Taming Black Swans: Risk Premia, Wealth Effects, and Monetary Policy under Tail Risk *

C. Christopher Hyland†

March 25, 2026

Abstract

This paper studies how downside risk affects monetary policy through risk premia and wealth effects. I develop a New Keynesian macro-finance framework in which increases in tail risk raise required returns on risky assets, depress asset valuations, reduce private wealth, and lower aggregate demand. With nominal rigidities, this mechanism lowers the neutral real rate relevant for policy. Using option-implied disaster probabilities, I show empirically that tail-risk shocks behave like negative aggregate demand shocks: output, consumption, investment, employment, inflation, and safe yields all fall following an increase in disaster risk. Quantitatively, tail risk is a first-order driver of the neutral rate: a one percentage-point increase in disaster probability lowers $r_t^\star$ by about 16 basis points. Policy counterfactuals imply that a central bank that tracks this risk-adjusted neutral rate stabilizes output and inflation much more effectively than one that ignores it. Finally, I construct a new historical estimate of the risk-adjusted neutral rate, showing that conventional measures miss sharp crisis-time declines associated with elevated risk premia and financial stress.

Keywords: Macro-finance, Monetary Economics, Higher Order Moments.

*I am grateful to Andrey Alexandrov, Francesco Bianchi, Javier Bianchi, Lukas Boehnert, Markus Brunnermeier, Sitong Ding, Martin Ellison, Sergio de Ferra, Alessandro Ferrari, Andrea Ferrero, Marco Garofalo, Goutham Gopalakrishna, Lars Hansen, Kathryn Holston, Miyu Haraguchi, Xinran He, Rustam Jamilov, Tobias Kawalec, Rohan Kekre, Paymon Khorrami, Paul Klemperer, Alex Kohlhas, Jiri Knesl, Seung Joo Lee, J. Zachary Mazlish, Michael McMahon, Fernando Mendo, Enrique Mendoza, Sebastian Merkel, Silvia Miranda-Agrippino, Theofanis Papamichalis, Ricardo Reis, Giovanni Rosso, Francisco Ruge-Murcia, Kevin Schneider, Andrew Scott, Dejanir Silva, Vikor Todorov, Dimitrios Tsomocos, Tim Willems, Jonathan Wright, Nancy Xu for comments. All errors are mine.

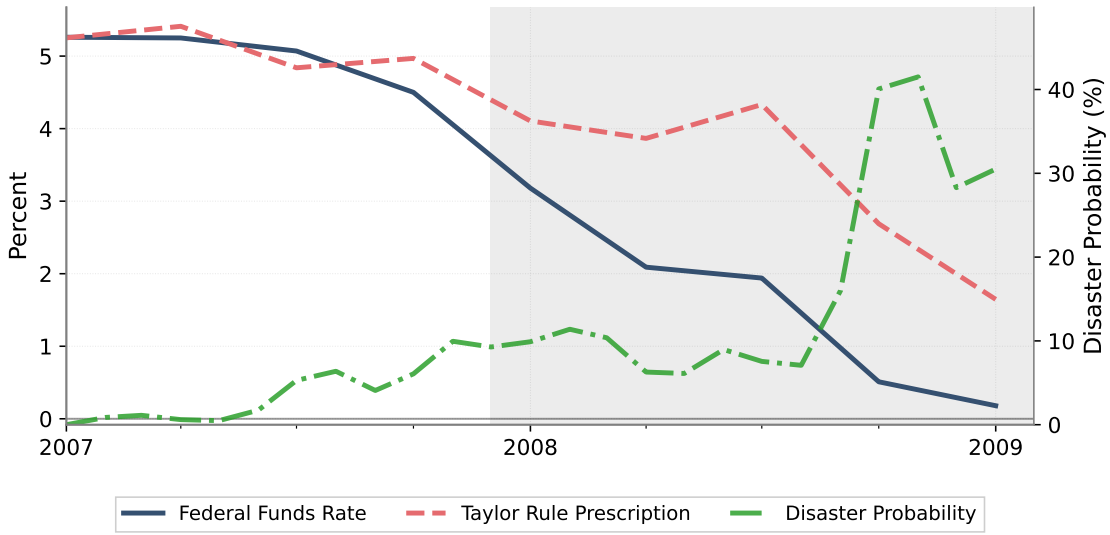
†University of Oxford, Saïd Business School, Exeter College, and Oxford-Man Institute. Email: christopher.hyland@sbs.ox.ac.uk.

1 Introduction

In the third quarter of 2007, the Federal Reserve began cutting interest rates at a pace that standard monetary policy benchmarks could not justify. Figure 1 shows the puzzle. Starting in late 2007, the federal funds rate fell sharply from above 5% to near zero by the end of 2008, while the original Taylor (1993) rule prescribed holding rates near 4–5% for much of the same period. The divergence is large, sustained, and economically consequential.

A natural first explanation is that the standard inputs into the Taylor rule—inflation and the output gap—must already have deteriorated enough to warrant aggressive easing. Figure 2 shows why that explanation is incomplete. Through the middle of 2007, inflation remained elevated, driven in part by commodity-price movements, and the output gap was closing only gradually in real time.[1] Taken at face value, those signals called for patience, not a rapid sequence of cuts. The same conclusion holds across the usual rule variants: balanced-approach, inertial, first-difference, and effective-lower-bound-adjusted specifications all prescribe policy rates materially above the realised federal funds rate for much of 2007 and into 2008; see Appendix A.1.

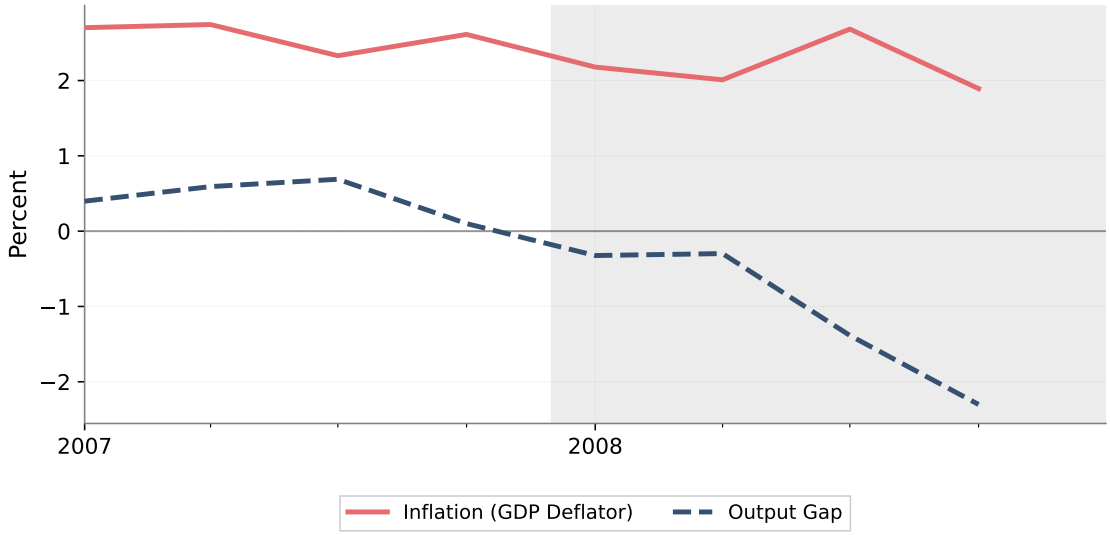
Figure 1: Federal Funds Rate vs. Taylor Rule During the Financial Crisis



Notes: The figure plots the effective federal funds rate, the prescription from the original Taylor (1993) rule using real-time data, and the disaster probability implied by out-of-the-money S&P 500 put options following Barro and Liao (2021). Shaded region indicates NBER-dated recession (December 2007–June 2009). Taylor rule: $i^\star = r^\star + \pi + 0.5(\pi - \pi^\star) + 0.5y$, where $r^\star = 2\%$, $\pi^\star = 2\%$, and the output gap uses real-time estimates.

[1] Real-time estimates of the output gap draw on Greenbook projections where available and CBO estimates otherwise.

Figure 2: Taylor Rule Components: Inflation and Output Gap



Notes: The figure plots the GDP deflator inflation rate and the output gap (100 times the log deviation of real GDP from potential GDP) using real-time data. Shaded region indicates NBER-dated recession (December 2007–June 2009).

What changed in late 2007 was not only the outlook for realised inflation and activity. Financial markets also began pricing a sharp increase in downside macroeconomic risk. The third series in Figure 1 gives one measure of that repricing: the probability of a rare macroeconomic disaster inferred from out-of-the-money S&P 500 put options following Barro and Liao (2021). This measure rose sharply just as the federal funds rate began its sustained departure from the Taylor prescription and peaked during the acute phase of the crisis. Other market-based indicators tell a similar story. As Appendix A shows, the VIX and the Excess Bond Premium of Gilchrist and Zakrajšek (2012) spiked well ahead of the worst deterioration in the conventional Taylor-rule inputs. FOMC statements and speeches from this period are explicit: policymakers repeatedly cited heightened downside risks, financial-market disruptions, and unusual uncertainty as reasons for accommodation.[2][3][4][5][6]

This observation points to a broader question. What should monetary policy do when financial markets reprice downside risk in a way that moves asset prices, private wealth, and aggregate demand? Standard New Keynesian models offer only a partial answer. In the benchmark framework, risk enters through the mean and variance of shocks, while the

[2] https://www.federalreserve.gov/newsevents/pressreleases/monetary20070807a.htm
[3] https://www.federalreserve.gov/newsevents/pressreleases/monetary20070817b.htm
[4] https://www.federalreserve.gov/newsevents/pressreleases/monetary20070918a.htm
[5] https://www.federalreserve.gov/newsevents/speech/kroszner2007116a.htm
[6] https://www.federalreserve.gov/newsevents/speech/bernanke20071015a.htm

policy problem is typically summarised by inflation and the output gap. But crisis episodes are precisely periods in which the left tail thickens, risk premia rise, and asset valuations move sharply. If those valuation changes affect spending through wealth effects, then they also move the equilibrium real rate relevant for policy. A central bank that ignores this mechanism can misread the stance of policy precisely when the economy is most exposed.

