + D$, and λ_i is leverage. Liabilities represent any risk-free claim on the bank (insured deposits, senior debt, or repo funding) supplied elastically at the normalized net rate of 0.³ When θ exceeds the equity-financed position the bank issues liabilities ($D > 0$, $\lambda_i > 1$). When it holds less, $D < 0$ and the bank places the difference in the risk-free asset.

Because liabilities are risk free and senior to equity, equity absorbs all fluctuation in asset value. Assets pay θv at maturity and liabilities are repaid at par, so terminal equity is

$$\widetilde{W}_i = \theta v - D = W_i + \theta(v - p). \quad (3)$$

2.1.2 Bank objectives and capital provision

Banks that do not perceive ambiguity in the signal and who will not be prone to panic in equilibrium have mean-variance preferences over terminal equity,

$$U = \mathbb{E}\widetilde{W}_i - \frac{\gamma}{2} \text{Var}(\widetilde{W}_i) = W_i + \theta(m(s) - p) - \frac{\gamma}{2} \theta^2 \sigma_I^2, \quad (4)$$

where γ is the risk-aversion coefficient. Capital enters additively, so the optimal position and indirect utility are, absent the constraint introduced below,

$$\theta_e = R_0^{-1}[m(s) - p(s)], \quad U_e(p, s) = W_e + p(s) + \frac{1}{2} R_0^{-1}[m(s) - p(s)]^2. \quad (5)$$

³The analysis does not depend on the specifics of the funding source for intermediaries, only that it is a fixed claim that is senior to equity.

Following the benchmark in which all banks are identical Bayesians and the market clears at $p(s) = m(s) - \gamma\sigma_I^2$, the *fundamental risk premium* is

$$R_0 = \gamma\sigma_I^2. \quad (6)$$

The quantity R_0 is the amount of risk in the market, normalized by risk aversion and expressed in units of expected future wealth.

Banks that in equilibrium may be subject to panic are modeled as having equity holders who are ambiguity averse. These banks evaluate terminal equity under the worst-case belief,

$$U = W_a + p(s) + \min_{z \in [-\delta, \delta]} \left\{ \theta(m_z(s) - p(s)) - \frac{\gamma}{2}\theta^2\sigma_I^2 \right\}, \quad (7)$$

which yields the following demand for the risky bank portfolio,

$$\theta_a = \begin{cases} R_0^{-1}(m(s) - \Delta - p(s)) & m(s) - p(s) > \Delta, \\ 0 & -\Delta \leq m(s) - p(s) \leq \Delta, \\ R_0^{-1}(m(s) + \Delta - p(s)) & m(s) - p(s) \leq -\Delta. \end{cases} \quad (8)$$

Banks with decision rules of this type have an inaction region. When ambiguity is high relative to the current premium available for bank capital provision, $|m(s) - p(s)| \leq \Delta$, the ambiguity-averse banks exit the market and hold no position. When the premium is high enough, $|m(s) - p(s)| > \Delta$, they hold a long position in the portfolio.

If ambiguity-averse banks are to hold a long position, the expected payoff must be positive under the worst-case belief. They require the additional *ambiguity premium* Δ to stay in the market (Gilboa & Schmeidler 1989, Dow & Werlang 1992).

2.1.3 A general risk-based capital requirement

A risk-based capital requirement obliges a bank to hold equity that is at least as large as a multiple $\kappa \in (0, 1]$ of the downside risk of its position, measured by a risk measure ρ_α (e.g. value at risk under Basel II or expected shortfall under the Fundamental Review of the Trading Book (Basel Committee on Banking Supervision 2019)). The requirement is modeled as

$$W_i \geq \kappa \rho_\alpha(\theta L) = \kappa \theta \rho_\alpha(L), \quad (9)$$

where L is the loss per unit of the portfolio over the position's holding horizon. The risk measure is assumed positively homogeneous (both VaR and ES are), which gives the second equality for $\theta \geq 0$. Equivalently, the requirement caps the position, the leverage, and the

capital held per unit at

$$\bar{\theta}_{\alpha,i} = \frac{W_i}{\kappa \rho_{\alpha}(L)}, \quad \bar{\lambda}_{\alpha} = \frac{p}{\kappa \rho_{\alpha}(L)}, \quad K_{\alpha} \equiv \kappa \rho_{\alpha}(L). \quad (10)$$

The random variable L governing loss over the relevant holding horizon is $L = p - v$ in the interim period (Section 2.2), since banks invest at time 1 and hold to maturity. In the two-period model (Section 2.3) the requirement on the time-0 book is set against the interim mark-to-market loss $L = p_0 - p_1$. Constrained demand for a Bayesian bank is the smaller of the frictionless demand and the cap,

$$\theta_e^c = \min\{R_0^{-1}[m(s) - p], \bar{\theta}_{\alpha,e}\}. \quad (11)$$

In the interim period (time 1), losses over the next period are Gaussian. For a Gaussian loss $X \sim N(\mu_X, \sigma_X^2)$, both risk measures take the form

$$\rho_{\alpha}(X) = \mu_X + \sigma_X c_{\alpha}, \quad c_{\alpha} = \begin{cases} z_{\alpha} \equiv \Phi^{-1}(\alpha) & \text{(value at risk),} \\ \frac{\phi(z_{\alpha})}{1 - \alpha} & \text{(expected shortfall),} \end{cases} \quad (12)$$

so a single constant $c_{\alpha} > 0$ indexes the *conservatism* of the measure, with $c_{\alpha}^{\text{ES}} > c_{\alpha}^{\text{VaR}}$ at any common α . Value at risk and expected shortfall are the two natural anchors of this dial rather than rival measures to be adjudicated. Writing results in terms of c_{α} keeps them in one set of formulas and isolates where the conservatism of the measure matters, in the capital results of Section 2.4.

Trading book, banking book, and what transfers. The value-at-risk and expected-shortfall apparatus above originates in the market-risk framework for the trading book, namely value at risk under Basel II and expected shortfall under the Fundamental Review of the Trading Book (Basel Committee on Banking Supervision 2019). The assets the empirics measure (the return on bank assets, loan-loss provisions) instead live in the banking book, where regulatory capital is set by credit-risk rules and losses are governed by provisioning standards (the incurred-loss allowance and, since 2020, CECL). The market-risk measures serve here as the sharpest available lens on a mechanism that is not specific to the trading book, and three features make the transfer legitimate. First, the object the requirement is written against is the same in both regimes: a distribution of a mark-to-market or fair-value loss over a holding horizon, summarized by a downside quantile (value at risk) or a tail average (expected shortfall). The $\rho_{\alpha}(L)$ used here is such a functional of the interim

loss $L = p_0 - p_1$, and nothing in the equilibrium derivation uses a trading-book-specific convention. Second, the economics the paper isolates is a property of the loss distribution and the measure, not of the regime. A coherent tail measure loads on a possible panic continuously and ahead of the event, while a quantile measure is flat and then jumps, and this holds wherever capital is a monotone function of a tail risk statistic. That includes the expected-loss logic of CECL, under which forward-looking loss estimates must be booked before losses are realized. Third, the bank-data ambiguity measure here is built from the very banking-book object (loan-loss provisioning) that CECL makes forward-looking, so the empirical and theoretical tail objects are aligned even though the calibration figures in Section 2.3.6 are expressed in the market-risk idiom. The model therefore reads as a statement about how any risk-sensitive capital rule, whether trading-book expected shortfall or banking-book expected-loss provisioning, responds to a possible ambiguity-driven panic. Points where the mapping is interpretive rather than literal are flagged as they arise.

2.2 Equilibrium and the panic

The market clears when total demand equals supply,

$$x_e \theta_e(p) + x_a \theta_a(p) = 1. \quad (13)$$

Because $L = p - v$ with $v \mid s \sim N(m(s), \sigma_I^2)$, the per-unit capital charge is

$$\rho_\alpha(p - v) = \sigma_I c_\alpha - (m(s) - p), \quad (14)$$

which *falls* as the price falls: a portfolio bought at a larger risk premium $m(s) - p$ carries a smaller regulatory loss, because the expected gain offsets more of the tail. This price dependence lets the market clear even when capital is scarce. There are four possible states, ordered by the realized ambiguity Δ and by whether the expected-utility sector is well enough capitalized to absorb the portfolio when the ambiguity-averse banks exit: no panic, an unconstrained panic, a capital-constrained panic, and a capital-constrained panic in which the ambiguity-averse banks re-enter. The last state arises because the price dependence just noted lets a fire sale deepen until it overshoots the ambiguity premium.

Regime 1: no panic. When ambiguity is low enough that ambiguity-averse banks are willing to provision capital, both bank types hold interior positions and (13) gives

$$p^\circ(s) = m(s) - \Delta x_a - R_0, \quad \theta_e^\circ = 1 + \frac{\Delta x_a}{R_0}, \quad \theta_a^\circ = 1 - \frac{\Delta x_e}{R_0}. \quad (15)$$

The equilibrium price is a wealth-weighted blend of the ambiguity premium and the risk premium. Expected-utility banks hold more than their unit endowment and are thus levered. The constraint is slack provided their capital covers the position,

$$\theta_e^\circ \leq \bar{\theta}_{\alpha,e} \iff W_e \geq \kappa \left(1 + \frac{\Delta x_a}{R_0}\right) [\sigma_I c_\alpha - (\Delta x_a + R_0)], \quad (16)$$

which holds for all but very thinly capitalized banks. This is the normal-times case.

Regime 2: unconstrained panic. When ambiguity is high the ambiguity-averse banks exit ($\theta_a = 0$). If expected-utility capital is ample they absorb the whole supply frictionlessly, $\theta_e = 1/x_e$, and

$$p^*(s) = m(s) - y(x_e), \quad y(x_e) \equiv \frac{R_0}{x_e}. \quad (17)$$

The fundamental premium is scaled up by the inverse of the expected-utility share because there is less risk sharing. That is, x_e measures the balance-sheet capacity that must absorb the portfolio when others exit, so $y(x_e)$ is the premium that scarce capacity commands. This is the same force that makes risk prices spike as the marginal intermediary's capacity shrinks in the intermediary asset-pricing literature (He & Krishnamurthy 2013, Adrian & Boyarchenko 2012). This regime holds when the constraint is slack at p^* ,

$$\frac{1}{x_e} \leq \bar{\theta}_{\alpha,e} \iff x_e W_e \geq \kappa [\sigma_I c_\alpha - y(x_e)]. \quad (18)$$

Regime 3: capital-constrained panic. If (18) fails, so that aggregate expected-utility capital is scarce, $x_e W_e < \kappa [\sigma_I c_\alpha - y(x_e)]$, the remaining banks cannot absorb the portfolio at p^* without violating (9). The price falls further. By (14) the fall shrinks the capital charge and expands capacity until the constraint binds *with* market clearing. Setting $\theta_e = 1/x_e$ in (9) at equality, $\rho_\alpha(p - v) = x_e W_e / \kappa$, and substituting (14),

$$p^{\text{cap}}(s) = m(s) - y_\alpha^{\text{cap}}, \quad y_\alpha^{\text{cap}} = \sigma_I c_\alpha - \frac{x_e W_e}{\kappa}, \quad (19)$$

provided the ambiguity-averse banks still prefer to stay out at this lower price, $\Delta \geq y_\alpha^{\text{cap}}$. The constrained premium is decreasing in expected-utility capital, decreasing in the leniency κ^{-1} of the rule, and increasing in the stringency c_α of the risk measure. As $W_e \rightarrow 0$ this premium approaches its ceiling $\sigma_I c_\alpha$, the largest premium at which the regulatory loss on the position is nonpositive.

The unified panic premium. Regimes 2 and 3 combine into one expression for the amount by which the panic price sits below the posterior mean,

$$y_\alpha^* = \max \left\{ \underbrace{\frac{R_0}{x_e}}_{\text{risk sharing}}, \underbrace{\sigma_I c_\alpha - \frac{x_e W_e}{\kappa}}_{\text{capital scarcity}} \right\}, \quad p^{\text{panic}}(s) = m(s) - y_\alpha^*. \quad (20)$$

The first term is the risk-sharing premium, set by how few banks remain. The second is a capital-scarcity premium, set by how little equity they have and how conservative the risk measure is. The panic is “unconstrained” when the former binds and “capital-constrained” when the latter does. In the normal-times benchmark of ample capital, $y_\alpha^* = y(x_e)$, and every object derived below reduces to the frictionless case.

For an equilibrium in which both types remain in the market, one verifies as usual that neither type takes a short position while the other is long, so market clearing with interior demands delivers (15). This is an equilibrium if and only if ambiguity-averse banks choose a positive position, $y_\alpha^* > \Delta$. The panic equilibrium, in which they hold zero, requires $y_\alpha^* \leq \Delta$. These two conditions partition the parameter space.

Result 1 (Exit). *Ambiguity-averse banks exit the market if and only if the premium available when they are absent falls short of the ambiguity premium they require to stay, $y_\alpha^* \leq \Delta$. Exit is more likely when perceived ambiguity Δ is high, the fundamental premium R_0 is low, the expected-utility share x_e is high, and, through the capital-scarcity term, when expected-utility capital $x_e W_e$ is high, since well-capitalized banks can absorb the portfolio at a smaller premium, making the market less attractive to those who would otherwise be paid to stay.*

Intuitively, when a small, well-capitalized set of banks perceive no ambiguity while others perceive a great deal in a precise signal, the ambiguity-averse banks are spooked and exit, selling their holdings to the banks that remain. This exit is the market analogue of a funding run: the ambiguity-averse banks withdraw and the portfolio is absorbed by the rest only at a fire-sale mark.

The capital comparative static in Result 1 runs opposite to Result 2 and could be misread. Result 1 holds Δ fixed and asks when exit *is triggered*: better-capitalized buyers hold the portfolio at a smaller premium, so the market pays the ambiguity-averse banks less to stay and the exit threshold y_α^* falls. Result 2 concerns *severity*: conditional on a constrained panic, scarcer capital deepens the fire sale. The two are consistent. Ample capital makes exit less likely to be triggered, yet a panic triggered under scarce capital is more severe. The reading is not that capital is destabilizing, but that the same equity prevents exit and cushions it, entering the two margins with opposite signs.

Result 2 (Capital scarcity amplifies the panic). *The panic price $m(s) - y_\alpha^*$ is weakly below the frictionless panic price $m(s) - y(x_e)$, strictly so when expected-utility capital is scarce. In that region the fire-sale mark deepens as capital falls and as the measure becomes more stringent,*

$$\frac{\partial y_\alpha^*}{\partial W_e} = -\frac{x_e}{\kappa} < 0, \quad \frac{\partial y_\alpha^*}{\partial c_\alpha} = \sigma_I > 0.$$

Because $c_\alpha^{\text{ES}} > c_\alpha^{\text{VaR}}$, a binding expected-shortfall requirement produces a deeper constrained-panic price than value at risk at the same confidence level.

Regime 4: Panic where ambiguity-averse banks remain. The capital-constrained price (19) was derived on the premise that ambiguity-averse banks prefer to stay out at that price, $\Delta \geq y_\alpha^{\text{cap}}$. When expected-utility capital is scarce enough that this premise fails, the fire sale overshoots their reservation premium and they remain in the market.⁴ With expected-utility banks pinned at their capital limit and ambiguity-averse banks at an interior optimum, and writing $\pi \equiv m(s) - p$ for the premium, market clearing is

$$x_e \frac{W_e}{\kappa(\sigma_I c_\alpha - \pi)} + x_a R_0^{-1}(\pi - \Delta) = 1. \quad (21)$$

Result 3 (Re-entry). *Suppose expected-utility capital is scarce, $x_e W_e < \kappa(\sigma_I c_\alpha - R_0/x_e)$, so that $y_\alpha^{\text{cap}} > y(x_e)$. For realized ambiguity in the intermediate range*

$$y(x_e) \leq \Delta < y_\alpha^{\text{cap}} = \sigma_I c_\alpha - \frac{x_e W_e}{\kappa},$$

neither the interior equilibrium nor the full-exit panic exists: at the full-exit price $m(s) - y_\alpha^{\text{cap}}$ the premium exceeds Δ , so ambiguity-averse banks strictly prefer to hold a long position. The unique equilibrium has the expected-utility banks at their capital limit and the ambiguity-averse banks holding a positive interior position at the premium π^ , the unique root of (21) in $(\Delta, y_\alpha^{\text{cap}})$. Since $\pi^* > \Delta \geq y(x_e)$ the realized fire sale is deeper than the frictionless panic, and since $\pi^* < y_\alpha^{\text{cap}}$ the re-entering ambiguity-averse capital places a floor under it: the price stops short of the full-exit capital-constrained mark.*

The clearing condition (21) has no closed form, but it can be evaluated numerically. Collecting the four regimes, the realized time-1 premium as a function of the ambiguity

⁴Regimes 3 and 4 are non-empty only if $\sigma_I c_\alpha > R_0/x_e$. Otherwise the capital charge at the unconstrained panic price, $\sigma_I c_\alpha - y(x_e)$, is non-positive, the requirement (9) is slack at every capital level, and every panic is unconstrained (Regime 2). The condition does not involve κ , and since $c_\alpha^{\text{ES}} > c_\alpha^{\text{VaR}}$ it can hold under an expected-shortfall rule while failing under value at risk.

draw is

$$\pi(\Delta) = m(s) - p_1 = \begin{cases} R_0 + x_a \Delta, & \Delta < y(x_e) \text{ (Regime 1, no panic),} \\ y(x_e), & \Delta \geq y(x_e), \text{ ample (Regime 2),} \\ \pi^*(\Delta) \in (\Delta, y_\alpha^{\text{cap}}), & y(x_e) \leq \Delta < y_\alpha^{\text{cap}}, \text{ scarce (Regime 4),} \\ y_\alpha^{\text{cap}}, & \Delta \geq y_\alpha^{\text{cap}}, \text{ scarce (Regime 3),} \end{cases} \quad (22)$$

plotted in Figure 2. The time-0 analysis of Section 2.3 integrates this premium across the four regimes.⁵ The four types are collected in Figure 1, in the plane of aggregate capital and ambiguity.




Figure 1: The four equilibrium regimes in the plane of aggregate expected-utility capital $\mathcal{W}_e = x_e W_e$ and perceived ambiguity Δ . The horizontal line $\Delta = y(x_e) = R_0/x_e$ is the panic trigger: below it both types hold interior positions (Regime 1), above it a panic occurs. When capital is ample, $\mathcal{W}_e \geq \mathcal{W}_e^* = \kappa[\sigma_I c_\alpha - y(x_e)]$, the panic is unconstrained at premium $y(x_e)$ (Regime 2). When it is scarce, the downward-sloping line $\Delta = y_\alpha^{\text{cap}} = \sigma_I c_\alpha - \mathcal{W}_e/\kappa$ splits the panic: above it the exiting banks stay fully out (Regime 3, price $m(s) - y_\alpha^{\text{cap}}$), and below it the fire sale overshoots their reservation premium and they re-enter (Regime 4). A stricter risk measure (larger c_α) raises the sloped line and enlarges the re-entry region.

⁵Regime 1 presumes the mild leverage of the no-panic position (15) is within capital, which holds except under extreme thin capitalization. See (16).




Figure 2: The realized time-1 premium $\pi(\Delta) = m(s) - p_1$ of (22) against the ambiguity draw Δ , holding the posterior mean fixed. All curves coincide on the no-panic region $\Delta < y(x_e)$, where $\pi(\Delta) = R_0 + x_a \Delta$. Above $y(x_e)$ the well-capitalized market (blue) settles at $y(x_e)$ (Regime 2). Under scarce capital the fire sale is deeper: the full-exit path (dashed) jumps to the ceiling y_α^{cap} , but for $y(x_e) \leq \Delta < y_\alpha^{\text{cap}}$ the ambiguity-averse banks re-enter and clear at $\pi^*(\Delta)$ (red, Regime 4), rising continuously to y_α^{cap} and capping the crash below the full-exit line. For $\Delta \geq y_\alpha^{\text{cap}}$ they stay out at y_α^{cap} (Regime 3). Parameters: $R_0 = 0.08$, $x_a = x_e = 0.5$ so $y(x_e) = 0.16$, $\sigma_I c_\alpha = 0.266$, $\kappa = 1$. The scarce curve is at $\mathcal{W}_e = 0.09$ ($y_\alpha^{\text{cap}} = 0.176$) and the ample curve at $\mathcal{W}_e = 0.14$.

Three features of this map are worth drawing out. First, the ambiguity premium Δ is both the *trigger* of the panic and the *backstop* that bounds its severity. Ambiguity-averse banks exit when the premium on offer falls short of Δ and re-enter when a deep enough fire sale lifts it back above Δ , so Δ is their reservation premium in both directions. Second, this phenomenon is *stabilizing*. The intermediary-capital channel of Result 2 says that scarcer equity deepens the fire sale. Regime 4 shows that the deepening is bounded, because once the price has fallen far enough the ambiguity-averse banks become the marginal buyers and clear the market at $\pi^* < y_\alpha^{\text{cap}}$. Their capital acts as a buyer of last resort that caps how far scarce intermediary capital can drive the price. Third, a more conservative risk measure makes re-entry *more likely*. The re-entry interval $[y(x_e), y_\alpha^{\text{cap}})$ widens in c_α , and $c_\alpha^{\text{ES}} > c_\alpha^{\text{VaR}}$, so an expected-shortfall requirement, by forcing a deeper constrained fire sale (Result 2), draws the ambiguity-averse banks back over a wider range of ambiguity than value at risk.

2.3 Valuing the possibility of panics before they happen

This section asks how banks price the possibility of a panic before it occurs. There are two decision periods $t \in \{0, 1\}$: origination at $t = 0$ and an interim mark-to-market date at $t = 1$, before final gains and losses are realized at maturity. Everyone begins at $t = 0$ endowed with one unit of the risky asset and equity capital W_i . At $t = 0$ the ambiguity Δ that will be in the signal at $t = 1$ is not known. Its CDF $F(\Delta)$ is common knowledge. After trading at $t = 0$, ambiguity is realized and the market proceeds as in Section 2.2 at $t = 1$. Ambiguity here is triggered by the arrival of a signal that is hard to interpret. Ambiguity-averse banks perceive no ambiguity at $t = 0$, before any signal has arrived, and become ambiguity-averse only at $t = 1$.⁶ This structure keeps the time-0 market tractable and, if anything, understates the ambiguity discount in the current price. Their time-0 utility is therefore the expectation, under the objective distribution of Δ , of their time-1 utility across the two regimes:

- the ambiguity level Δ is low enough that they remain, with probability $F(y_\alpha^*) = 1 - q^*$ (*no panic*).
- the ambiguity level is high enough that they exit, with probability $1 - F(y_\alpha^*) = q^*$ (*panic*).

2.3.1 Time-1 indirect utility and the time-0 price

A bank entering $t = 1$ holding θ_0 units and $(1 - \theta_0)p_0$ in risk-free wealth rebalances optimally at the time-1 price p_1 . Its indirect utility, for belief $\mathbb{E}_i[v]$, is

$$V_i(\theta_0, s, \Delta) = W_i + (1 - \theta_0)p_0 + \theta_0 p_1 + \frac{(\mathbb{E}_i[v] - p_1)^2}{2R_0}, \quad (23)$$

the frictionless expression plus the additive constant W_i . The trading benefit $\frac{1}{2R_0}(\mathbb{E}_i[v] - p_1)^2$ does not depend on the time-0 holding θ_0 , which affects utility only through the endowment-value term $\theta_0 p_1$. The time-1 price is $p_1 = m(s) - \pi(\Delta)$, with the realized premium $\pi(\Delta)$ of (22) spanning the four regimes. It is convenient to write it as

$$p_1 = m(s) - R_0 - x_a Z, \quad Z \equiv \frac{\pi(\Delta) - R_0}{x_a}, \quad (24)$$

⁶Such banks are therefore not dynamically consistent in the sense of Epstein & Schneider (2003): their preferences are expected utility at $t = 0$ but maxmin at $t = 1$, rather than a single recursive multiple-priors preference over a rectangular set of priors. A dynamically consistent ambiguity-averse bank would already carry the worst-case belief at $t = 0$, value the asset's future resale below $p_0 = \mathbb{E}[p_1]$, and pare back its time-0 holding accordingly. Because that recursive pessimism would only deepen the discount, the two-stage formulation adopted here is conservative in that sense.

where Z is the *realized-premium shock*, the amount of downside the ambiguity draw adds to the fundamental premium R_0 . When capital is ample it reduces to the censored ambiguity shock $Z = \min(\Delta, y(x_e))$: on the no-panic region $\pi(\Delta) = R_0 + x_a \Delta$ gives $Z = \Delta$, and in the unconstrained panic $\pi(\Delta) = y(x_e)$, $Z = (y(x_e) - R_0)/x_a$. When capital is scarce, Z tracks the re-entry path of Figure 2: it rises continuously through the re-entry region and is censored at $(y_\alpha^{\text{cap}} - R_0)/x_a$ once the ambiguity-averse banks are fully out. Everything below is written in terms of Z and so holds for either capital regime. The ample capital case, in which $Z = \min(\Delta, y(x_e))$, is characterized in closed form, while the scarce case is evaluated from (22) numerically.

Because V_i is affine in θ_0 and the trading benefit is independent of it, time-0 demand is a discontinuous (bang-bang) function of $\bar{p}_1 - p_0$, and the market must clear at $p_0 = \bar{p}_1 = \mathbb{E}[p_1] = \mu_0 - \mathbb{E}[\pi(\Delta)]$. In the ample-capital case, where the interim outcome is either no panic or full exit, this is the two-regime average

$$p_0 = (1 - q^*)(\mu_0 - \mathbb{E}[\Delta \mid \Delta \leq y_\alpha^*] x_a - R_0) + q^*(\mu_0 - y_\alpha^*), \quad (25)$$

while under scarce capital the panic branch is replaced by the re-entry-weighted average of (22).