The main contribution of this paper is to develop and document a *risk-premium channel of monetary policy through wealth effects*. The core mechanism is simple. When downside risk rises, investors demand greater compensation for holding risky claims. Risk premia increase and asset prices fall. Lower asset prices reduce household wealth and weaken aggregate demand. With nominal rigidities, this decline in demand lowers the equilibrium real interest rate, so the appropriate policy intercept falls as well. In this environment, monetary policy should respond not only to inflation and slack, but also to movements in a *risk-adjusted neutral rate*.

The mechanism. The argument can be summarised in five steps. First, higher-order features of the aggregate shock distribution—in particular negative skewness and excess kurtosis—raise required returns on risky assets. Second, those movements in required returns depress equity prices and broader asset valuations. Third, lower valuations reduce private wealth and contract spending through wealth effects. Fourth, weaker demand lowers the equilibrium real rate of interest, $r_t^\star$. Fifth, monetary policy should accommodate that decline in $r_t^\star$ if it is to stabilise output and inflation. The key economic object is therefore not tail risk in isolation, but the way tail risk works through *risk premia and wealth effects* to move the neutral rate.

This logic implies a simple generalisation of the standard neutral-rate concept. In the model developed below, the risk-adjusted neutral rate can be written as

$$r_t^\star = \underbrace{g_t + \rho}_{\text{Ramsey benchmark}} - \underbrace{\mathcal{W}_t}_{\text{Risk wedge}}, \tag{1}$$

where g_t is expected productivity growth, ρ is the rate of time preference, and $\mathcal{W}_t$ is a wedge that depends on the full shape of the shock distribution, including variance, skewness, kurtosis, and higher-order moments. The same forces that raise compensation for bearing downside risk in asset markets also lower the equilibrium real rate relevant for monetary policy. In that sense, the paper links two objects that are usually studied separately: the equity premium and the neutral rate.[7]

[7] The paper focuses on the demand-side channel running from changes in downside risk through asset prices to

What this paper does. I develop this argument in three steps. First, I build a tractable New Keynesian macro-finance model in which a representative household with Epstein–Zin preferences values the full distribution of future consumption. The model nests the standard Gaussian benchmark as a special case, but it allows the same risk wedge that drives asset premia to enter the IS relation and the neutral rate. This delivers a transparent theoretical result: shocks to downside risk shift $r_t^\star$ through their effects on equilibrium asset prices and private wealth.

Second, I take the mechanism to the data. Using the disaster-probability series of Barro and Liao (2021), constructed from deep out-of-the-money S&P 500 put options, I show that increases in downside risk behave empirically like negative aggregate-demand shocks. Identified shocks reduce output, consumption, investment, and employment simultaneously on impact, while inflation and safe yields fall alongside real activity. This joint response is exactly what the wealth-effect channel predicts: a rise in required returns depresses valuations, contracts demand, and lowers the equilibrium real rate.

Third, I use the theory to measure the policy-relevant object. I construct a new risk-adjusted estimate of the neutral rate that augments the standard Laubach and Williams (2003) framework with time variation in downside risk, disciplined by option-implied disaster probabilities and the VIX. This series provides a market-informed measure of movements in $r_t^\star$ that standard trend-based estimates smooth away. I then embed the mechanism in a quantitative continuous-time New Keynesian model with disaster jumps and use it to study policy rules and counterfactuals around the onset of the Global Financial Crisis.

Main results. The paper establishes four results.

First, risk premia and wealth effects are quantitatively important drivers of the neutral rate. A one-percentage-point increase in disaster probability lowers $r_t^\star$ by roughly 16 basis points, an effect much larger than that of an equivalent increase in diffusion volatility. Crisis-time movements in the neutral rate are therefore not explained primarily by slow-moving trend forces. They also reflect sharp changes in the price of downside risk and the associated revaluation of wealth.

Second, the model implies that monetary policy should track a risk-adjusted neutral rate. In the analytical model, the same wedge that widens the equity premium lowers $r_t^\star$. In the quantitative model, a policy rule that responds only to conventional macro variables sets rates too high when downside risk rises. A rule that tracks the risk-adjusted neutral rate

spending and the neutral rate. Other manifestations of tail risk, especially cost-push or supply-side channels, are important but left for future work.

stabilises output and inflation much more effectively, even though it does not eliminate the rise in financial stress itself.

Third, the historical counterfactual around 2007–08 is economically large. A calibrated model matched to the onset of the crisis implies that the Federal Reserve accommodated a substantial but incomplete part of the tail-risk-driven fall in $r_t^\star$. Had policy ignored that decline entirely, the contraction in output would have been much more severe. The interpretation is not that the Fed directly targeted option prices or higher moments; it is that policy partially responded to a fall in the neutral rate induced by rising downside risk.

Fourth, conventional estimates of $r_t^\star$ miss crisis dynamics. The risk-adjusted neutral rate falls sharply in the Global Financial Crisis and again during COVID-19, by amounts that are largely absent from standard benchmarks such as Laubach and Williams (2003). This changes the interpretation of the policy stance. A real rate that appears accommodative relative to a smooth trend-based estimate can be close to neutral, or even restrictive, once financial repricing and wealth effects are taken into account.

Related literature. This paper sits at the intersection of three literatures.

The first is the asset-pricing literature on rare disasters and higher-order risk. A large body of work shows that investors demand compensation not only for variance but also for skewness, kurtosis, and disaster exposure, and that these risks help account for the equity premium and the low risk-free rate (Barro, 2006a; Gabaix, 2012; Gourio, 2012; Wachter, 2013). That literature has been highly successful in explaining asset prices, but it has largely taken monetary policy as given. My contribution is to bring those same risk-pricing forces into a monetary environment and show that they move the neutral rate through aggregate-demand effects.

The second is the macro-finance literature on wealth effects and demand-driven recessions. Declines in asset prices reduce household net worth and compress spending, with potentially large consequences for output and employment (Brunnermeier et al., 2024; Caballero and Simsek, 2020; Del Negro et al., 2017a). This paper builds on that insight but adds a monetary-policy margin: once wealth effects are combined with nominal rigidities, changes in risk premia become changes in the policy-relevant equilibrium real rate.

The third is the literature on the natural rate of interest. Existing work has emphasised secular forces such as trend growth, demographics, and safety premia in explaining the decline in $r^\star$ (Brunnermeier et al., 2024; Del Negro et al., 2017b; Laubach and Williams, 2003). I show that time-varying downside risk generates large, rapid, and crisis-timed

movements in $r_t^\star$ that are missed by those approaches. In this respect, the paper also relates to the uncertainty literature (Bloom, 2009; Bloom et al., 2018; Kozeniauskas et al., 2018; Kozlowski et al., 2019), but with a sharper mechanism: the relevant object is not uncertainty in general, but the component of risk that commands a premium in financial markets, revalues wealth, and feeds back into demand.

The remainder of the paper proceeds as follows. Section 2 presents a two-period analytical model and derives the risk-premium wedge linking asset prices and the neutral rate. Section 3 documents the macroeconomic effects of downside-risk shocks. Section 4 develops the continuous-time New Keynesian model with jump-diffusion technology. Section 5 calibrates the model, evaluates the quantitative importance of downside risk for the neutral rate, and studies policy counterfactuals around the Global Financial Crisis. Section 6 estimates the risk-adjusted neutral rate and presents its historical decomposition. Section 7 concludes.

2 Stylised Model

The model builds on Caballero and Simsek (2020), who develop a New Keynesian framework with wealth effects in which fluctuations in perceived risk generate demand recessions through their effect on asset prices and household portfolio allocation. I modify their framework to characterise the full distribution of aggregate risk rather than its second moment alone. This extension allows me to derive how each higher-order cumulant of the shock distribution enters equilibrium asset prices, and how the resulting movements in risk premia feed through wealth effects into the neutral rate relevant for monetary policy. The central result of the section is that the same risk wedge that widens the equity premium also lowers the equilibrium real rate, providing a simple analytical link between asset pricing, wealth effects, and monetary policy.

2.1 Environment

Consider a two-period economy, $t \in \{0, 1\}$, with a single consumption good that serves as the numeraire. Production uses a single factor, capital, with supply fixed and normalised to unity.

Production and New Keynesian firms. Firms face fully rigid nominal prices at $t = 0$; prices become flexible at $t = 1$. Let z_t denote capital productivity, which equals potential output. At $t = 0$, firms meet consumption demand c_0 up to potential output, so actual output satisfies:

$$y_0 = c_0 \leq z_0. \tag{2}$$

I normalise $z_0 = 1$, so that period-0 potential output serves as the natural unit of account. At $t = 1$, prices are flexible and output reaches potential, $y_1 = z_1$. Since $z_0 = 1$, the period-1 productivity level z_1 can be interpreted simultaneously as potential output at $t = 1$ and as the gross growth rate of potential output between the two periods. I therefore define $G \equiv \log z_1$, which is the log productivity level at $t = 1$ and equivalently the log growth rate of potential output, or TFP growth. The variable G is stochastic, and its distribution is the central object of the model.

Monetary policy. With fully sticky prices at $t = 0$, the real interest rate equals the nominal policy rate. The central bank sets r^f to target the neutral rate, defined as the rate consistent with output at potential, $y_0 = z_0 = 1$.[8] The question is therefore how the distribution of G changes the value of r^f that clears goods and asset markets at full employment.

Distribution of G: the cumulant-generating function. Rather than imposing a parametric distribution on G, I characterise it through its cumulant-generating function (CGF). For a random variable G, the CGF is defined as $c_z(\theta) \equiv \log \mathbb{E}[e^{\theta G}]$ and encodes the entire distribution of G in a single analytic object (Billingsley, 1995). Its key property is that successive derivatives evaluated at $\theta = 0$ deliver the cumulants of G: the first cumulant is the mean, the second is the variance, the third is related to skewness, the fourth to excess kurtosis, and so on. Expanding around zero gives:

$$c_z(\theta) \equiv \log \mathbb{E}\left[e^{\theta G}\right] = \kappa_1 \theta + \frac{\kappa_2}{2}\theta^2 + \frac{\kappa_3}{6}\theta^3 + \frac{\kappa_4}{24}\theta^4 + \cdots, \tag{3}$$

where κ_i denotes the i-th cumulant of G. The CGF is particularly convenient for asset pricing because expected exponential payoffs, which appear naturally in the pricing of any asset with log-linear returns, can be expressed directly in terms of the CGF evaluated at an argument equal to the investor's risk aversion. Martin (2013) exploits precisely this structure to characterise consumption-based asset prices under an arbitrary return distribution, and the approach has since become standard in the macro-finance literature

[8] I abstract from the zero lower bound throughout the stylised model.

for pricing assets whose returns exhibit skewness and heavy tails. I follow the same approach here. Crucially, the representation in equation (3) imposes no parametric restriction on the shock distribution: skewness, kurtosis, and tail thickness all enter equilibrium outcomes directly, and the standard Gaussian case, in which $\kappa_n = 0$ for all $n \geq 3$, emerges as a special case.

2.2 Asset Market Structure

Households allocate savings between two assets.

Asset menu. The first asset is a risk-free bond in zero net supply, which pays one unit of the consumption good per unit invested and carries log return r^f. The bond being in zero net supply means that in equilibrium all bond positions are held domestically and net out: the bond serves to transfer resources across states but does not represent net claims on the productive sector.

The second asset is the market portfolio, a claim on period-1 output y_1. Let $\mathcal{Q}$ denote its equilibrium price at $t = 0$. The log return on the market portfolio is:

$$r^m = \log\left(\frac{y_1}{\mathcal{Q}}\right). \tag{4}$$

At $t = 1$, prices are flexible, so firms produce at potential and $y_1 = z_1$. Making this substitution and recalling the definition $G \equiv \log z_1$, the market return becomes:

$$r^m = G - \log \mathcal{Q}, \tag{5}$$

where $\log \mathcal{Q}$ is determined at $t = 0$ and is therefore non-stochastic from the perspective of the household. The market return is consequently an affine function of the TFP growth shock G, shifted by a constant known at the time the portfolio decision is made.

The CGF of the market return. Since $r^m = G - \log \mathcal{Q}$ is an affine shift of G, the CGF of the market return inherits the CGF of TFP growth directly.

Lemma 2.1 (CGF of the market portfolio)**.** *Let $c_{r^m}(\theta) \equiv \log \mathbb{E}[e^{\theta r^m}]$ denote the CGF of the market return. Then:*

$$c_{r^m}(\theta) = c_z(\theta) - \theta \log \mathcal{Q}. \tag{6}$$

Proof. See Online Appendix B. □

This result has an important implication. The log price $\log \mathcal{Q}$ enters equation (6) only through the linear term $\theta \log \mathcal{Q}$, which affects the first derivative of the CGF at $\theta = 0$ and hence only the first cumulant. All higher-order cumulants of the market return are therefore identical to those of TFP growth:

$$\kappa_n(r^m) = \kappa_n(G), \qquad n \geq 2. \tag{7}$$

This means that the variance, skewness, kurtosis, and all higher cumulants of asset returns in this economy are determined entirely by the assumed distribution of TFP growth. Equilibrium asset pricing pins down the expected return on the market portfolio, but it cannot alter the shape of the return distribution beyond the mean. Whatever distributional assumptions I place on G therefore translate directly and without distortion into the higher-order cumulants of market returns. The CGF representation in equation (6) thus provides a tight link between the underlying macroeconomic shock distribution and the asset return distribution that households face, with no parametric restriction imposed on either.

2.3 The Household Problem

The economy is populated by a representative household endowed with current output y_0 and the market portfolio. At $t = 0$, the household allocates wealth between consumption c_0 and savings a_0. Savings are invested in a portfolio composed of a share ω in the risky market asset and a share $1 - \omega$ in the risk-free asset.

Preferences take the form of Epstein-Zin utility with unitary elasticity of intertemporal substitution and risk aversion $\gamma > 1$.[9] The household solves:

$$\begin{aligned}
\max_{c_0,\, a_0,\, \omega} \quad & \log c_0 + e^{-\rho} \log U_1, \\
\text{where} \quad & U_1 = \left(\mathbb{E}\left[c_1^{1-\gamma}\right]\right)^{\frac{1}{1-\gamma}}, \qquad \gamma > 1, \\
\text{subject to} \quad & c_0 + a_0 = \underbrace{y_0 + \mathcal{Q}}_{\equiv W_0}, \\
& c_1 = a_0 \underbrace{\left[\omega e^{r^m} + (1-\omega) e^{r^f}\right]}_{\equiv R^p(\omega)}.
\end{aligned}$$

The period-0 budget constraint has two components. The term y_0 is the household's

[9] Unitary EIS delivers the log-linear consumption rule used below and simplifies the aggregation of higher-order risk. See Online Appendix B.1 for the full derivation of the preference specification.

endowment of current output. The term $\mathcal{Q}$ is the date-0 market value of the aggregate risky claim: it is the price at which the household sells its endowment of the market portfolio. Together they constitute spendable wealth, $W_0 \equiv y_0 + \mathcal{Q}$. The period-1 constraint reflects the gross return on the savings portfolio $R^p(\omega)$, which is stochastic through the market return r^m.

The homotheticity of preferences allows this problem to decompose into two stages: an intratemporal portfolio choice and an intertemporal savings decision.

Portfolio choice. Conditional on savings a_0, the household selects the portfolio weight ω to maximise the certainty-equivalent utility of terminal wealth. The optimal ω satisfies the Euler equation:

$$\mathbb{E}\left[\mathcal{M}(\omega)\, e^{r^m}\right] = 1, \tag{8}$$

where $\mathcal{M}(\omega)$ is the stochastic discount factor.[10] Equation (8) equates the risk-adjusted returns on the two assets. The SDF $\mathcal{M}$ is high in bad states and low in good states, so optimality requires that expected payoffs, weighted by marginal utility, are equalised across assets.

Intertemporal savings decision. Recall $W_0 \equiv y_0 + \mathcal{Q}$ denotes total household wealth at $t = 0$. Given the optimal portfolio weight, solving the savings problem with log utility delivers a consumption-wealth ratio that is constant and independent of investment opportunities:

$$c_0 = \frac{1}{1 + e^{-\rho}} W_0, \tag{9}$$

$$a_0 = \frac{e^{-\rho}}{1 + e^{-\rho}} W_0. \tag{10}$$

Asset market clearing requires that the entire market portfolio is held, so $\omega = 1$. Imposing this condition and combining equation (9) with goods market clearing $y_0 = c_0$ yields, following Caballero and Simsek (2020):

$$y_0 = e^{\rho}\, \mathcal{Q}. \tag{11}$$

Equation (11) is the key result of the household block and makes the wealth transmission channel explicit. A fall in the asset price $\mathcal{Q}$ reduces household wealth W_0 directly. With the consumption-wealth ratio fixed by log utility, lower wealth passes through into lower consumption c_0. Because prices are rigid at $t = 0$, goods market clearing converts this

[10] The derivation of $\mathcal{M}$ and the full first-order conditions are in Online Appendix B.3.

demand shortfall into lower output rather than lower prices, giving $y_0 = c_0$. Equation (11) therefore summarises the central link in the paper: asset prices move wealth, and wealth moves aggregate demand. What remains to be shown is how the distribution of G feeds into $\mathcal{Q}$ through equilibrium asset pricing. That is the task of the next subsection.

2.4 Asset Prices

Equity Risk Premium. To see how higher-order cumulants of the return distribution affect equilibrium asset prices, it is useful to isolate the component of the market return that is purely about risk. Subtracting the first moment of r^m constructs a mean-zero pure risk shock:

$$\tilde{\varepsilon} \equiv r^m - \mathbb{E}[r^m]. \tag{12}$$

By construction, $\kappa_1(\tilde{\varepsilon}) = 0$, while all higher-order cumulants are inherited unchanged from r^m: $\kappa_n(\tilde{\varepsilon}) = \kappa_n(r^m)$ for all $n \geq 2$. Stripping out the first moment isolates the second cumulant and above, leaving a shock that captures purely the shape of the return distribution with no contribution from expected returns. The CGF of $\tilde{\varepsilon}$ is then $c_{\tilde{\varepsilon}}(\theta) = c_m(\theta) - \theta\,\mathbb{E}[r^m]$.

We can now characterise how these higher-order cumulants determine the equity premium. The result shows that the equity premium equals a risk wedge $W(\{\kappa_n\})$, a weighted sum over the cumulants of $\tilde{\varepsilon}$ with coefficients pinned entirely by the household's risk aversion.

Theorem 1 (Equity Risk Premium)**.** *Let r^m denote the log market return with CGF $c_m(\theta) \equiv \log \mathbb{E}[e^{\theta r^m}]$, and let $\tilde{\varepsilon} \equiv r^m - \mathbb{E}[r^m]$ be the pure risk shock with CGF $c_{\tilde{\varepsilon}}(\theta) = c_m(\theta) - \theta\,\mathbb{E}[r^m]$ and cumulants $\{\kappa_n\}_{n\geq 2}$. Define the* risk-wedge coefficients*:*

$$\Upsilon_n \equiv \frac{(-\gamma)^n - (1-\gamma)^n}{n!}, \qquad n \geq 1,$$

and the risk wedge*:*

$$W(\{\kappa_n\}) \equiv \sum_{n=1}^{\infty} \Upsilon_n\, \kappa_n = c_{\tilde{\varepsilon}}(-\gamma) - c_{\tilde{\varepsilon}}(1-\gamma). \tag{13}$$

Given $\gamma > 1$, the market-clearing allocation $\omega = 1$ is supported if and only if:

$$\underbrace{\mathbb{E}[r^m] - r^f}_{\textit{Equity premium}} \;=\; \underbrace{W(\{\kappa_n\})}_{\textit{Risk wedge}} \;=\; \sum_{n=1}^{\infty} \Upsilon_n\, \kappa_n, \tag{14}$$

where $\mathbb{E}[r^m] = \kappa_1(G) - \log \mathcal{Q}$.