2.3.2 The panic premium

If there were no possibility of ambiguity, the time-1 price would always be $m(s) - R_0$ and the time-0 price $p_0^{\text{no amb}} = \mu_0 - R_0$. Since $p_0 = \mu_0 - \mathbb{E}[\pi(\Delta)]$ and $\pi(\Delta) = R_0 + x_a Z$,

$$p_0 = \mu_0 - R_0 - x_a \mathbb{E}[Z], \quad (26)$$

so the panic premium relative to the no-ambiguity benchmark is

$$\Omega = p_0^{\text{no amb}} - p_0 = x_a \mathbb{E}[Z] \geq 0, \quad (27)$$

with $Z = \min(\Delta, y(x_e))$ under ample capital and the re-entry-adjusted shock under scarce capital. The panic premium is increasing in the panic-susceptible share x_a . When capital is scarce it is, in addition, increasing in the stringency c_α of the risk measure and decreasing in expected-utility capital W_e , both of which raise the censoring point y_α^{cap} and so the mean shock $\mathbb{E}[Z]$.

Part of the expected return on the bank portfolio is thus compensation specifically for bearing the risk of a future exit, and that compensation is larger when the banking sector is thinly capitalized. Writing $\bar{p}_1^\circ = \mu_0 - \mathbb{E}[\Delta \mid \Delta \leq y_\alpha^*] x_a - R_0$ as the expected price conditional

on a panic not occurring, and $\bar{p}_1^* = \mu_0 - y_\alpha^*$ as the expected price conditional on a panic, $p_0 = (1 - q^*)\bar{p}_1^\circ + q^*\bar{p}_1^*$ lies between them,

$$\bar{p}_1^* \leq p_0 \leq \bar{p}_1^\circ. \quad (28)$$

That is, the time-0 price incorporates the expected damage of a panic but not the full realized damage.

2.3.3 The role of the panic probability

Since the expected time-1 price in times of panic $\bar{p}_1^*$ is less than the expected time-1 price in normal times, ($\bar{p}_1^* \leq \bar{p}_1^\circ$), the time-0 price is decreasing in the panic probability q^* . As a panic becomes more likely, weight shifts to the lower panic-regime price. Conditional on no panic the realized price exceeds p_0 ,

$$p_1^\circ - p_0 = q^*(\bar{p}_1^\circ - \bar{p}_1^*) > 0,$$

so a bank that holds from 0 to 1 earns a positive expected capital gain when no panic occurs. The positive pre-event expected return parallels the pre-announcement return premium in Lucca & Moench (2015). The gap between the crash the time-0 price anticipates and the milder movement usually realized parallels the variance risk premium (Carr & Wu 2009, Bollerslev, Tauchen & Zhou 2009), attributed to rarely realized tail events (Bollerslev & Todorov 2011) and to ambiguity aversion (Miao et al. 2019).

2.3.4 The distribution of time-1 prices

Since $s = v + \epsilon$ is a sum of independent normals, $m(s) \sim N(\mu_0, \sigma_m^2)$ with

$$\sigma_m^2 \equiv \text{Var}(m(s)) = \frac{1}{\rho_0} - \frac{1}{\rho_I} = \frac{\rho_\epsilon}{\rho_0 \rho_I}. \quad (29)$$

Writing the time-1 price (24) as $p_1 = m(s) - R_0 - x_a Z$ with the realized-premium shock $Z = (\pi(\Delta) - R_0)/x_a$ and noting $m(s) \perp Z$, the unconditional density is a location mixture of normals,

$$f_{p_1}(p) = \int \phi_N(p; \mu_0 - x_a z - R_0, \sigma_m^2) dF_Z(z). \quad (30)$$

In the *ample-capital* case $Z = \min(\Delta, y(x_e))$, so F_Z has a density on $[0, y(x_e))$ (the no-panic set of outcomes) and a mass $q^* = 1 - F(y(x_e))$ at $y(x_e)$ (the panic). In this case, (30) splits into the no-panic mixture plus a normal centered at the panic price $\mu_0 - y(x_e)$. Under scarce

capital, the same expression holds with F_Z following the re-entry path of Figure 2. Here, the atomic mass moves out to the censoring point $(y_\alpha^{\text{cap}} - R_0)/x_a$, but the re-entry region spreads part of that mass continuously below it, so the left tail is heavier than in the ample capital case. The mean is $\mathbb{E}[p_1] = p_0$, confirming (26), and the variance decomposes into a fundamental and an ambiguity component,

$$\text{Var}(p_1) = \sigma_m^2 + x_a^2 \text{Var}(Z), \quad \text{Var}(Z) = 2 \int_0^{y_\alpha^*} z \bar{F}(z) dz - \left(\int_0^{y_\alpha^*} \bar{F}(z) dz \right)^2, \quad (31)$$

where $\bar{F} = 1 - F$. The ambiguity component vanishes when ambiguity is absent and is non-diversifiable: it originates in the exit of one's counterparties. Writing $p_1 - p_0 = A - x_a B$ with $A = m(s) - \mu_0$ symmetric and $B = Z - \mathbb{E}[Z]$,

$$\mathbb{E}[(p_1 - p_0)^3] = -x_a^3 \mu_3(Z), \quad (32)$$

so the price distribution is left-skewed iff Z is right-skewed, which holds whenever a panic is possible: the point mass q^* at the boundary y_α^* pushes the right tail of Z out. An analogous fourth-moment calculation gives positive excess kurtosis whenever Z does. Because the boundary is y_α^* , *a more thinly capitalized banking sector, or a stricter risk measure, moves the censoring point out and thereby deepens the left skew and fattens the left tail of the interim price distribution.* This comparative static is absent from the frictionless model. The left skewness and excess kurtosis reproduce well-documented features of returns. The panic plays the role of a rare disaster (Barro 2006, Gabaix 2012) whose threshold is endogenous, since it occurs iff $\Delta > y_\alpha^*$ and both its probability and its severity depend on market composition, capital, and the risk measure.

2.3.5 Return distribution with exponential ambiguity

This subsection derives the return distribution under the assumption of exponential ambiguity in order to obtain tractable closed-form expressions for the panic probability, panic premium, and return distribution. Specializing to $\Delta \sim \text{Exp}(\psi)$, with $\mathbb{E}[\Delta] = 1/\psi$, yields closed forms in the ample-capital baseline, where $Z = \min(\Delta, y_\alpha^*)$ with $y_\alpha^* = y(x_e)$ and no re-entry. In the capital-scarce case, the corresponding quantities are evaluated numerically from (22) and reported in Figures 4–5 and 6 and Table 1. In the baseline the panic probability is

$$q^* = e^{-\psi y_\alpha^*}, \quad (33)$$

so scarce capital or a stricter measure (larger y_α^*) *lowers* the panic probability. The deeper the fire sale the market can reach, the more attractive it is for ambiguity-averse banks to stay. The panic, when it comes, is then more severe. The panic premium and time-0 price are

$$\Omega = \frac{x_a}{\psi} (1 - e^{-\psi y_\alpha^*}), \quad p_0 = \mu_0 - R_0 - \frac{x_a}{\psi} (1 - e^{-\psi y_\alpha^*}). \quad (34)$$

The total return $r = v - p_0$, measured from purchase at time 0 to terminal payoff, is normal,

$$r \sim N(R_0 + \Omega, \rho_0^{-1}), \quad \mathbb{E}[r] = \underbrace{R_0}_{\text{fundamental}} + \underbrace{\Omega}_{\text{panic premium}}. \quad (35)$$

Here, the expected return is increasing in x_a , the fraction of banks who are susceptible to panic, and in expected ambiguity $1/\psi$. When the panic is capital-constrained, it is larger when the sector is thinly capitalized or the binding requirement is conservative. In the limit as $\psi \rightarrow \infty$ (no ambiguity) $\mathbb{E}[r] \rightarrow R_0$. Likewise, as $\psi \rightarrow 0$ (panic nearly certain) the expected return approaches $R_0 + x_a y_\alpha^*$. The variance of the total return is ρ_0^{-1} , independent of the ambiguity parameters. The non-normal features are a property of the *interim* price return, generated by the trading dynamics, not of the terminal payoff.

Figure 3 plots the panic premium as a fraction of total expected return, $\Omega/(R_0 + \Omega)$, against ψ and x_a , for both the ample-capital baseline ($y_\alpha^* = y(x_e)$) and the capital-scarce case with Regime 4 re-entry (at the reference $\mathcal{W}_e = 0.09$). The panic premium is a larger share when expected ambiguity is high (low ψ) and when more banks are panic-susceptible (high x_a). Scarce capital raises the share wherever the capital constraint binds, because the deeper fire sale a thinly capitalized market can reach lifts the realized panic premium and hence Ω . It does so only there, however: once x_a is large enough that the risk-sharing premium $y(x_e)$ already exceeds the constrained level y_α^{cap} , capital is effectively ample and the two surfaces coincide.




Figure 3: Panic premium as a fraction of total expected return, $\Omega/(R_0 + \Omega)$, against the ambiguity parameter ψ and the panic-susceptible share x_a . The grey surface is the ample-capital baseline, and the red surface is the capital-scarce case with Regime 4 re-entry ($\mathcal{W}_e = 0.09$). The panic premium is a larger share when expected ambiguity is higher (lower ψ) and more banks are susceptible (higher x_a). Re-entry raises the share where the capital constraint binds, at high ambiguity and low-to-moderate x_a , with the deeper fire sale lifting the realized premium. It coincides with the baseline once $y(x_e)$ exceeds the constrained level y_α^{cap} and capital is effectively ample.

2.3.6 Expected shortfall and value at risk

The previous section used the capital requirements at $t = 1$ to determine the panic threshold y_α^* and the realized premium $\pi(\Delta)$. The question now is how a risk manager at $t = 0$ would measure the downside risk of holding the asset from $t = 0$ to $t = 1$, before the ambiguity shock is realized.

Consider a bank that holds one unit of the risky portfolio from $t = 0$ to $t = 1$, purchased

at p_0 . The interim mark-to-market loss is

$$L \equiv p_0 - p_1 = x_a(Z - \mathbb{E}[Z]) - A, \quad A = m(s) - \mu_0 \sim N(0, \sigma_m^2). \quad (36)$$

The random variable A is mean zero because $p_0 = \mathbb{E}[p_1]$. Downside risk therefore measures the dispersion of p_1 around a price that already embeds the anticipated panic. Conditioning on the ambiguity shock, VaR_α solves

$$(1 - q^*) \mathbb{E} \left[\Phi \left(\frac{\text{VaR}_\alpha + \Omega - x_a \Delta}{\sigma_m} \right) \mid \Delta \leq y_\alpha^* \right] + q^* \Phi \left(\frac{\text{VaR}_\alpha + \Omega - x_a y_\alpha^*}{\sigma_m} \right) = \alpha. \quad (37)$$

To understand the role of ambiguity, consider two benchmarks and how ambiguity relates to them. If ambiguity is absent the loss is Gaussian and $\text{VaR}_\alpha^0 = \sigma_m z_\alpha$. A risk manager who acknowledges the ambiguity-driven variance but still assumes normality would instead report the variance-matched $\text{VaR}_\alpha^N = z_\alpha \sqrt{\sigma_m^2 + x_a^2 \text{Var}(Z)}$. The behavior of the true value at risk (37) changes at $\alpha^\dagger = 1 - q^*$. For $\alpha < \alpha^\dagger$ the quantile lies in the no-panic region and value at risk is close to what it would be if the distribution were normal, while for $\alpha > \alpha^\dagger$ it falls inside the panic regime and, when the regimes are well separated,

$$\text{VaR}_\alpha \approx \underbrace{x_a(y_\alpha^* - \mathbb{E}[Z])}_{\text{expected panic drop}} + \sigma_m \Phi^{-1} \left(\frac{\alpha - (1 - q^*)}{q^*} \right), \quad (38)$$

with leading term $x_a(y_\alpha^* - \mathbb{E}[Z]) = y_\alpha^* - R_0 - \Omega$, the price decline net of the panic premium already in p_0 . Holding α fixed, value at risk is governed by fundamental risk while $q^* < 1 - \alpha$ and jumps by approximately the panic drop once q^* crosses $1 - \alpha$. It is non-monotonic in q^* , receding as $q^* \rightarrow 1$ because a near-certain panic is fully anticipated in p_0 .

Comparing this true value at risk with the variance-matched normal benchmark isolates the error a Gaussian model makes in the tail.