Proof. See Online Appendix B.4. □

Equation (14) is exact for any distribution for which the CGF exists, exhaustive in that every cumulant of $\tilde{\varepsilon}$ enters the equity premium, and interpretable in that the coefficients Υ_n measure how much risk aversion amplifies each cumulant, growing in magnitude with both n and γ so that a more risk-averse household is disproportionately sensitive to the shape of the left tail. The standard mean-variance result, where only κ_2 matters, is the Gaussian special case with $\kappa_n = 0$ for $n \geq 3$. Expanding the first few terms:

$$W(\{\kappa_n\}) = (1-\gamma)\,\kappa_1 + \frac{2\gamma-1}{2}\,\kappa_2 + \frac{3\gamma-3\gamma^2-1}{6}\,\kappa_3 + \frac{4\gamma^3-6\gamma^2+4\gamma-1}{24}\,\kappa_4 + \cdots \tag{15}$$

The $n = 1$ term vanishes by construction: since $\tilde{\varepsilon}$ is mean-zero, $\kappa_1(\tilde{\varepsilon}) = 0$, so expected returns contribute nothing to the risk wedge and the sum is determined entirely by κ_2 and above. For $\gamma > 1$, the wedge is increasing in variance and excess kurtosis but decreasing in skewness, so a distribution with a thicker left tail raises the required equity premium. To restore equilibrium after an increase in the risk wedge, the equity premium must widen. Since $\mathbb{E}[r^m] = \kappa_1(G) - \log \mathcal{Q}$, this occurs through a fall in the equilibrium asset price $\mathcal{Q}$. Through equation (11), that fall in asset prices reduces wealth and contracts output directly. The higher-order moments of G therefore matter for macroeconomic outcomes because they first move risk premia and valuations. As shown in Online Appendix B.11, this mechanism also generates countercyclical Sharpe ratios: when tail risk rises, the reward per unit of effective risk must increase to clear the market, producing the flight-to-safety dynamics observed in the data.

Neutral Rate of Interest. The central bank seeks to hold output at potential. Given the normalization of the capital stock and initial productivity ($z_0 = 1$), potential output at $t = 0$ is unity, so the monetary authority achieves $y_0 = y_0^*$ by setting the risk-free rate r^f equal to the neutral rate.

Proposition 2.1 (Risk-Adjusted Neutral Rate)**.** *The neutral real rate consistent with full employment is:*

$$\boxed{r^{f*} = \underbrace{g+\rho}_{\textit{Ramsey term}} - \underbrace{\sum_{n=1}^{\infty} \Upsilon_n \kappa_n}_{\textit{Risk Wedge}}} \tag{16}$$

where the risk-wedge coefficients Υ_n are as defined in Theorem 1, and $\{\kappa_n\}$ are the cumulants of the mean-zero risk shock $\tilde{\varepsilon}$.

Equation (16) has two parts. The Ramsey term $g+\rho$ is the familiar deterministic benchmark: the neutral rate rises with trend productivity growth g, since households willingly defer consumption when future income is high, and with the discount rate ρ, since more impatient households require a higher return to save. The risk wedge $\sum_{n\geq 1} \Upsilon_n \kappa_n$ subtracts from this benchmark, capturing the full cumulant structure of the shock distribution. As precautionary saving rises with left-tail risk, higher-order uncertainty depresses r^{f*} below the Ramsey level. Standard New Keynesian models, which either linearize around a non-stochastic steady state or restrict shocks to be Gaussian, retain only the variance term $\Upsilon_2\kappa_2$ in this wedge; they therefore miss the contributions of skewness and kurtosis that become quantitatively important during crisis episodes.

The key point is that equation (16) is not an independent condition from equation (14). The same tail-risk variation that widens the equity premium also lowers the neutral rate. Asset pricing and monetary policy are thus linked by a common wedge: when disaster risk rises, risky-asset prices must fall to clear the market, wealth declines, and the equilibrium safe rate consistent with full employment falls as well. This is the analytical version of the risk-premium channel emphasized in the introduction. At the same time, declining risk-free rates and falling risky-asset prices replicate the flight-to-safety dynamics documented by Lustig and Verdelhan (2012) and Campbell and Cochrane (1999), generating a countercyclical market price of risk endogenously rather than through an assumption about preference dynamics. Online Appendix B.11 formalizes this via a generalized Sharpe ratio.

Analytical Examples. We evaluate equation (14) and equation (16) under two distributions for the risk shock $\tilde{\varepsilon}$, illustrating how the shape of the shock distribution enters both the equity risk premium and the neutral rate through the common risk wedge $W(\{\kappa_n\})$.

Example 1: Gaussian benchmark. If $\tilde{\varepsilon} \sim \mathcal{N}(0, \sigma^2)$, all cumulants beyond the variance vanish and the risk wedge reduces to $\Upsilon_2\sigma^2 = (\gamma - \frac{1}{2})\sigma^2$. The equity premium and neutral rate are then:

$$\mathbb{E}[r^m] - r^f = \left(\gamma - \tfrac{1}{2}\right)\sigma^2, \tag{17}$$

$$r^{f*} = g + \rho - \left(\gamma - \tfrac{1}{2}\right)\sigma^2. \tag{18}$$

The coefficient $\gamma - \frac{1}{2}$ combines the precautionary saving motive, proportional to γ, with the Jensen's inequality correction of $-\frac{1}{2}$. The two results share the same wedge: whatever

depresses the neutral rate below the Ramsey benchmark is precisely the compensation investors require to hold the risky asset.

Example 2: Normal + Bernoulli disaster jump. Let $\tilde{\varepsilon} = Z + (B - p)u$ where $Z \sim \mathcal{N}(0, \sigma^2)$, $B \sim \text{Bernoulli}(p)$, $u < 0$, and $Z \perp B$. The Bernoulli component captures the possibility of a rare disaster of size u occurring with probability p. Figure 3 shows the resulting mixture density for three values of p. The distribution is bimodal: the tall central peak reflects the normal-times component, while a smaller left hump corresponds to the disaster state. As p rises from 1% to 5%, mass shifts from the center toward the left tail, increasing both negative skewness ($\kappa_3 < 0$) and excess kurtosis ($\kappa_4 > 0$). Both enter the risk wedge with positive weight, amplifying the upward pressure on the equity premium and the downward pressure on the neutral rate.

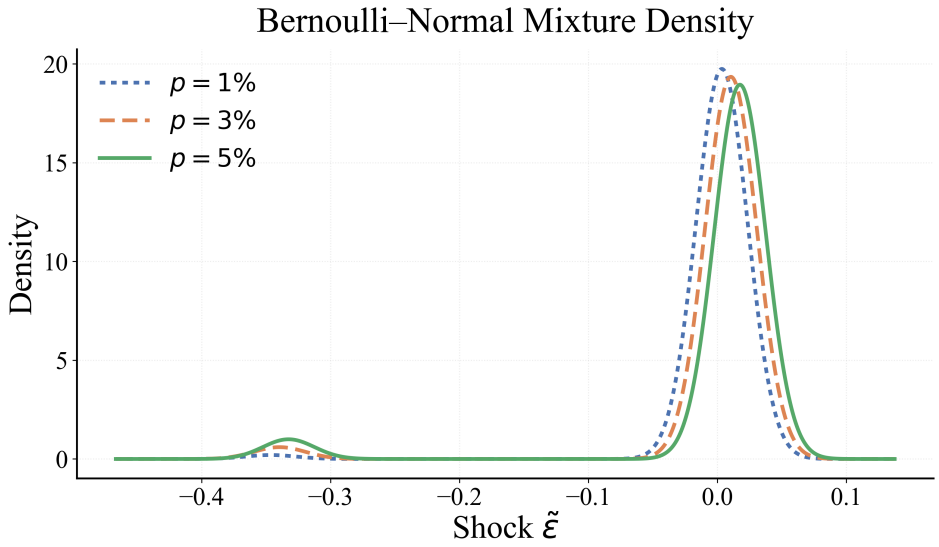


Figure 3: **Bernoulli–Normal Mixture Density.** Density of the demeaned shock $\tilde{\varepsilon} = Z + (B - p)u$ for $Z \sim \mathcal{N}(0, \sigma^2)$, $B \sim \text{Bernoulli}(p)$, and $u < 0$, under three values of the disaster probability p. The central peak corresponds to normal-times outcomes; the left hump at $\tilde{\varepsilon} \approx u$ corresponds to the disaster state. Increasing p shifts mass toward the left tail, raising negative skewness and excess kurtosis and thereby amplifying the disaster risk wedge in (20).

Evaluating the cumulant-generating function of this mixture yields the equity premium and neutral rate:

$$\mathbb{E}[r^m] - r^f = \left(\gamma - \tfrac{1}{2}\right)\sigma^2 + \underbrace{\log \frac{1 - p + pe^{-\gamma u}}{1 - p + pe^{(1-\gamma)u}} + pu}_{\text{disaster premium}}, \tag{19}$$

$$r^{f*} = g + \rho - \left[\left(\gamma - \tfrac{1}{2}\right)\sigma^2 + \log \frac{1 - p + pe^{-\gamma u}}{1 - p + pe^{(1-\gamma)u}} + pu \right]. \tag{20}$$

The disaster premium in (19), equivalently the bracketed term in (20), is strictly positive for $u < 0$ and $\gamma > 1$, and increases monotonically in both the disaster probability p and severity $|u|$. As $p \to 0$ or $u \to 0$, it vanishes and both expressions collapse to the Gaussian benchmarks (17)–(18). The mixture also generates non-zero higher cumulants, in particular negative skewness and excess kurtosis, all of which contribute positively to the wedge.[11]

These examples convey a broader point: any distributional feature that raises the marginal utility of wealth in bad states simultaneously widens the equity premium above the Gaussian level and depresses the neutral rate below the Ramsey benchmark. The two results are two sides of the same wedge. Taken together with the asset-pricing mechanism of Section 2.4, which connects equilibrium valuations to output through equation (11), these results complete the transmission channel: higher-order risk raises required returns, lowers asset prices, depresses wealth and demand, and therefore shifts the neutral rate relevant for monetary policy.