Result 4 (Gaussian value at risk understates the tail). *In the panic-dominated regime $\alpha > 1 - q^*$, the true value at risk (38) exceeds the variance-matched normal $\text{VaR}_\alpha^N = z_\alpha \sqrt{\sigma_m^2 + x_a^2 \text{Var}(Z)}$ by*

$$\text{VaR}_\alpha - \text{VaR}_\alpha^N = \underbrace{x_a(y_\alpha^* - \mathbb{E}[Z])}_{\text{panic drop}} - z_\alpha \sqrt{\sigma_m^2 + x_a^2 \text{Var}(Z)} + \sigma_m \Phi^{-1} \left(\frac{\alpha - (1 - q^*)}{q^*} \right), \quad (39)$$

which in the well-separated limit $\sigma_m \rightarrow 0$ reduces to

$$\text{VaR}_\alpha - \text{VaR}_\alpha^N = x_a \left[(y_\alpha^* - \mathbb{E}[Z]) - z_\alpha \sqrt{\text{Var}(Z)} \right]. \quad (40)$$

The gap is positive, meaning the normal model understates the true value at risk, whenever the panic drop $y_\alpha^* - \mathbb{E}[Z]$ exceeds $z_\alpha \sqrt{\text{Var}(Z)}$, that is, when the panic is a small-to-moderate-probability event and the loss is highly skewed. As the panic becomes near-certain, $q^* \rightarrow 1$, the drop is priced into p_0 and $\mathbb{E}[Z] \rightarrow y_\alpha^*$, so the gap narrows, crosses zero, and turns negative.

The variance-matched normal is therefore most reassuring where it is most wrong: in the panic-dominated tail, where the skewness it discards is largest. Section 2.3.7 scales this shortfall into the capital a normal model misses and gives the analogous gap for expected shortfall.

Expected shortfall, the average loss in the worst $1 - \alpha$ of outcomes,

$$\text{ES}_\alpha = \mathbb{E}[L \mid L \geq \text{VaR}_\alpha] = \frac{1}{1 - \alpha} \mathbb{E}[L \mathbf{1}\{L \geq \text{VaR}_\alpha\}], \quad (41)$$

behaves very differently. Because the panic loss is centered at the panic drop, above the no-panic loss for any common fundamental draw, the panic lies inside the tail that expected shortfall averages, and

$$\text{ES}_\alpha \approx \frac{q^*}{1 - \alpha} \underbrace{x_a(y_\alpha^* - \mathbb{E}[Z])}_{\text{panic drop}} + \left(1 - \frac{q^*}{1 - \alpha}\right) \underbrace{\text{ES}_\alpha^\circ}_{\text{fundamental tail}}. \quad (42)$$

The panic enters with weight $q^*/(1 - \alpha)$ at *every* confidence level a panic can reach. Expected shortfall loads on the panic as soon as one is possible and responds continuously to q^* . Over the range in which a panic is plausible but not yet certain, it is increasing in the panic probability, in the panic's severity y_α^* , and in the susceptible share x_a . Under $\Delta \sim \text{Exp}(\psi)$ in the well-separated limit, expected shortfall is closed-form,

$$\text{ES}_\alpha = \begin{cases} \frac{x_a}{\psi} \left[-\ln(1 - \alpha) - \frac{\alpha q^*}{1 - \alpha} \right], & \alpha < 1 - q^*, \\ x_a(y_\alpha^* - \mathbb{E}[Z]) = y_\alpha^* - R_0 - \Omega, & \alpha \geq 1 - q^*, \end{cases} \quad (43)$$

with $q^* = e^{-\psi y_\alpha^*}$ and $\mathbb{E}[Z] = (1 - q^*)/\psi$. The two branches agree at $\alpha = 1 - q^*$, so expected shortfall is continuous in the confidence level, unlike value at risk, which jumps there. Two channels are at work: a frequency channel, in which more weight falls on the panic outcome as q^* rises, and an anticipation channel, in which the panic drop shrinks as $\mathbb{E}[Z] \rightarrow y_\alpha^*$ when a panic becomes near-certain and fully priced. Together they make expected shortfall hump-shaped in the panic probability, and hence in ψ . It rises from the Gaussian floor as ambiguity builds, peaks at an intermediate panic probability, and recedes only in the limit where the panic is all but certain. Figures 4 and 5 display both measures at $\alpha = 0.99$ for

the ample-capital baseline and for a capital-scarce market with re-entry, computed from the exact four-regime loss (22). Scarce capital lifts both curves, because the deeper fire sale is a heavier interim tail, while re-entry keeps them below what the full-exit approximation would report. Figure 6 traces the same tail risk against the level of capital.⁷ This continuity of the tail average against the jump of the quantile is a property of the equilibrium loss distribution, not of any regulatory convention, though it aligns with the move to expected shortfall under the Fundamental Review of the Trading Book (Basel Committee on Banking Supervision 2019).

⁷With fundamental noise retained ($\sigma_m > 0$), expected shortfall reduces to a one-dimensional integral of Gaussian partial expectations against the ambiguity density, $ES_\alpha = (1 - \alpha)^{-1} \mathbb{E}_Z[\mu(Z) \Phi(\frac{\mu(Z) - \text{VaR}_\alpha}{\sigma_m}) + \sigma_m \phi(\frac{\text{VaR}_\alpha - \mu(Z)}{\sigma_m})]$ with $\mu(z) = x_a(z - \mathbb{E}[Z])$, which the figures compute.




Figure 4: Expected shortfall (solid) and value at risk (dashed) of the interim position at $\alpha = 0.99$ as a function of the onset panic probability $q^* = \mathbb{P}(\Delta \geq y(x_e))$ (log scale), from the exact four-regime loss distribution (22). Blue curves are the ample-capital baseline, and red curves are a capital-scarce market ($\mathcal{W}_e = 0.09$) in which Regime 4 re-entry is active. Scarce capital raises both measures, since a deeper fire sale is a heavier interim tail, while re-entry holds them below the full-exit approximation. The faint horizontal line is the no-ambiguity Gaussian floor, and the vertical dotted line marks $q^* = 1 - \alpha$. Both measures are hump-shaped and recede as $q^* \rightarrow 1$, the panic becoming fully anticipated in p_0 . Parameters: $R_0 = 0.08$, $x_a = x_e = 0.5$, $\sigma_{IC\alpha} = 0.266$, $\kappa = 1$, $\sigma_m = 0.015$, $\Delta \sim \text{Exp}(\psi)$. The $\alpha = 0.95$ picture is qualitatively identical.




Figure 5: Expected shortfall (solid) and value at risk (dashed) at $\alpha = 0.99$ as a function of the exponential ambiguity rate ψ (lower ψ is higher expected ambiguity $\mathbb{E}[\Delta] = 1/\psi$), for the ample-capital baseline (blue) and a capital-scarce market with re-entry (red, $\mathcal{W}_e = 0.09$). Vertical dotted lines mark $\psi^\dagger = -\ln(1 - \alpha)/y_\alpha^*$ for each capital regime. Because the scarce market censors at the deeper $y_\alpha^{\text{cap}} > y(x_e)$, its threshold sits to the left. For $\psi > \psi^\dagger$ value at risk tracks the Gaussian floor while expected shortfall already reflects the panic. For $\psi < \psi^\dagger$ both rise, and scarce capital lifts them above the baseline. Parameters as in Figure 4.




Figure 6: Interim tail risk at $\alpha = 0.99$ as a function of aggregate expected-utility capital $\mathcal{W}_e = x_e W_e$, at fixed ambiguity ($\psi = 35$). For $\mathcal{W}_e \geq \mathcal{W}_e^*$ the market is ample and re-entry never occurs, so tail risk is flat in capital. For $\mathcal{W}_e < \mathcal{W}_e^*$ (shaded) capital is scarce: the fire sale deepens as capital falls, raising both expected shortfall (red) and value at risk (blue). The dashed line is the expected shortfall the full-exit approximation would report, pinning the premium at y_α^{cap} throughout the panic. The solid red line lies below it because re-entering ambiguity-averse capital caps the crash (Regime 4), so accounting for re-entry *lowers* the required interim capital relative to the conservative approximation. Parameters as in Figure 4.

2.3.7 How much capital does the normal approximation miss?

A risk manager who acknowledges the ambiguity-driven variance (31) but still assumes the interim loss is normal would set capital against a normal with the *correct total variance* $\text{Var}(L) = \sigma_m^2 + x_a^2 \text{Var}(Z)$, obtaining, for either measure,

$$K_\alpha^{\text{N}} = \kappa \rho_\alpha^{\text{N}}(L) = \kappa c_\alpha \sqrt{\sigma_m^2 + x_a^2 \text{Var}(Z)}, \quad c_\alpha = \begin{cases} z_\alpha, & \text{VaR,} \\ \phi(z_\alpha)/(1 - \alpha), & \text{ES.} \end{cases} \quad (44)$$

The true requirement is $K_\alpha^{\text{T}} = \kappa \rho_\alpha(L)$ computed from the actual, exit-bearing loss distribution (30). The question, for each measure, is how much capital the normal approximation

misses,

$$\mathcal{G}_\alpha^\rho \equiv K_\alpha^T - K_\alpha^N = \kappa [\rho_\alpha(L) - \rho_\alpha^N(L)], \quad (45)$$

positive when the normal model *understates* the capital the true tail requires. Both the benchmark and the gap are proportional to κ , so the comparison is a statement about the risk measures themselves.

In the ample-capital baseline, under $\Delta \sim \text{Exp}(\psi)$, the ingredients are closed-form. The censored shock $Z = \min(\Delta, y_\alpha^*)$ with $y_\alpha^* = y(x_e)$ has

$$\mathbb{E}[Z] = \frac{1 - q^*}{\psi}, \quad \text{Var}(Z) = \frac{2(1 - q^*(1 + \psi y_\alpha^*))}{\psi^2} - \frac{(1 - q^*)^2}{\psi^2}, \quad q^* = e^{-\psi y_\alpha^*}, \quad (46)$$

which fixes the benchmark (44). For the true measure the well-separated limit is used ($\sigma_m \rightarrow 0$, fundamental noise negligible next to the panic), in which the loss is $L = x_a(Z - \mathbb{E}[Z])$ and both true measures are explicit:

$$\rho_\alpha(L) = \begin{cases} \frac{x_a}{\psi} [-\ln(1 - \alpha) - (1 - q^*)], & \text{VaR, } \alpha < 1 - q^*, \\ \frac{x_a}{\psi} [-\ln(1 - \alpha) - \frac{\alpha q^*}{1 - \alpha}], & \text{ES, } \alpha < 1 - q^*, \\ x_a(y_\alpha^* - \mathbb{E}[Z]), & \text{either, } \alpha \geq 1 - q^*. \end{cases} \quad (47)$$

The three lines say something economically specific. Below the threshold $1 - q^*$ the loss quantile lies in the body of the ambiguity distribution, and value at risk and expected shortfall differ because expected shortfall already averages in the panic mass while value at risk does not. *At or above* the threshold the entire averaged tail is the panic atom, a single mass point at the panic drop, so the two measures coincide at $x_a(y_\alpha^* - \mathbb{E}[Z]) = y_\alpha^* - R_0 - \Omega$, the equilibrium price decline net of the panic premium already in p_0 . Substituting (44) and (47) into (45) gives the two capital gaps in closed form. In the panic-dominated regime they take the transparent form

$$\mathcal{G}_\alpha^{\text{VaR}} = \kappa x_a [(y_\alpha^* - \mathbb{E}[Z]) - z_\alpha \sqrt{\text{Var}(Z)}], \quad \mathcal{G}_\alpha^{\text{ES}} = \kappa x_a [(y_\alpha^* - \mathbb{E}[Z]) - \frac{\phi(z_\alpha)}{1 - \alpha} \sqrt{\text{Var}(Z)}]. \quad (48)$$

Two forces set the sign of the gap. The true tail is the fixed panic drop $x_a(y_\alpha^* - \mathbb{E}[Z])$, while the normal benchmark scales with the standard deviation $x_a \sqrt{\text{Var}(Z)}$. When the panic is a small-to-moderate probability, the drop is far out relative to its dispersion, so that the distribution is highly skewed and the true requirement exceeds the variance-matched normal. It does so by more for value at risk than for expected shortfall, since $z_\alpha < \phi(z_\alpha)/(1 - \alpha)$ makes the normal benchmark for value at risk the more slender of the two. As the panic

probability rises toward one, the drop $x_a(y_\alpha^* - \mathbb{E}[Z])$ shrinks, because it is priced into p_0 and $\mathbb{E}[Z] \rightarrow y_\alpha^*$, while $\text{Var}(Z)$ and hence the normal benchmark keep rising. The gap therefore narrows, crosses zero, and turns negative: once a panic is near-certain the normal model *overstates* the capital that the (now largely anticipated) tail requires.

Beyond the baseline, the gap depends on capital through the depth of the fire sale. Table 1 evaluates the true measures and the gap from the exact four-regime loss (22), holding ambiguity fixed at $\psi = 35$ (a small onset panic probability, $q^* \approx 0.4\%$) and sweeping aggregate expected-utility capital $\mathcal{W}_e$ from scarce to ample, at $\kappa = 1$ so entries are per-unit risk measures in percent. The understatement is large throughout the region a supervisor faces. In the ample market ($\mathcal{W}_e \geq \mathcal{W}_e^*$) the true 99% expected shortfall is 6.92% against a variance-matched 5.47%, a gap of +1.45. The gap *widens as capital falls*, reaching +2.08 at $\mathcal{W}_e = 0.082$, because scarce capital deepens the panic drop while the variance-matched normal responds only to the second moment. Re-entry works the other way at the margin: the true expected shortfall at $\mathcal{W}_e = 0.082$ is 7.63% with re-entry against 7.76% under the full-exit approximation, so ignoring Regime 4 would overstate the required buffer by about a tenth of a percent of the position. The gap is largest for expected shortfall, which loads on the panic at every reachable confidence level. Matching the variance is not enough, and the shortfall grows as the banking sector thins.

Table 1: Capital under the true, exit-bearing loss distribution versus a variance-matched normal, as a function of aggregate expected-utility capital $\mathcal{W}_e = x_e W_e$, from the exact four-regime loss (22). Entries are per-unit risk measures $\rho_\alpha(L)$ in percent (capital at $\kappa = 1$). Here ρ^N is the variance-matched normal benchmark (44), ρ^T the true measure with re-entry, and $\mathcal{G}$ their difference (45). Positive $\mathcal{G}$ means the normal model understates required capital. As capital falls into the scarce region ($\mathcal{W}_e < \mathcal{W}_e^* = 0.106$) the fire sale deepens and both ρ^T and the gap rise. The last column reports the true expected shortfall under the full-exit approximation (no Regime 4), which re-entry shaves. Parameters $R_0 = 0.08$, $x_a = x_e = 0.5$, $\sigma_{IC_\alpha} = 0.266$, $\sigma_m = 0.015$, $\Delta \sim \text{Exp}(35)$.

$\mathcal{W}_e$	y_α^{cap}	Value at risk			Expected shortfall			ES ^T
		ρ^N	ρ^T	$\mathcal{G}$	ρ^N	ρ^T	$\mathcal{G}$	no re-ent.
<i>Panel A: $\alpha = 0.99$</i>								
0.082	0.184	4.84	6.01	+1.17	5.55	7.63	+2.08	7.76
0.090	0.176	4.82	5.97	+1.16	5.52	7.36	+1.84	7.41
0.100	0.166	4.79	5.91	+1.12	5.49	7.08	+1.59	7.09
$\geq \mathcal{W}_e^*$	≤ 0.160	4.77	5.85	+1.08	5.47	6.92	+1.45	6.92
<i>Panel B: $\alpha = 0.95$</i>								
0.082	0.184	3.42	3.63	+0.21	4.29	5.13	+0.83	5.15
0.090	0.176	3.41	3.64	+0.23	4.27	5.07	+0.80	5.08
0.100	0.166	3.39	3.64	+0.25	4.25	5.01	+0.76	5.01
$\geq \mathcal{W}_e^*$	≤ 0.160	3.37	3.64	+0.26	4.23	4.97	+0.74	4.97

The gap also grows with the composition of the market. Figure 7 plots it against the ambiguity-averse share x_a at a realistic requirement $\kappa = 10\%$. The gap is negligible when few traders can panic, since the interim loss is then almost purely fundamental and nearly Gaussian. It widens as x_a rises, because a larger panic-susceptible mass makes the fire-sale tail deeper and more skewed, the feature the variance-matched normal cannot see. At $x_a = 0.5$ the true 99% expected shortfall exceeds its normal benchmark by about 0.15% of the position, and by roughly 0.4% as the panic-susceptible share approaches the whole market.




Figure 7: The capital gap $\mathcal{G}_\alpha = \kappa [\rho_\alpha^T - \rho_\alpha^N]$, the capital the variance-matched normal model misses relative to the true, exit-bearing interim loss, as a function of the ambiguity-averse share x_a , for a capital requirement $\kappa = 10\%$, in percent of position value. Curves are expected shortfall (red) and value at risk (blue) at $\alpha = 0.99$ (solid) and $\alpha = 0.95$ (dashed). The gap vanishes as $x_a \rightarrow 0$ (the interim loss becomes purely fundamental and Gaussian) and rises with x_a as the panic-susceptible mass deepens and skews the fire-sale tail. It is largest for expected shortfall, which loads on the panic at every reachable confidence level. Ample-capital baseline (no re-entry), so κ scales the gap linearly. Parameters $R_0 = 0.08$, $\sigma_m = 0.015$, $\Delta \sim \text{Exp}(35)$.

2.4 Leverage, capital, and procyclicality

The downside risk measures above are, by (10), the capital a bank holds. The capital carried per unit of the portfolio,

$$K_\alpha = \kappa \rho_\alpha(L), \quad L = p_0 - p_1, \quad (49)$$

is *literally* κ times the value at risk or expected shortfall of (37)–(43). Its cyclical behavior is the risk measure of Section 2.3.6, rescaled. Figures 4 and 5, read on a κ -scaled axis, are the required-capital schedule as a function of the panic probability and the amount of ambiguity, and the maximum leverage $\bar{\lambda}_\alpha = p_0/K_\alpha$ is their reciprocal.

Result 5 (Procyclicality is inherited from the risk measure). *Fix a bank's capital W and let it*

carry the largest position (9) permits, $\bar{\theta}_\alpha = W/(\kappa\rho_\alpha(L))$. Under a value-at-risk requirement $\bar{\theta}_\alpha$ is flat in the panic probability for $q^* < 1 - \alpha$ and drops discontinuously at $q^* = 1 - \alpha$, because $\rho_\alpha(L) = \text{VaR}_\alpha$ is flat then jumps. Under an expected-shortfall requirement $\bar{\theta}_\alpha$ declines continuously in q^* , because ES_α rises continuously. A bank near a value-at-risk constraint is therefore forced to deleverage abruptly just as the panic becomes likely, selling into the price decline (20) the panic produces. An expected-shortfall constraint instead induces gradual, anticipatory deleveraging that builds the buffer while it is still cheap.

The procyclicality is doubly grounded. Its sign and shape are inherited from the risk measure of the interim loss, as in a single price-taking bank's problem. In addition, the model makes the feedback explicit: forced sales lower p_1 , which by (36) enlarges L and by (19) deepens the capital-constrained panic, tightening (9) further. This is a loss spiral that operates through equity (Adrian & Shin 2014, Brunnermeier & Pedersen 2009, Danielsson et al. 2012). It acquires a systemic character when that ambiguity is common across banks, since a rise in ambiguity then pushes many across the value-at-risk threshold at once and synchronizes deleveraging.

Result 6 (The conservatism of the risk measure is an equilibrium object). *A more conservative measure (larger c_α) has two opposing effects. Ex ante it accumulates capital earlier and more gradually against a rising panic probability (Result 5), avoiding the abrupt threshold deleveraging that a quantile forces. Ex post, once the requirement binds, it deepens the constrained-panic price (Result 2). The choice of measure is therefore an equilibrium object, trading an ex ante timing benefit against an ex post amplification cost. That tension is one the statistical comparison of value at risk and expected shortfall cannot see, since it concerns the feedback of the measure onto the price rather than the properties of a fixed loss distribution. The balance favors the more conservative measure when the panic is a small-to-moderate-probability event, when the timing benefit accrues while the amplification is rarely realized.*

This two-sided effect, together with the fact that ambiguity rather than volatility is the state variable driving it, is what the countercyclical buffer of Section 2.5 is built to navigate.

2.5 An ambiguity-indexed countercyclical buffer

The model's results carry a direct implication for the design of countercyclical capital. The externality is standard for a fire-sale economy: each bank, when it delevers into a falling market, ignores that its sales deepen the price decline others face, so private deleveraging is excessive from the system's point of view (Brunnermeier & Pedersen 2009, Danielsson

et al. 2012). Two features of the model sharpen what a supervisor should lean against. First, the state variable that governs the fire sale is not realized volatility but *ambiguity*: the panic probability is $q^* = e^{-\psi y_\alpha^*}$ and enters expected shortfall continuously (Section 2.3.6), whereas volatility need not rise with ambiguity and the bank measure of Section 3.2 is elevated in banking crises where volatility-based stress indices are muted. A buffer indexed to volatility, the implicit target of a value-at-risk rule, therefore relaxes just when the tail risk the model identifies is building. Second, the severity of a triggered panic scales with the scarcity of the remaining banks' capital and with the capital-weighted share of panic-susceptible holders (Result 2), so the buffer should be larger when the banking sector is thinly capitalized and when a large, capital-weighted mass of holders is prone to exit.

These two features suggest a concrete supervisory rule. First, *index the buffer to a measure of ambiguity, alongside the volatility and credit aggregates on which current rules rely*. The Basel III countercyclical capital buffer is anchored to the credit-to-GDP gap (Basel Committee on Banking Supervision 2010, Borio & Lowe 2002, Drehmann & Juselius 2014), a broad and slow-moving credit stock, while a value-at-risk rule is implicitly anchored to volatility. Neither speaks to the banking sector's own assessment of its assets. Volatility need not rise with ambiguity, and an aggregate credit gap does not register banks' disagreement about the value of what they hold. The provisioning-dispersion measure of Section 3.2, read directly from the Call Reports on which capital is already assessed, is elevated in banking crises where equity-market stress indices are muted, and it forecasts the deepening of the left tail one to four quarters ahead. It supplies a forward-looking, banking-sector signal that a volatility or credit-growth trigger does not. Second, *prefer a coherent tail measure (expected shortfall) to a quantile (value at risk) for the buffer's construction*, because the former accumulates capital gradually and ahead of the event while the latter is flat and then forces synchronized, procyclical deleveraging at the threshold (Result 5). Third, *weight the buffer by the capital-weighted panic-susceptible share x_a* , proxied for banks by reliance on runnable funding (uninsured deposits, wholesale and repo funding) and by the concentration of risk-bearing capacity, so that the requirement is heaviest where an exit would transfer the most risk onto the fewest balance sheets. The one countervailing consideration the model itself supplies is that a more conservative measure deepens the constrained-panic price ex post (Result 2). The buffer's benefit, anticipatory accumulation, must be weighed against this cost, and the balance tilts toward the buffer when the panic is a small-to-moderate-probability event, which is when it is being built.⁸

⁸This paper stops short of a full welfare analysis. The planner's problem that would pin down the optimal buffer size, trading the ex ante insurance value of anticipatory capital against the ex post amplification of a binding conservative rule, is a natural next step. The point here is qualitative: the *state variable* the buffer should track is ambiguity, and the *measure* it should be built on is a coherent tail statistic.

A scope condition attends these systemic and macroprudential readings. The model is one asset in one market, with exogenous fundamentals, and it does not derive comovement across assets or across banks. The systemic language above rests on a common ambiguity shock reaching many banks at once, a commonality the model imposes rather than generates. A full treatment would embed the mechanism in a multi-asset or aggregate block, in which a common shift in ambiguity produces synchronized exit and a fire sale in general equilibrium, and in which the cross-sectional composition of runnable funding shapes which banks propagate it. That extension is left for future work. The single-market results, namely the equilibrium panic, the capital-constrained fire sale, and the continuous-versus-jump behavior of the two risk measures, do not depend on it. The empirical identification does not depend on it either, since the bank-level tests compare banks within the same quarter and so hold the aggregate shock fixed.

3 Empirical calibration

The model makes sharp predictions that can be confronted with data: periods of high ambiguity should display more left-skewed and fatter-tailed returns, a larger gap between the tail risk an investor faces and what a Gaussian model reports, and higher expected returns reflecting the panic premium. These predictions are confronted first, and most directly, in the assets banks hold, using the quarterly U.S. Call Reports. The Call Reports serve three purposes: to document the shape of the bank-asset loss distribution, to construct from banks' own loss provisioning a measure of the ambiguity the model places at its center, and to show that this measure forecasts depletion of the regulatory capital the model concerns. The exercise is illustrative rather than a structural estimation: the aim is to show that the quantities the model identifies are measurable, economically sizable, and move with ambiguity in the predicted directions.

3.1 Data and mapping to the model

The centerpiece of the data is the quarterly U.S. Call Reports, which report the income, losses, and balance sheets of every U.S. commercial bank and make it possible to examine the distribution and dynamics of bank assets. From the Call Reports each bank's return on assets is constructed, along with a new measure of ambiguity. This new measure is formed using the cross-sectional dispersion of banks' loan-loss provisioning, net of the differences explained by their nonperforming loans and realized net charge-offs. The measure captures disagreement among banks about the value of their assets and is consistent with a world in which the true

distribution of asset quality is difficult to measure. In the model this measure stands in for $1/\psi$ (equivalently $E[\Delta]$), where a high value of the empirical ambiguity measure corresponds to a low ψ and a high panic probability $q^* = e^{-\psi y_\alpha^*}$. Appendix B details its construction. As a comparison for the bank measure, a monthly equity-market ambiguity index is also used, the formal measure of Brenner & Izhakian (2018) in the spirit of Epstein & Schneider (2008), which captures uncertainty about the probability distribution of returns for the equity market as a whole rather than return volatility itself. Section 3.2 contrasts the two. The remaining model primitives are mapped to observables as follows. The baseline risk premium R_0 and the fundamental price volatility σ_m are calibrated from unconditional return moments in low-ambiguity periods, where the model implies the distribution is approximately Gaussian. The panic-susceptible share x_a is proxied by measures of investor-base composition and the concentration of risk-bearing capacity. The panic probability q^* is read off the ambiguity measure as described above.