2.5 Comparative Statics of Disaster Risk

This section examines how macroeconomic quantities respond as disaster risk varies, using the Normal–Bernoulli framework derived in the analytical examples above. The exercise holds mean growth g fixed throughout and asks: as the probability of a disaster p rises, how do asset prices, output, inflation, and welfare move in equilibrium? Because mean fundamentals are unchanged, every response in what follows reflects the pricing of distributional shape alone.

The key object is the risk wedge $W(\{\kappa_n\})$, which aggregates the higher-order effects of varying p into a single sufficient statistic that simultaneously determines the equity premium, the neutral rate, and equilibrium output. I also compute an option-implied variance measure, VIX^2, defined as the entropy of the market return. Following Martin (2017), the entropy of the portfolio return can be recovered directly from the CGF of the market return derived in Proposition 2.1, providing a model-consistent counterpart to the option-implied variance observed in the data.

[11] Online Appendix B.4 provides the complete derivation and verifies that the Bernoulli–Normal CGF reproduces (19) and (20) exactly from the general formulae (14) and (16).

2.5.1 Numerical Illustration

Due to the analytical tractability of the model, I solve for a global nonlinear equilibrium. Table 1 displays the calibrated parameters; the key objects of interest are disaster size u and probability p.[12]

Parameter	Description	Value	Source
γ	Relative risk aversion coefficient	4	Standard
ρ	Time preference rate	0.005	Martin (2008)
g	Steady-state growth rate	0.025	Martin (2008)
σ	Volatility parameter	0.02	Martin (2008)
u	Disaster size	-0.2 to -0.5	Barro (2006b)

Table 1: Model parameters.

Figure 4 displays comparative statics as disaster probability p varies from 0 to 5%, under three calibrations of disaster size ranging from a conservative 20% drop ($u = -0.2$) to a severe 50% drop ($u = -0.5$).

Interpretation. The top-left panel shows the equity premium rising sharply with p. Under the intermediate calibration ($u = -0.35$), the premium climbs from near zero to roughly 3.8 percentage points as p reaches 5%; under the severe calibration ($u = -0.5$), it exceeds 9 percentage points. This steepening is not merely a variance effect: raising p worsens skewness and kurtosis simultaneously, and both enter $W(\{\kappa_n\})$ with positive weight, so the amplification is inherently nonlinear in the tail parameters. The top-middle panel shows VIX^2 rising in lockstep with the premium and steepening more sharply under larger $|u|$, reflecting that worsening tail risk registers primarily as a rise in the market price of risk per unit of effective volatility rather than as an increase in diffusion volatility alone.

The top-right panel shows the neutral real rate moving in the opposite direction and by comparable magnitudes. Starting from roughly 5.3% at $p = 0$, r^{f*} falls to about 1.8% under the moderate calibration and to approximately -3.8% under the severe one. The comparison between the top-left and top-right panels is the key insight: the same increase in downside risk that raises required compensation for holding risky assets also lowers the return households demand on safe assets. The model therefore links equity premia and the neutral rate through one underlying force rather than treating them as separate objects.

This repricing transmits into the real economy through the valuation channel. The bottom-left panel shows current output y_0 declining with p, reflecting the asset-pricing

[12] Online Appendix C.5 describes the exact analytical solutions.

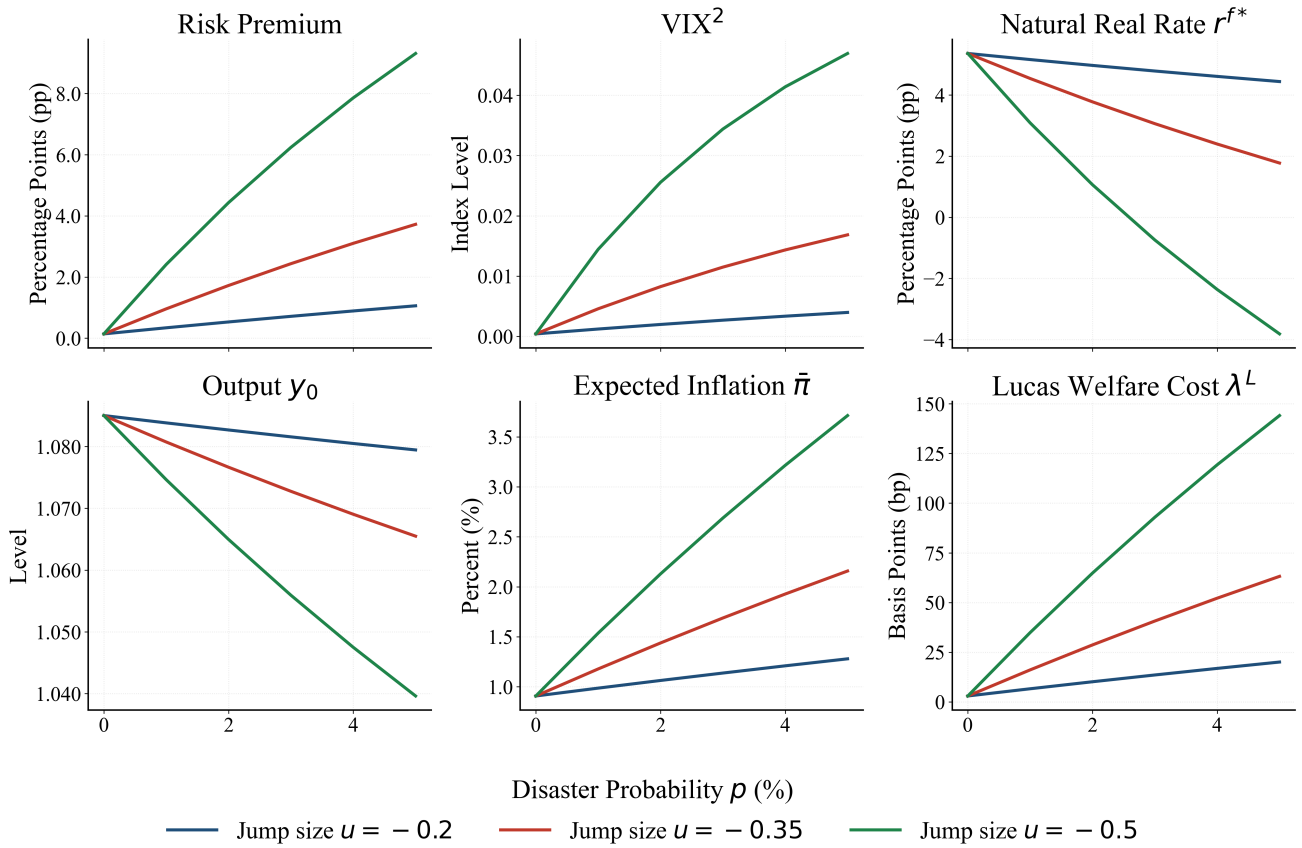


Figure 4: Comparative statics from varying disaster risk. Each panel shows the equilibrium response of key variables to changes in disaster probability p under three calibrations of disaster size u. The figure illustrates how rare-disaster risk propagates through risk premia, asset prices, wealth, and macroeconomic aggregates.

identity $y_0 = e^{\rho}\mathcal{Q}$: when tail risk rises, equilibrium asset prices fall to restore market clearing, and lower valuations contract aggregate demand one-for-one. Put differently, the movement in risk premia matters for monetary policy because it revalues wealth.

The bottom-middle panel shows expected inflation $\bar{\pi}$ rising alongside the output contraction. The mechanism is intertemporal and specific to the structure of the stylised model. Because prices are fully rigid at $t = 0$, the price level today cannot adjust to the shock; the entire burden of intertemporal goods-market clearing falls on the future price level P_1. As disaster risk rises, the precautionary saving motive intensifies: households shift demand away from period 0 and toward period 1. With period-1 output at potential and today's price level fixed, goods-market clearing at $t = 1$ requires P_1 to rise, pushing expected inflation up. In this sense, $\bar{\pi}$ is not a Phillips-curve object but an intertemporal relative-price adjustment: it measures the gap between depressed current consumption and unchanged expected future output, and the model clears that gap entirely through a higher future

price level. This mechanism stands in deliberate contrast to the inflation dynamics in the empirical evidence of Section 3 and the quantitative model of Section 4. There, disaster risk operates as a negative demand shock under nominal rigidities, and inflation is governed by the New Keynesian Phillips curve: falling output generates slack, marginal costs decline, and inflation falls unless policy fully accommodates the drop in the risk-adjusted neutral rate. The two results are complementary rather than contradictory: the stylised model isolates the intertemporal price-level adjustment required when P_0 is frozen, while the quantitative model and the data reveal that under more general nominal rigidities, the dominant effect of disaster risk is deflationary through the output gap.

Finally, the bottom-right panel reports the Lucas-style welfare cost λ, defined as the fraction of consumption the household would sacrifice to replace the risky plan with a deterministic one growing at the mean rate $\bar{R}$; the formal derivation is in Appendix **??**. This cost rises monotonically in both p and $|u|$, reaching roughly 60 basis points under the moderate calibration and about 145 basis points under the severe one at $p = 5\%$. These magnitudes are an order of magnitude larger than the Gaussian estimate of Lucas (1983) because CRRA marginal utility places disproportionate weight on the low-consumption disaster state, making welfare highly sensitive to left-tail thickness in a way that a second-order approximation cannot capture.

2.6 Truncated-Moment Monetary Policy

In practice, central banks observe only a finite set of moments when setting policy. This subsection examines the consequences of such truncation.[13] The key insight is that the optimal rate formula depends on the infinite sum of cumulants through the risk wedge $W(\{\kappa_n\}) = \sum_{n\geq 1} \Upsilon_n \kappa_n$. A central bank that targets only the first N moments uses the truncated wedge:

$$W^{(N)} \equiv \sum_{n=1}^{N} \Upsilon_n \kappa_n,$$

and sets the policy rate accordingly:

$$r^{f,(N)} \equiv g + \rho - W^{(N)}.$$

[13] I thank Lukas Boehnert for suggesting this exercise.