3.2 The distribution of bank asset returns

The model concerns the distribution of the asset portfolio that banks hold. Bank equity is an imperfect proxy for it, because equity is a levered claim with an implicit government guarantee. Its return distribution is roughly symmetric, because the asymmetry of the underlying assets has been absorbed by leverage and the bailout option.⁹ This analysis turns to bank assets themselves, using the quarterly Consolidated Reports of Condition and Income (“Call Reports”) filed by all U.S. commercial banks. From the income statements each bank’s quarterly return on assets is constructed, defined as net income over average assets and de-cumulated from the year-to-date filings, along with the corresponding loss items. The merged panel covers 1.24 million bank-quarters for 18,364 banks over 1984–2026.

The distribution of bank asset returns is left-skewed and leptokurtic, as the model predicts for the underlying risky asset. Pooling across banks and quarters (Table 2, Panel A), the quarterly return on assets has a skewness of -0.38 and an excess kurtosis near 60, and the asymmetry is concentrated in the tails. Measured from the median of $+0.26\%$, the first percentile lies at -1.18% while the ninety-fifth reaches only $+0.56\%$, so the downside tail is roughly five times the upside. The loss items confirm this from the other side: loan-loss provisions and securities losses are strongly right-skewed and fat-tailed. The asymmetry that is absent from bank *equity* is thus plainly present in bank *assets*.

The asymmetry is, moreover, a crisis phenomenon, the realized counterpart of the model’s

⁹In the Fama–French financial-industry portfolios, bank and broker-dealer equity returns display the same fat tails as the broad market but essentially no left skew (high-ambiguity skewness near zero, against -0.6 for the market), consistent with the bailout option truncating the lower tail of the equity claim.

panic. Panel B of Table 2 reports the cross-sectional return-on-assets distribution by year. In calm and boom years the cross-section is right-skewed, as a few exceptional banks pull up the upper tail. In the financial crisis the distribution inverts: skewness turns sharply negative (-2.6 in 2009), the mean collapses to zero, and the first percentile of the cross-section falls to -3.3% . The left tail in bank asset returns seems to materialize at times when uncertainty is highest, consistent with the model’s ambiguity-driven exit and the left-skewed price distribution it generates. Figure 8 shows the contrast directly across three quarters: the savings-and-loan crisis (1985 Q3), a calm quarter (1995 Q3), and the financial crisis (2009 Q3). The return on assets is only mildly left-skewed in the calm quarter (skewness -0.8) but develops a pronounced left tail in both crises (-3.8 in 1985 and -3.5 in 2009), while the provision rate and the nonperforming-loan ratio, which are right-skewed throughout, are most dispersed in the crisis quarters.

This left tail can, moreover, be predicted by a measure of ambiguity built from the bank data themselves. Because ambiguity is uncertainty about the distribution of asset values, banks that face it provision for losses in divergent ways. Here, ambiguity is measured each quarter by the cross-sectional dispersion of banks’ loan-loss provision rates, after controlling for the dispersion in their nonperforming-loan ratios and net charge-offs. What remains is disagreement about future losses beyond what realized credit quality and realized losses justify. The bottom row of Figure 8 plots this residual. Its dispersion is roughly three times as wide in the two banking crises as in the calm quarter (standard deviation 0.33 in 1985 and 0.37 in 2009, against 0.12 in 1995), widening just as the left tail of bank asset returns deepens. This measure has several helpful properties. First, unlike measures of ambiguity that are built from a general equity-market index, which is elevated in market-stress episodes such as 2020 that did not become banking crises, this bank measure is elevated through the 1980s banking crisis and again in 2008–2009, and it is, if anything, negatively correlated with the equity index (-0.23). Figure 9 plots the two series over time. Second, it forecasts the shape of the bank-asset distribution: a one-standard-deviation increase in provisioning disagreement is associated with a more left-skewed cross-sectional return on assets and a first percentile lower by about two-thirds of a percentage point (Newey–West $t = -7.9$, allowing for eight quarters of autocorrelation). Third, the relationship is forward-looking. It holds when the measure is lagged one to four quarters (HAC t between -6.3 and -10.1 at leads of up to a year), so it is not the mechanical consequence of provisions reducing current income. It also holds on the pre-2020 incurred-loss subsample, before the CECL transition altered provisioning (first-percentile and skewness forecasts at $t = -6.1$ and -9.9), so it is not an artifact of the change in accounting regime. It is also reinforced by the precautionary behavior the model attributes to ambiguity aversion, since quarters of heavy reserve-building precede




Figure 8: The distribution of bank asset returns and their determinants in three quarters: the savings-and-loan crisis (1985 Q3, the first quarter for which nonperforming loans are reported), a calm quarter (1995 Q3), and the financial crisis (2009 Q3). Each panel is a cross-sectional histogram across U.S. commercial banks with average assets above \$25 million, from the Call Reports. Densities are plotted, each series is winsorized at the 0.5 and 99.5 percentiles within the quarter, and the dashed line marks the mean. The rows, top to bottom, are the return on assets (net income over average assets), then the loan-loss provision rate and the nonperforming-loan ratio (both as a percent of loans), and finally the residual from the within-quarter regression of the provision rate on the nonperforming-loan ratio and the net charge-off rate. That residual is provisioning disagreement beyond what realized fundamentals, the stock of bad loans and the realized loss flow, explain, the bank-data ambiguity measure of Section 3.2. In both banking crises the return on assets develops a pronounced left tail (skewness -3.8 in 1985 and -3.5 in 2009, against -0.8 in the calm quarter) and the dispersion of the provisioning residual is roughly three times its calm-quarter level (standard deviation 0.33 and 0.37 , against 0.12). Provisions and nonperforming loans, right-skewed throughout, fatten on the right in the crises.

a deeper left tail (HAC t between -2 and -5), as, more weakly, do quarters of liquidity hoarding. Where an equity-market ambiguity measure cannot see the asymmetry of bank assets, a measure of ambiguity constructed from banks' own disagreement can. The assets that bank capital is held against are thus left-skewed and leptokurtic, most so in crises, and their asymmetry deepens when banks disagree most about what those assets are worth.




Figure 9: A bank-data measure of ambiguity over time. The solid line is the bank-asset ambiguity measure, the cross-sectional dispersion of banks' loan-loss provisioning purged of nonperforming loans and net charge-offs (Section 3.2, Appendix B), and the dashed line is the equity-market ambiguity index. Both are standardized to zero mean and unit variance over their samples. The equity-market index begins in 1993. Shaded bands mark the 2007–2009 and 2020 NBER recessions. The bank measure is elevated through the 1980s banking crisis (savings-and-loan and farm/energy failures) and again in 2008–2009, where the equity-market index is muted, and the two are negatively correlated (-0.23).

The construction places the measure in a recognized tradition while keeping it distinct from existing bank work. Treating cross-sectional dispersion as a proxy for uncertainty follows Bloom (2009), who uses the spread of firm-level outcomes to measure uncertainty, and the differences-of-opinion literature in which dispersion stands for disagreement (Diether, Malloy & Scherbina 2002, Patton & Timmermann 2010). The equity-market index compared against is the formal ambiguity measure of Brenner & Izhakian (2018). The measure is, however, different from the closest accounting work on bank provisioning. Bushman & Williams (2012), Bushman & Williams (2015), and Beatty & Liao (2011) characterize the *timeliness* of a given bank's provisioning, how promptly it recognizes expected losses, as

a property of the bank estimated over time, whereas the object here is the *cross-sectional disagreement* among banks within a quarter. The economic link is nonetheless shared. Bushman & Williams (2015) find that delayed loss recognition raises bank tail risk specifically by fattening the left tail of the outcome distribution, and Andreou, Cooper, Louca & Philip (2017) find that conservative loss provisioning reduces crash risk. Both corroborate the provisioning–tail-risk channel exploited here. The two existing bank-level measures of ambiguity or uncertainty in this area are built from equity returns (Garcia-Feijoo & Viale 2023) or from the text of earnings calls (Soto 2021), not from provisioning, and the one paper linking uncertainty to provisions runs the other way, using an external policy-uncertainty index to explain them (Ng, Saffar & Zhang 2020). No prior work appears to construct a measure of ambiguity from the dispersion of bank provisioning.

Whether this dispersion measures *ambiguity* (uncertainty about the distribution of asset values) rather than heterogeneous discretion or risk deserves comment, since cross-sectional spreads are known to contain variation unrelated to uncertainty (Jurado, Ludvigson & Ng 2015). Three features argue for the ambiguity reading. First, purging the provision rate of the nonperforming-loan ratio and the realized net charge-off rate removes the component of provisioning that realized credit quality and realized losses explain, so what remains is disagreement about losses not yet evident in the loan book or in current write-offs rather than dispersion in realized credit experience. Second, the measure is forward-looking: it forecasts the left tail one to four quarters ahead, which a contemporaneous artifact of earnings or capital management would not. Third, it discriminates between kinds of stress as an ambiguity measure should, peaking in the 1980s and 2008–2009 banking crises but not in the 2020 market disruption, and correlating negatively with an equity-market ambiguity index. This is a validation against known episodes in the spirit of Baker, Bloom & Davis (2016), and it is one that an index of generic provisioning discretion or of market uncertainty would not produce.

Table 2: The distribution of bank asset returns. Quarterly return on assets (net income over average assets, in percent) and the loan-loss provision rate from U.S. Call Reports, 1984–2026, for banks with assets above \$25 million (1.24 million bank-quarters). Panel A pools all bank-quarters. Panel B reports the cross-sectional distribution of the return on assets in selected years.

<i>Panel A: pooled distribution</i>					
	Skewness	Excess kurtosis	1st pctile	Median	95th pctile
Return on assets	−0.38	59.0	−1.18	+0.26	+0.56
Provision rate	+7.5	77.6	−0.09	+0.03	+0.37
<i>Panel B: cross-sectional return on assets by year</i>					
Year	Mean	Skewness	Excess kurtosis	1st pctile	Median
2006 (boom)	+0.28	+2.9	71.3	−0.83	+0.27
2009 (crisis)	−0.01	−2.6	18.1	−3.31	+0.14
2020 (COVID)	+0.27	+5.4	76.4	−0.48	+0.25

3.3 Provisioning ambiguity and future bank capital

The model makes a prediction about future bank capital. This section asks whether the ambiguity a bank’s provisioning reveals forecasts its subsequent capital. To do this, the object underlying the aggregate measure is taken to the bank level. For each bank-quarter, the residual ε_{it} from the within-quarter regression of the provision rate on the nonperforming-loan ratio and the net charge-off rate (Appendix B) is formed, and its magnitude $|\varepsilon_{it}|$, provisioning that a bank’s own realized credit quality does not explain, serves as a bank-level reading of ambiguity, standardized over the panel.¹⁰ The forecasting targets are three capital outcomes $k = 1, 2, 4$ quarters ahead: the probability that a dividend-paying bank cuts its trailing-year common dividend by more than half, the change in its Tier 1 risk-based capital ratio, and the change in its book equity-to-assets ratio. The risk-based ratio is the object the model’s requirement is literally written against. It is taken, together with risk-weighted assets, from the FDIC financials (Schedule RC-R), merged to the Call Report panel

¹⁰The proxy is the magnitude $|\varepsilon_{it}|$ rather than the signed residual, because ambiguity is disagreement with the cross-section in either direction. The signed residual is not a clean signal of retreat. A bank that over-provisions relative to its realized credit quality tends, if anything, to *expand* its loan book over the following year, since provisioning ahead of realized losses is largely a prudent-growth signal, reinforced under expected-loss accounting. Retreat from the market is instead forecast by the magnitude, with over- and under-provisioning alike preceding it. That symmetric, disagreement-based pattern is what the ambiguity reading implies and what a level-of-distress reading would not produce.

by the FDIC certificate–RSSD crosswalk, and available from 1990.¹¹

The design is a bank-level panel with bank *and* quarter fixed effects. The quarter effects absorb the aggregate cycle, so identification is within-quarter across banks, comparing banks that provision ambiguously to others facing the same aggregate conditions, and the bank effects absorb permanent differences in business model and portfolio. Adding controls for size and loan composition (the C&I, commercial-real-estate, residential, and consumer shares) leaves the estimates essentially unchanged, and they survive clustering by both bank and quarter (Table 3).¹² A one-standard-deviation increase in provisioning ambiguity forecasts a 4.0-percentage-point higher probability of a dividend cut within the year, against a 15% base rate, a Tier 1 ratio lower by 0.13 points, and a book equity-to-assets ratio lower by 0.08 points. Each effect grows with the horizon, and each is identified purely off within-quarter cross-bank variation. This is the sharpest form of the paper’s claim: the ambiguity read from banks’ provisioning forecasts depletion of the regulatory capital the model concerns, not merely the cycle and not merely differences in portfolios.

One specification choice is worth stating plainly. The effect on the risk-based ratio is absorbed when contemporaneous return on assets is added: the risk-based coefficient falls to -0.02 and loses significance ($t = -1.0$).¹³ This reads as confirmation of the mechanism, not a failure of it: the model’s channel is that ambiguity induces precautionary provisioning that depresses earnings and so depletes capital, and current ROA is the mediator on that path. Conditioning on the mediator mechanically removes the effect it transmits, so the correct specification for a forecasting statement includes the composition controls that address confounding but not the earnings variable that lies on the causal path. Reported that way, the ambiguity a bank’s provisioning reveals is a leading, within-quarter signal of the capital it will have a year hence.

¹¹The time-consistent Call Report extract omits Schedule RC-R. The FDIC financials supply the Tier 1 and total risk-based ratios and risk-weighted assets, computed and de-annualized, for every insured bank. Book equity, assets, loan composition, and dividends are from the Call Reports. The dividend-cut sample is restricted to banks paying a positive trailing-year dividend, from 2000 on, when dividend reporting is complete.

¹²The estimates are also essentially unchanged on the pre-2020 subsample, before the transition to expected-loss (CECL) accounting altered loan-loss provisioning. On that incurred-loss sample the four-quarter dividend-cut, Tier 1, and book-equity coefficients are $+4.1$, -0.13 , and -0.06 , each significant under two-way clustering and matching the full-sample estimates. Expected-loss provisioning, if anything, sharpens the reading of the residual as disagreement about losses not yet realized.

¹³The book equity-to-assets effect is more robust to the ROA control, at -0.04 ($t = -1.9$), because book equity moves with more than the current quarter’s earnings.

Table 3: Provisioning ambiguity and future bank capital. Each cell is the coefficient on the standardized bank-level ambiguity proxy $|\varepsilon_{it}|$ (provisioning unexplained by the bank’s own nonperforming loans and net charge-offs) in a panel regression of the capital outcome k quarters ahead, with bank and quarter fixed effects. Bank-clustered t -statistics in parentheses, and the last column reports the two-way (bank and quarter) clustered t for the specification with controls. Controls are log assets and the C&I, commercial-real-estate, residential, and consumer loan shares. Panel A is a linear probability model ($\times 100$, so entries are percentage-point changes in the cut probability) on dividend payers, 2000–2026. Panels B and C are in percentage points of the ratio, RC-R sample from 1990. U.S. Call Reports and FDIC financials, banks with assets above \$25 million.

Horizon	Time FE	Two-way FE	Two-way FE + controls	+ controls (2-way cl. t)
<i>Panel A. $Pr(\text{dividend cut} > 50\% \text{ within } k), \times 100$ (base rate 7–15%)</i>				
$k = 1$	+2.38 (13.3)	+1.84 (6.6)	+1.85 (6.7)	[+6.5]
$k = 2$	+4.36 (11.1)	+3.48 (7.1)	+3.51 (7.2)	[+7.2]
$k = 4$	+5.34 (8.3)	+3.99 (7.2)	+4.02 (7.3)	[+7.4]
<i>Panel B. Tier 1 risk-based capital ratio, change (pp)</i>				
$k = 1$	−0.058 (−8.4)	−0.056 (−7.9)	−0.051 (−6.8)	[−4.0]
$k = 2$	−0.104 (−8.9)	−0.102 (−9.2)	−0.093 (−7.9)	[−4.1]
$k = 4$	−0.142 (−8.0)	−0.140 (−8.7)	−0.126 (−7.1)	[−3.6]
<i>Panel C. Book equity/assets, change (pp)</i>				
$k = 1$	−0.023 (−8.6)	−0.012 (−4.6)	−0.022 (−5.5)	[−4.2]
$k = 2$	−0.047 (−10.2)	−0.028 (−7.4)	−0.046 (−8.2)	[−4.6]
$k = 4$	−0.079 (−6.9)	−0.050 (−5.4)	−0.079 (−5.6)	[−4.1]

The same signal reaches the terminal outcome. Matching the FDIC failure records to the panel and coding each bank-quarter by whether the bank fails within the next k quarters, provisioning ambiguity forecasts a bank’s own failure one to two years ahead.¹⁴ With quarter fixed effects and the standard supervisory controls (capital, nonperforming loans, earnings, and size), a one-standard-deviation increase in $|\varepsilon_{it}|$ raises the probability of failure within a year by 0.50 percentage points ($t = 9.3$) and within two years by 0.69 points ($t = 10.5$), a large effect against base rates of 0.4 and 0.9 percent. The estimate is identified within-quarter across banks and survives clustering by both bank and quarter. Provisioning ambiguity is thus the leading edge of a single chain: it forecasts, in turn, a bank’s dividend cuts, its capital depletion, and its failure. The measure is an independent predictor rather than a superior one. Almost every failure is already visible in nonperforming loans and thin capital well ahead of time, so provisioning ambiguity adds little to the failure ranking a standard

¹⁴Failure dates are the FDIC’s, matched to the panel by the FDIC certificate–RSSD crosswalk. The sample contains 1,524 banks that subsequently failed. The specification is a linear probability model with the same fixed effects and controls as Table 3, adding the standard supervisory ratios as regressors.

supervisory model produces (the area under the ROC curve rises from 0.962 to 0.964). Its value is to confirm, from banks' own forward-looking loss recognition, the tail risk the model places at its center.

4 Conclusion

This paper has developed an equilibrium model in which a market panic is the exit of ambiguity-averse banks from a market whose signal they can no longer confidently interpret, and has used it to derive the downside risk measures, value at risk and expected shortfall, and the regulatory capital built on them as equilibrium objects. The capital requirement is slack in normal times but binds in the panic, where it makes the depth of the fire sale a function of banks' capital and of the risk measure itself. If the potential crash is believed to be deep enough, ambiguity-averse banks do not exit since the premium associated with staying is so high. This places a floor under the severity of the potential crash. Expected shortfall loads on the possibility of a panic continuously and ahead of the event, whereas value-at-risk-based capital is flat and then jumps, forcing deleveraging into the very price decline the panic produces. A variance-matched normal model, in turn, understates the buffer by about a quarter in the model's calibration.

The paper's empirical contribution is to make this ambiguity observable in the assets banks actually hold. In the U.S. Call Reports the return on bank assets is left-skewed and leptokurtic, an asymmetry that leverage and the government guarantee hide in bank equity, and its left tail deepens most in banking crises. A new measure of ambiguity, built from the cross-sectional dispersion of banks' own loss provisioning and purged of what realized credit quality explains, rises in the banking crises of the 1980s and of 2008–2009 where equity-market ambiguity measures are muted, is negatively correlated with them, and forecasts the deepening of that left tail one to four quarters ahead. Because it is read from the reports on which supervisory capital is assessed, the measure supplies a forward-looking, bank-level signal of the tail risk that a variance-based capital rule is least able to see. In a bank-level panel with bank and quarter fixed effects it forecasts, within-quarter across banks, both dividend cuts and declines in the Tier 1 risk-based ratio one to four quarters ahead. Mapping it through the model's closed-form expressions to produce bank-level, time-varying value at risk and expected shortfall, and thereby a structural estimate of the capital a possible panic requires, is the natural next step that these first-pass results motivate.

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The bank data are the public U.S. Call Reports (FFIEC) and FDIC BankFind financials. The equity-market ambiguity index used for comparison follows Brenner & Izhakian (2018), and the Fama–French financial-industry portfolios are from the Fama–French data library. Replication code for the ambiguity measures and figures is available from the author on request.

A Proofs

Throughout, the per-unit regulatory loss on a position bought at premium $\pi = m(s) - p$ is $\rho_\alpha(p - v) = \sigma_I c_\alpha - \pi$ by (14), so the expected-utility capital constraint (9) at a per-bank position θ_e reads $\kappa \theta_e (\sigma_I c_\alpha - \pi) \leq W_e$.

Proof of Result 1. Two candidate equilibria must be compared. In the interior (no-panic) candidate both types hold positive positions, and market clearing (13) with the demands (5) and (8) gives the price (15), at which the premium is $\pi^\circ = R_0 + x_a \Delta$ and the ambiguity-averse position is $\theta_a^\circ = R_0^{-1}(\pi^\circ - \Delta) = 1 - x_e \Delta / R_0$. This is nonnegative iff $\Delta \leq R_0 / x_e = y(x_e) \leq y_\alpha^*$. In the panic candidate the ambiguity-averse banks hold zero and the expected-utility banks absorb the unit supply, so the premium is the unified y_α^* of (20). By the demand rule (8), holding zero is optimal for an ambiguity-averse bank if and only if the available premium does not exceed its reservation premium, $y_\alpha^* \leq \Delta$. The two conditions are mutually exclusive and exhaustive, so exit occurs iff $y_\alpha^* \leq \Delta$.

For the comparative statics, exit occurs on the set $\{\Delta \geq y_\alpha^*\}$, which enlarges as $y_\alpha^* = \max\{R_0/x_e, \sigma_I c_\alpha - x_e W_e/\kappa\}$ falls. The first argument is increasing in R_0 and decreasing in x_e , and the second is decreasing in $x_e W_e$. Hence a lower R_0 , a higher x_e , or a higher $x_e W_e$ each weakly lowers y_α^* and makes exit more likely, and a higher Δ makes exit more likely directly. $\square$

Proof of Result 2. By (20), $y_\alpha^* = \max\{y(x_e), \sigma_I c_\alpha - x_e W_e/\kappa\} \geq y(x_e)$, with strict inequality iff the capital-scarcity term is the larger, i.e. iff $x_e W_e < \kappa[\sigma_I c_\alpha - y(x_e)]$. On that scarce region $y_\alpha^* = \sigma_I c_\alpha - x_e W_e/\kappa$, so $\partial y_\alpha^*/\partial W_e = -x_e/\kappa < 0$ and $\partial y_\alpha^*/\partial c_\alpha = \sigma_I > 0$. Because

$c_\alpha^{\text{ES}} > c_\alpha^{\text{VaR}}$ at a common confidence level, the constrained premium, and hence the depth of the fire sale, is larger under expected shortfall. $\square$

Proof of Result 3. Assume $x_e W_e < \kappa(\sigma_I c_\alpha - R_0/x_e)$, so $y_\alpha^{\text{cap}} = \sigma_I c_\alpha - x_e W_e/\kappa > y(x_e)$, and fix $\Delta \in [y(x_e), y_\alpha^{\text{cap}}]$. First, neither pure-regime equilibrium survives. At the full-exit price $m(s) - y_\alpha^{\text{cap}}$ the premium is $y_\alpha^{\text{cap}} > \Delta$, so by (8) the ambiguity-averse banks strictly prefer a long position, contradicting full exit. At the interior price the ambiguity-averse position $1 - x_e \Delta/R_0$ is nonpositive because $\Delta \geq y(x_e) = R_0/x_e$, and moreover under scarce capital the expected-utility constraint is violated at that price, so the interior candidate fails as well.