This mis-setting creates a gap relative to the optimal Ramsey rate:

$$\Delta r_N^f \equiv r^{f,(N)} - r^{f*} = \sum_{n>N} \Upsilon_n \kappa_n, \tag{21}$$

which equals the sum of all cumulant contributions the central bank is ignoring. Whenever the shock distribution has fat tails or negative skewness, this residual wedge is strictly positive, meaning the central bank sets rates too high relative to the Ramsey optimum. Interpreted through the mechanism above, the policy mistake is a failure to recognise how much rising downside risk has depressed the neutral rate through risk premia and wealth effects. The associated welfare cost of policy misspecification is derived in Online Appendix D.3.

2.6.1 Quantitative Implications

Figure 5 illustrates the implications under the disaster calibration $u = -0.5$ as disaster probability p varies from 0 to 5 per cent.[14] The top-left panel plots the policy rate implied by rules incorporating successively more moments. Under the Ramsey benchmark (dashed line), the optimal rate falls sharply as disaster risk increases, fully accommodating the precautionary savings motive and the decline in asset valuations. A risk-neutral central bank ($N = 1$) ignores risk entirely and holds the rate fixed at $g + \rho$. As the central bank incorporates variance ($N = 2$), skewness ($N = 3$), and kurtosis ($N = 4$), the policy rate progressively converges toward the Ramsey path. Including skewness provides a meaningful correction beyond variance alone, but a residual gap remains until the fourth moment is included.

The top-right panel shows the consequences for real activity. Under the Ramsey rule, output is successfully stabilised at potential. Under the risk-neutral rule, output falls significantly as p rises, because the central bank fails to lower rates in response to rising tail risk. Each additional moment narrows this output gap, with kurtosis effectively closing it.

The bottom panels report welfare outcomes. The bottom-right panel shows the fundamental Lucas welfare cost of risk λ, which is independent of the policy rule and rises with disaster probability. The bottom-left panel isolates the additional welfare loss from policy misspecification λ_N^{pol}. Strikingly, for a risk-neutral central bank, this policy-induced loss can exceed the fundamental cost of risk itself, underscoring that the failure to accommodate tail risk in setting the policy rate is not a second-order concern. There is a clear hierarchy

[14] Explicit formulas under the Normal-Bernoulli mixture are in Online Appendix D.4.

across moments: variance alone reduces the loss relative to full risk-neutrality but leaves it substantial; skewness closes much of the remaining gap; and kurtosis effectively eliminates the policy-induced welfare cost.

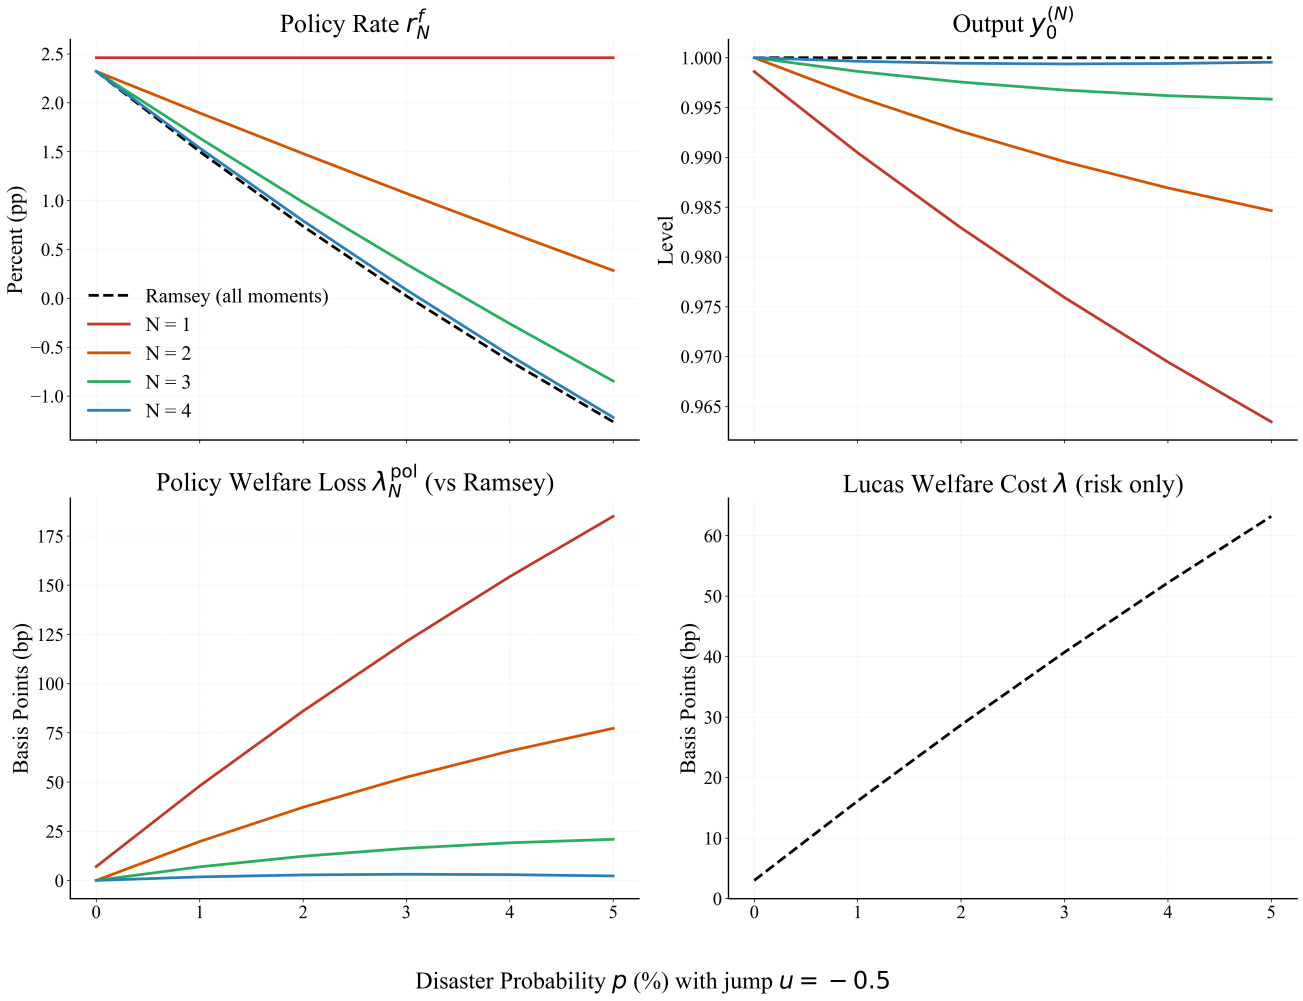


Figure 5: **Policy and Welfare Outcomes under Truncated-Moment Rules.** Top-left: policy rate as a function of disaster probability. Top-right: output relative to potential. Bottom-left: welfare loss from policy misspecification (λ_N^{pol}). Bottom-right: fundamental Lucas welfare cost of risk (λ). Disaster size is set to $u = -0.5$.

The results establish a clear hierarchy for monetary policy in a tail-risk environment. A standard mean-variance framework is insufficient when risks are asymmetric and fat-tailed: a central bank that monitors and responds to skewness and kurtosis comes much closer to tracking the true risk-adjusted neutral rate and therefore achieves welfare outcomes close to the Ramsey optimum. The broader lesson is the same as in the rest of the paper: the policy problem is not only about tail risk per se, but about how tail risk moves risk premia, asset prices, private wealth, and ultimately aggregate demand. This motivates the

richer quantitative policy analysis of Section 5, where the central bank's ability to track the risk-adjusted neutral rate is examined in a fully dynamic general equilibrium setting.

3 Empirical Evidence

The stylised model of Section 2 delivers a set of testable predictions about the macroeconomic consequences of time-varying downside risk. A rise in disaster probability raises risk premia, depresses asset valuations, lowers household wealth, and contracts aggregate demand. Through this wealth-effect channel, the neutral real rate falls. This section examines whether these predictions find support in the data. I extract a time-varying measure of disaster probability from deep out-of-the-money equity put options and use a structural vector autoregression to trace the macroeconomic effects of tail-risk shocks. The empirical analysis serves two purposes. First, it tests the directional implications of the model: do shocks to downside risk generate the joint movement in asset valuations, real activity, inflation, and safe yields implied by the theory? Second, the responses of inflation and Treasury yields help identify the relevant transmission mechanism. If disaster risk primarily operates through risk premia and wealth effects under nominal rigidities, then the data should look like a negative aggregate demand shock. Establishing that result motivates the New Keynesian framework developed in Section 4, and the co-movement between disaster risk and the policy rate documented in Online Appendix A.4 informs the specification of the monetary policy rule in the quantitative analysis.

3.1 Measuring Disaster Risk

To measure time-varying tail risk, I extract the objective probability of rare disasters, p_t, from the prices of far out-of-the-money put options on the S&P 500 index, following Barro and Liao (2021). The key insight is that for sufficiently deep out-of-the-money options with short maturities, the diffusion component of the option price is negligible and the value is driven almost entirely by the probability of a discrete jump. Under a power-law distribution for disaster sizes and CRRA utility, the option price is approximately linear in p_t, so that variation in disaster probability can be recovered from time variation in option prices after controlling for moneyness and maturity. I implement this using S&P 500 put options, filtering for deep out-of-the-money contracts with maturities up to six months; full details of the construction are provided in Online Appendix A.1.

Figure 6 plots the resulting disaster probability series alongside the VIX. The estimated series exhibits substantial time variation, spiking sharply during the 2008 Global Financial Crisis and the COVID-19 pandemic. While it co-moves with the VIX (correlation $\rho \approx 0.7$),

the two measures capture distinct objects: the VIX reflects risk-neutral entropy (Martin, 2011), whereas p_t more directly isolates changes in the perceived likelihood of extreme left-tail events under the objective measure. For the purposes of this paper, that distinction is important. The theory is not about uncertainty in general, but about the component of risk that commands a premium in financial markets and thereby moves wealth and the policy-relevant neutral rate.