The remaining candidate has the expected-utility banks at their binding cap, $\theta_e = W_e/[\kappa(\sigma_I c_\alpha - \pi)]$, and the ambiguity-averse banks at the interior optimum $\theta_a = R_0^{-1}(\pi - \Delta)$. Market clearing is $g(\pi) = 0$ with

$$g(\pi) \equiv x_e \frac{W_e}{\kappa(\sigma_I c_\alpha - \pi)} + x_a R_0^{-1}(\pi - \Delta) - 1.$$

On $\pi \in (\Delta, \sigma_I c_\alpha)$ both terms are continuous and strictly increasing in π (the first because its positive denominator $\sigma_I c_\alpha - \pi$ shrinks, the second linearly), so g is strictly increasing there and has at most one root. At $\pi = \Delta$, $g(\Delta) = x_e W_e/[\kappa(\sigma_I c_\alpha - \Delta)] - 1 < 0$, since $\Delta < y_\alpha^{\text{cap}}$ gives $\sigma_I c_\alpha - \Delta > x_e W_e/\kappa$. At $\pi = y_\alpha^{\text{cap}}$, the first term equals 1 (as $\sigma_I c_\alpha - y_\alpha^{\text{cap}} = x_e W_e/\kappa$) and the second is strictly positive, so $g(y_\alpha^{\text{cap}}) > 0$. By the intermediate value theorem there is a unique root $\pi^* \in (\Delta, y_\alpha^{\text{cap}})$. At this root the expected-utility frictionless demand $R_0^{-1}\pi^*$ exceeds the cap (equivalently the constraint binds), confirming the conjectured configuration. Since $\pi^* > \Delta \geq y(x_e)$ the fire sale is deeper than the frictionless panic, and since $\pi^* < y_\alpha^{\text{cap}}$ the re-entering ambiguity-averse capital keeps the price above the full-exit mark. $\square$

Proof of Result 4. In the panic-dominated regime $\alpha > 1 - q^*$ the loss quantile lies inside the panic branch, where (38) gives $\text{VaR}_\alpha = x_a(y_\alpha^* - \mathbb{E}[Z]) + \sigma_m \Phi^{-1}(\frac{\alpha - (1 - q^*)}{q^*})$. Subtracting the variance-matched normal $\text{VaR}_\alpha^N = z_\alpha \sqrt{\sigma_m^2 + x_a^2 \text{Var}(Z)}$ yields (39). Letting $\sigma_m \rightarrow 0$ (the well-separated limit) kills the last term and the radical's σ_m^2 , giving (40), $\text{VaR}_\alpha - \text{VaR}_\alpha^N = x_a[(y_\alpha^* - \mathbb{E}[Z]) - z_\alpha \sqrt{\text{Var}(Z)}]$. This is positive iff the panic drop $y_\alpha^* - \mathbb{E}[Z]$ exceeds $z_\alpha \sqrt{\text{Var}(Z)}$, which holds when the panic is a small-to-moderate-probability, highly skewed event. As $q^* \rightarrow 1$ the drop $y_\alpha^* - \mathbb{E}[Z] \rightarrow 0$ (the panic is priced into p_0 , so $\mathbb{E}[Z] \rightarrow y_\alpha^*$) while $\text{Var}(Z)$ stays bounded away from zero, so the gap crosses zero and turns negative. $\square$

Proof of Result 5. The maximal position is $\bar{\theta}_\alpha = W/[\kappa \rho_\alpha(L)]$, a strictly decreasing function of $\rho_\alpha(L)$, so its behavior in q^* is the inverse image of the risk measure's. Under a value-at-risk requirement, $\rho_\alpha(L) = \text{VaR}_\alpha$ is, by the analysis around (38), flat in q^* for $q^* < 1 - \alpha$ (the quantile sits in the fundamental region) and jumps up by the panic drop as q^* crosses

$1 - \alpha$, so $\bar{\theta}_\alpha$ is flat and then drops discontinuously. Under an expected-shortfall requirement, $\rho_\alpha(L) = \text{ES}_\alpha$ is continuous and increasing in q^* over the range in which a panic is plausible, by (43) and the continuity of its two branches at $\alpha = 1 - q^*$, so $\bar{\theta}_\alpha$ declines continuously. $\square$

B Construction of the bank-data ambiguity measure

This appendix documents the bank-data ambiguity measure of Section 3.2. The data are the quarterly Consolidated Reports of Condition and Income filed by all U.S. commercial banks. Income-statement items are reported year-to-date and reset each January, so quarterly flows are recovered by differencing within each bank-year. The sample retains banks with average assets above \$25 million. The measure is computed from 1984 onward, when banks first report nonperforming loans in the Call Reports, and only quarters with at least 500 reporting banks are used.

Each bank-quarter yields four quantities, all in percent. The *return on assets* is quarterly net income divided by average assets. The *provision rate* is quarterly loan-loss provisions divided by total loans, and the *nonperforming-loan ratio* is nonperforming loans divided by total loans. The *net charge-off rate* is quarterly net loan-and-lease charge-offs (gross charge-offs less recoveries) divided by total loans. Because the time-consistent income extract omits Schedule RI-B, this item is merged from the FDIC BankFind financials (net loan-and-lease charge-offs, already de-YTD'd), consolidated to the FED RSSD entity and keyed to bank and quarter. Ratios are winsorized at the 0.5 and 99.5 percentiles, and the return, provision, and net charge-off rates additionally at the 0.1 and 99.9 percentiles, to remove reporting errors.

The ambiguity measure is the provisioning dispersion that remains after accounting for realized credit quality. Each quarter t admits the cross-sectional regression

$$\text{prov}_{it} = a_t + b_t \text{npl}_{it} + c_t \text{nco}_{it} + \varepsilon_{it}$$

across all banks i (requiring at least 200 banks with all three items reported), and ambiguity is the standard deviation of the residuals ε_{it} , the disagreement in provisioning beyond what the stock of nonperforming loans and the realized net charge-off flow explain. Precautionary measures are also constructed as the cross-sectional medians of the quarterly change in the reserve ratio (loan-loss reserves over gross loans), the change in the cash ratio (cash over assets), and loan growth. A composite hoarding index sums the standardized cash and reserve changes and subtracts standardized loan growth.

To relate ambiguity to the shape of the bank-asset distribution, the cross-sectional skew-

ness and first percentile of the return on assets in quarter t are regressed on the standardized ambiguity measure at quarter $t - k$, for $k = 0, \dots, 4$. Reported t -statistics use Newey–West standard errors with a Bartlett kernel and a bandwidth of eight quarters, which accounts for the strong persistence of both series across the cycle.

References

- Acerbi, C. & Tasche, D. (2002), ‘Expected shortfall: A natural coherent alternative to value at risk’, *Economic Notes* **31**(2), 379–388.
- Adrian, T. & Boyarchenko, N. (2012), Intermediary leverage cycles and financial stability, Staff Report 567, Federal Reserve Bank of New York.
- Adrian, T. & Shin, H. S. (2014), ‘Procyclical leverage and value-at-risk’, *Review of Financial Studies* **27**(2), 373–403.
- Anderson, A. G. (2019), ‘Ambiguity in securitization markets’, *Journal of Banking & Finance* **102**, 231–255.
- Anderson, E. W., Ghysels, E. & Juergens, J. L. (2009), ‘The impact of risk and uncertainty on expected returns’, *Journal of Financial Economics* **94**(2), 233–263.
- Andreou, P. C., Cooper, I., Louca, C. & Philip, D. (2017), ‘Bank loan loss accounting treatments, credit cycles and crash risk’, *British Accounting Review* **49**(5), 474–492.
- Anthropelos, M. & Schneider, P. (2024), ‘Optimal investment and equilibrium pricing under ambiguity’, *Review of Finance* **28**(6), 1759–1805.
- Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. (1999), ‘Coherent measures of risk’, *Mathematical Finance* **9**(3), 203–228.
- Baker, S. R., Bloom, N. & Davis, S. J. (2016), ‘Measuring economic policy uncertainty’, *Quarterly Journal of Economics* **131**(4), 1593–1636.
- Barro, R. J. (2006), ‘Rare disasters and asset markets in the twentieth century’, *Quarterly Journal of Economics* **121**(3), 823–866.
- Basel Committee on Banking Supervision (2010), Guidance for national authorities operating the countercyclical capital buffer, Technical report, Bank for International Settlements.

- Basel Committee on Banking Supervision (2019), Minimum capital requirements for market risk, Standards, Bank for International Settlements. Fundamental Review of the Trading Book.
- Beatty, A. & Liao, S. (2011), ‘Do delays in expected loss recognition affect banks’ willingness to lend?’, *Journal of Accounting and Economics* **52**(1), 1–20.
- Bloom, N. (2009), ‘The impact of uncertainty shocks’, *Econometrica* **77**(3), 623–685.
- 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.
- Borio, C. & Lowe, P. (2002), Asset prices, financial and monetary stability: Exploring the nexus, BIS Working Paper 114, Bank for International Settlements.
- Boyarchenko, N. (2012), ‘Ambiguity shifts and the 2007–2008 financial crisis’, *Journal of Monetary Economics* **59**(5), 493–507.
- Brenner, M. & Izhakian, Y. (2018), ‘Asset pricing and ambiguity: Empirical evidence’, *Journal of Financial Economics* **130**(3), 503–531.
- Brunnermeier, M. K. & Pedersen, L. H. (2009), ‘Market liquidity and funding liquidity’, *Review of Financial Studies* **22**(6), 2201–2238.
- Bushman, R. M. & Williams, C. D. (2012), ‘Accounting discretion, loan loss provisioning, and discipline of banks’ risk-taking’, *Journal of Accounting and Economics* **54**(1), 1–18.
- Bushman, R. M. & Williams, C. D. (2015), ‘Delayed expected loss recognition and the risk profile of banks’, *Journal of Accounting Research* **53**(3), 511–553.
- Caballero, R. J. & Krishnamurthy, A. (2008), ‘Collective risk management in a flight to quality episode’, *Journal of Finance* **63**(5), 2195–2230.
- Caballero, R. J. & Simsek, A. (2013), ‘Fire sales in a model of complexity’, *The Journal of Finance* **68**(6), 2549–2587.
- Cao, H. H., Wang, T. & Zhang, H. H. (2005), ‘Model uncertainty, limited market participation, and asset prices’, *Review of Financial Studies* **18**(4), 1219–1251.

- Carr, P. & Wu, L. (2009), ‘Variance risk premiums’, *Review of Financial Studies* **22**(3), 1311–1341.
- Condie, S. & Ganguli, J. V. (2017), ‘The pricing effects of ambiguous private information’, *Economic Theory* **64**(3), 535–567.
- Danielsson, J., Shin, H. S. & Zigrand, J.-P. (2012), Procyclical leverage and endogenous risk, Working paper, London School of Economics.
- Dicks, D. L. & Fulghieri, P. (2019), ‘Uncertainty aversion and systemic risk’, *Journal of Political Economy* **127**(3), 1118–1155.
- Diether, K. B., Malloy, C. J. & Scherbina, A. (2002), ‘Differences of opinion and the cross section of stock returns’, *Journal of Finance* **57**(5), 2113–2141.
- Dow, J. & Werlang, S. R. d. C. (1992), ‘Uncertainty aversion, risk aversion, and the optimal choice of portfolio’, *Econometrica* **60**(1), 197–204.
- Drechsler, I. (2013), ‘Uncertainty, time-varying fear, and asset prices’, *Journal of Finance* **68**(5), 1843–1889.
- Drehmann, M. & Juselius, M. (2014), ‘Evaluating early warning indicators of banking crises: Satisfying policy requirements’, *International Journal of Forecasting* **30**(3), 759–780.
- Easley, D. & O’Hara, M. (2009), ‘Ambiguity and nonparticipation: The role of regulation’, *Review of Financial Studies* **22**(5), 1817–1843.
- El Ghaoui, L., Oks, M. & Oustry, F. (2003), ‘Worst-case value-at-risk and robust portfolio optimization: A conic programming approach’, *Operations Research* **51**(4), 543–556.
- Epstein, L. G. & Schneider, M. (2003), ‘Recursive multiple-priors’, *Journal of Economic Theory* **113**(1), 1–31.
- Epstein, L. G. & Schneider, M. (2008), ‘Ambiguity, information quality, and asset pricing’, *Journal of Finance* **63**(1), 197–228.
- Gabaix, X. (2012), ‘Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance’, *Quarterly Journal of Economics* **127**(2), 645–700.
- Garcia-Feijoo, L. & Viale, A. M. (2023), ‘Ambiguity and risk factors in bank stocks’, *Journal of Financial Research* **46**(4), 993–1019.

- Gilboa, I. & Schmeidler, D. (1989), ‘Maxmin expected utility with non-unique prior’, *Journal of Mathematical Economics* **18**(2), 141–153.
- He, Z. & Krishnamurthy, A. (2013), ‘Intermediary asset pricing’, *American Economic Review* **103**(2), 732–770.
- Illeditsch, P. K., Ganguli, J. V. & Condie, S. (2021), ‘Information inertia’, *The Journal of Finance* **76**(1), 443–479.
- Jurado, K., Ludvigson, S. C. & Ng, S. (2015), ‘Measuring uncertainty’, *American Economic Review* **105**(3), 1177–1216.
- Kang, Z., Li, X. & Li, Z. (2020), ‘Mean-CVaR portfolio selection model with ambiguity in distribution and attitude’, *Journal of Industrial and Management Optimization* **16**(6), 3065–3081.
- Lucca, D. O. & Moench, E. (2015), ‘The pre-FOMC announcement drift’, *Journal of Finance* **70**(1), 329–371.
- Maccheroni, F., Marinacci, M. & Rustichini, A. (2006), ‘Ambiguity aversion, robustness, and the variational representation of preferences’, *Econometrica* **74**(6), 1447–1498.
- Miao, J., Wei, B. & Zhou, H. (2019), ‘Ambiguity aversion and the variance premium’, *Quarterly Journal of Finance* **9**(2), 1–36.
- Ng, J., Saffar, W. & Zhang, J. J. (2020), ‘Policy uncertainty and loan loss provisions in the banking industry’, *Review of Accounting Studies* **25**(2), 726–777.
- Patton, A. J. & Timmermann, A. (2010), ‘Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion’, *Journal of Monetary Economics* **57**(7), 803–820.
- Routledge, B. R. & Zin, S. E. (2009), ‘Model uncertainty and liquidity’, *Review of Economic Dynamics* **12**(4), 543–566.
- Soto, P. E. (2021), ‘Breaking the word bank: Measurement and effects of bank level uncertainty’, *Journal of Financial Services Research* **59**(1), 1–45.