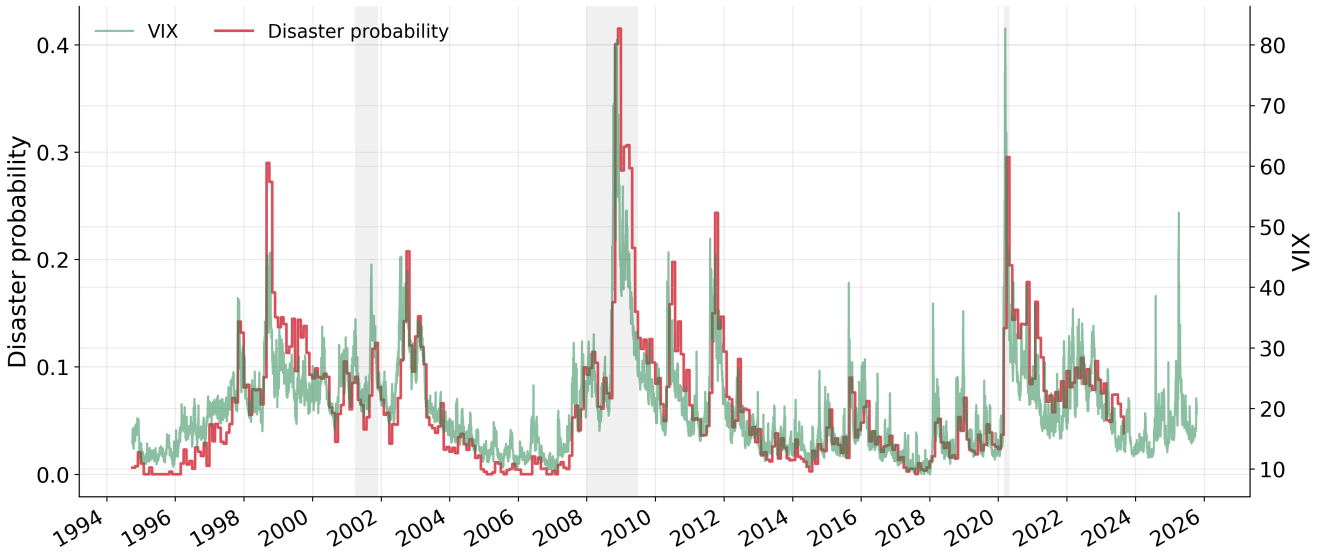


Figure 6: **Estimated Disaster Probability and the VIX.** The red line depicts the annualized disaster probability extracted from S&P 500 put options following Barro and Liao (2021). The green line depicts the VIX index. Shaded areas indicate NBER recessions.

3.2 Macroeconomic Effects of Disaster Risk Shocks

To identify the causal impact of tail risk on the business cycle, I employ a structural vector autoregression using the maximum forecast error variance identification proposed by Angeletos et al. (2020). I identify a single tail-risk shock as the structural shock that explains the maximum share of the forecast error variance of the disaster probability over a finite horizon of 0 to 40 quarters.[15] The VAR includes eight variables: the extracted disaster probability, real GDP, real consumption, real investment, employment, the real S&P 500 return, the 5-year Treasury yield, and the VIX, estimated on quarterly U.S. data.[16]

[15] This approach remains agnostic about the structural interpretation of the remaining shocks while rigorously isolating the innovation responsible for the bulk of variation in tail risk (Basu et al., 2021).

[16] Sample period: 1996Q1–2023Q4. Online Appendix A.2 provides full details of data sources and transformations.

Real activity. Figure 7 displays the impulse responses of real variables to a one standard deviation shock to disaster probability.

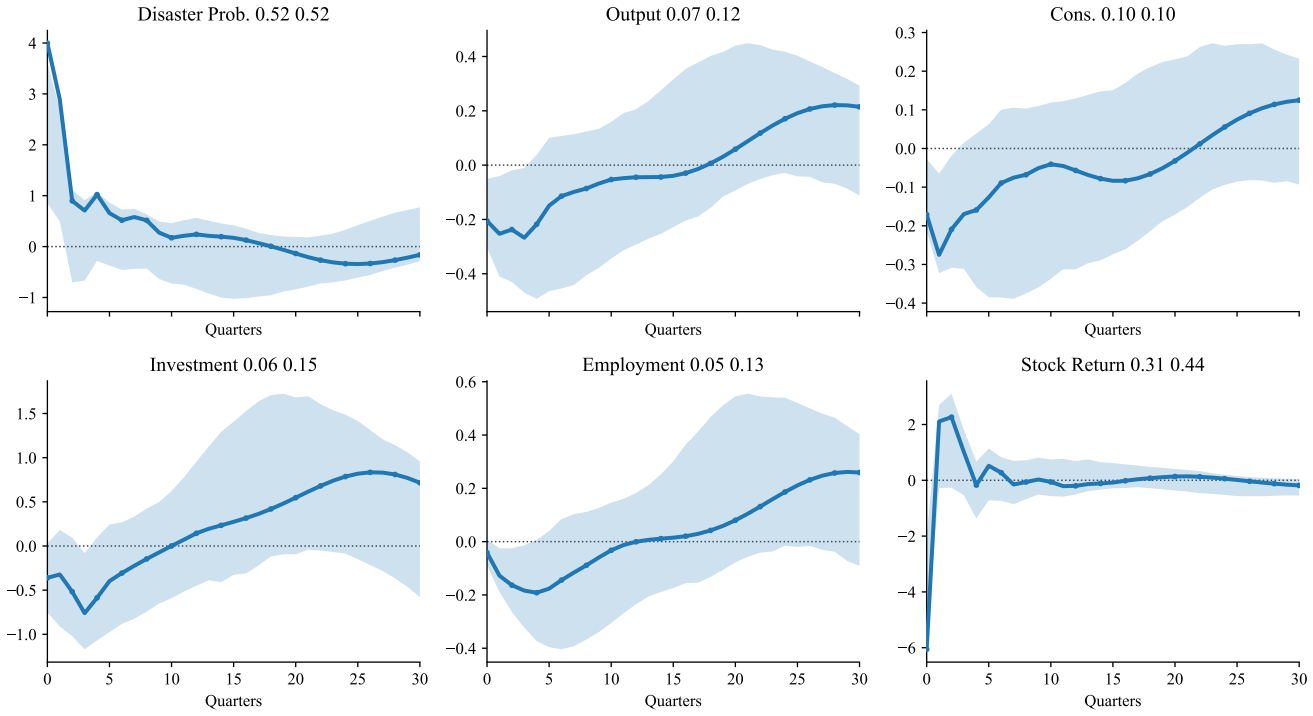


Figure 7: **Real Effects of a Disaster Probability Shock.** Impulse responses to a shock maximizing the forecast error variance of the disaster probability. Shaded areas represent 68% confidence intervals.

A shock that raises disaster probability generates a synchronized and persistent contraction in real activity: output, consumption, investment, and employment all fall significantly on impact and remain depressed for several quarters. This positive co-movement across components of spending is an important finding in its own right, since it is exactly what one would expect from a demand contraction operating through wealth effects and inconsistent with the mechanism predicted by standard flexible-price models. Barro and King (1984) show that in a time-separable flexible-price setting, a shock that depresses consumption must raise investment as households redirect precautionary saving into capital accumulation, generating a counterfactual negative co-movement. Bloom (2009) and Basu and Bundick (2017) document the same tension: flexible-price models struggle to match the synchronized contraction in real variables without additional mechanisms. The co-movement I find is instead consistent with a demand-driven recession in which financial repricing lowers wealth and nominal rigidities prevent prices from fully adjusting, forcing quantity reductions across all components of spending. Bloom et al. (2018) make a related point: models with only second-moment shocks struggle to generate this positive co-movement and typically

require first-moment disturbances to match the data. Higher-order risk shocks of the kind studied here offer a more direct route to the same result.

Nominal variables and the transmission mechanism. Figure 8 presents the responses of financial and nominal variables.

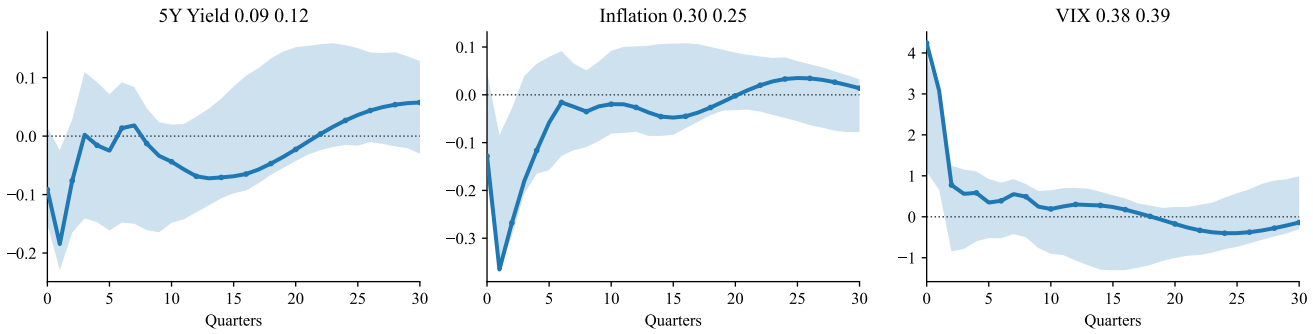


Figure 8: **Financial and Nominal Responses to a Tail Risk Shock.** Impulse responses of the 5-year Treasury yield, inflation (GDP deflator), and the VIX to a disaster probability shock. Shaded areas represent 68% confidence intervals.

The VIX rises sharply on impact, confirming that the identified shock captures a broad increase in perceived risk. More informative for the mechanism are the responses of inflation and the 5-year Treasury yield, both of which fall significantly. This pattern contrasts directly with the findings of Basu et al. (2021), who maximise the variance of the 5-year equity risk premium and find that inflation and Treasury yields rise under their identified risk shock. They interpret this as evidence against a New Keynesian channel and propose a real flight-to-safety model in which capital reallocation lowers TFP. When the shock is instead identified to maximise the variance of disaster probability specifically, both inflation and yields fall alongside output. That combination is hard to reconcile with a supply-side interpretation. Falling output, falling inflation, and declining safe yields are instead the hallmark of a negative aggregate demand shock under nominal rigidities.

This pattern also lines up closely with the mechanism in the stylised model. A rise in downside risk increases the compensation required to hold risky claims, pushes up market stress, lowers valuations, and weakens demand. The decline in the 5-year Treasury yield is especially informative from the perspective of monetary policy: it is consistent with a fall in the equilibrium real rate associated with the rise in tail risk. The empirical evidence therefore supports the view that the relevant object for policy is not simply higher uncertainty, but a risk-premium shock with macroeconomic effects operating through wealth and aggregate demand.

Together, the real and nominal responses paint a coherent picture: higher-order risk shocks, interacting with nominal rigidities, generate the joint movement in financial conditions, real activity, inflation, and safe yields predicted by the model. This provides direct empirical motivation for the New Keynesian framework augmented with disaster risk developed in Section 4.[17]

Co-movement with monetary policy. Online Appendix A.4 documents a complementary pattern in reduced-form Taylor-rule regressions. Augmenting a standard interest-rate rule with the Barro–Liao disaster probability reveals a robust negative conditional correlation between the shadow federal funds rate and measured tail risk: the implied long-run coefficient is approximately -33 basis points per percentage-point increase in annual disaster probability, conditional on inflation and the output gap. This co-movement is consistent with the model's prediction that rising disaster risk depresses the neutral rate and calls for accommodation. Read through the lens of the paper's mechanism, the evidence suggests that policymakers respond not to tail risk in the abstract, but to its implications for financial conditions, wealth, and the stance of policy relative to the risk-adjusted neutral rate. That observation informs the monetary policy rule specified in the quantitative model of Section 4.

4 Quantitative Model

This section develops a continuous-time New Keynesian framework that integrates rare disaster risk with nominal rigidities. The model extends the canonical New Keynesian framework of Woodford (2003) and Galí (2015) into continuous time, augmenting the standard Gaussian productivity process with a Poisson jump component that captures the possibility of rare, severe economic disasters. The purpose of the framework is to bring the mechanism from the stylised model into a fully dynamic setting: disaster risk raises risk premia, lowers asset prices and private wealth, and thereby shifts aggregate demand and the neutral real rate relevant for monetary policy. The model features Rotemberg price adjustment costs, monopolistic competition, and a representative household with log utility. Aggregate productivity follows a jump-diffusion process incorporating both Gaussian fluctuations and infrequent severe contractions, bridging the macro-finance literature on

[17] Online Appendix A.3 presents robustness checks using alternative identification schemes, including Cholesky ordering and identification via heteroskedasticity, which yield qualitatively similar results.

rare disasters (Barro, 2006b) with New Keynesian business-cycle theory. This delivers a risk-adjusted neutral rate whose disaster wedge is time-varying and state-dependent, and generates equilibrium dynamics for inflation and the output gap that differ sharply from the standard Gaussian New Keynesian benchmark.

4.1 Environment

The environment extends the intuition of the stylised model in Section 2 into a fully dynamic continuous-time setting with nominal rigidities. The key departure from the standard New Keynesian framework is that aggregate risk is characterised by a jump-diffusion process, so that disaster probability enters equilibrium asset prices, the neutral rate, and inflation dynamics simultaneously. In particular, the model is designed to isolate the common state variable that links asset-pricing outcomes and monetary policy: the same rise in disaster risk that widens risk premia also lowers the equilibrium real rate through its effect on wealth and aggregate demand.

Technology. Aggregate productivity A_t follows a jump-diffusion process. Writing log productivity $a_t \equiv \log A_t$:

$$da_t = \mu_A\, dt + \sigma_A\, dZ_t + J_A\, dN_t, \qquad J_A < 0, \tag{22}$$

where Z_t is a standard Brownian motion with volatility $\sigma_A > 0$, N_t is a Poisson process with constant arrival intensity $\lambda_A > 0$, and $J_A < 0$ is the log fall in TFP upon a disaster. The process has two distinct sources of uncertainty: a continuous Gaussian component capturing normal business-cycle fluctuations, and an infrequent discrete jump generating the rare, severe contractions central to the model's asset-pricing and neutral-rate implications. For the propositions in this section, λ_A and the drift μ_A are treated as constants; the formulas can be read as applying conditional on a given productivity state by replacing λ_A with $\lambda(a)$.

Households. The representative household has preferences over consumption C_t and labor supply L_t:

$$\mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \left(\log C_t - \frac{L_t^{1+\varphi}}{1+\varphi} \right) dt \right], \tag{23}$$

where $\rho > 0$ is the discount rate and $\varphi > 0$ is the inverse Frisch elasticity. The household trades a nominal risk-free bond B_t earning rate i_t, supplies labor at wage W_t, and receives

profits Π_t from intermediate firms. The nominal flow budget constraint is:

$$\dot{B}_t = i_t B_t + W_t L_t + \Pi_t - P_t C_t.$$

Defining real bond holdings $b_t = B_t/P_t$ and imposing a no-Ponzi condition $\lim_{T\to\infty} \mathbb{E}[M_T b_T] \geq 0$, this integrates to the lifetime present-value budget constraint:

$$\frac{B_0}{P_0} + \Pi_0 + \mathbb{E}_0\left[\int_0^\infty M_t \frac{W_t L_t}{P_t}\, dt\right] \geq \mathbb{E}_0\left[\int_0^\infty M_t C_t\, dt\right], \tag{24}$$

where $\Pi_0 \equiv \mathbb{E}_0[\int_0^\infty M_t(\Pi_t/P_t)\, dt]$ is the present value of profits.[18] Optimising over consumption and labor supply yields the intratemporal condition:

$$\frac{W_t}{P_t} = C_t\, L_t^{\varphi},$$

and, under log utility, the real SDF $M_t = e^{-\rho t}/C_t$. The evolution of the SDF in the jump-diffusion environment is characterised in Proposition 4.1 below.

Firms and Price Setting. A continuum of intermediate producers operate linear technology $y_t = A_t l_t$ and face Dixit–Stiglitz demand with elasticity $\varepsilon > 1$:

$$y_t = \left(\frac{p_t}{P_t}\right)^{-\varepsilon} Y_t.$$

To introduce nominal rigidities, each firm faces Rotemberg quadratic adjustment costs (Rotemberg, 1982) and chooses its price path p_t and labor input l_t to maximise:

$$\sup_{p,l}\ \mathbb{E}\left[\int_0^\infty M_t\left(\frac{p_t}{P_t} y_t - \frac{W_t l_t}{P_t} - \frac{1}{2\eta}\left(\frac{\dot{p}_t}{p_t}\right)^2 Y_t\right) dt\right], \tag{25}$$

where $\eta > 0$ governs the degree of price stickiness. As $\eta \to 0$, adjustment becomes prohibitively costly and prices approach full rigidity. The quadratic structure ensures P_t is absolutely continuous with $dP_t = \pi_t P_t\, dt$, though inflation can itself jump in response to disaster arrivals. The full derivation is in Online Appendix E.3.

Output Gap. Under flexible prices, profit maximisation yields a constant markup $\mu = \varepsilon/(\varepsilon - 1)$ over marginal cost and efficient output:

$$Y_t^n = \bar{Y} A_t, \qquad \bar{Y} := \left(\frac{\varepsilon - 1}{\varepsilon}\right)^{\frac{1}{1+\varphi}}.$$

[18] The full derivation is in Online Appendix E.1.

The log output gap measures the deviation of actual output from this flexible-price benchmark:

$$x_t \equiv \log\left(\frac{Y_t}{Y_t^n}\right) = \log C_t - a_t - \log \bar{Y}. \tag{26}$$

In the jump-diffusion environment, both the output gap and inflation follow jump-diffusions:

$$dx_t = \mu_{x,t}\,dt + \sigma_{x,t}\,dZ_t + \kappa_{x,t}\,dN_t, \qquad d\pi_t = \mu_{\pi,t}\,dt + \sigma_{\pi,t}\,dZ_t + \kappa_{\pi,t}\,dN_t.$$

A labor continuity condition is imposed: employment does not jump at disaster times ($\Delta L_t = 0$), so consumption inherits TFP jumps one-for-one and the output gap is continuous at disaster arrivals.

Asset Pricing. To connect the model to financial market observables, I introduce an aggregate equity claim with dividends modeled as a levered power of consumption:

$$D_t \equiv \chi\, C_t^{\phi_d}, \qquad \phi_d \geq 1, \tag{27}$$

where ϕ_d governs leverage and χ scales dividend levels.[19] This object is useful because it lets the model speak directly to risk premia and valuation changes in financial markets. The equity premium is characterised in Proposition 4.2 below. As λ_t rises, the disaster component widens, equity prices fall, and household wealth declines. The same state variable then enters the neutral rate through the IS curve, linking the asset-pricing and macroeconomic sides of the model through a common risk wedge. Dividend leverage ϕ_d amplifies both Gaussian and disaster components, connecting the model to the empirical equity premium in the quantitative analysis.

Monetary Policy. The central bank sets the nominal rate i_t following a Taylor rule in gap form:

$$i_t - \pi_t = r_t^\star + \phi_x x_t + \phi_\pi(\pi_t - \pi^\star), \tag{28}$$

where $\pi^\star$ is the inflation target and (ϕ_x, ϕ_π) are the policy response coefficients. The intercept $r_t^\star$ is the time-varying risk-adjusted neutral real rate characterised in Proposition 4.3 below. It is through this intercept that the central bank responds to time-varying disaster risk and the associated repricing of financial claims: a rule that accurately tracks $r_t^\star$ accommodates the decline in demand induced by rising risk premia and falling wealth, while a rule that ignores it generates systematic policy errors. The optimal specification of this inter-

[19] This specification is standard in the rare-disaster macro-finance literature (Barro, 2006b; Wachter, 2013).

cept and the welfare consequences of misspecification are the central questions addressed in Section 5.

4.2 Equilibrium Dynamics

Combining household optimality, firm pricing, and market clearing yields a continuous-time three-equation New Keynesian system augmented with disaster risk. The propositions below characterise the SDF, the equity premium, the IS curve, and the NKPC in turn.

Proposition 4.1 (Stochastic Discount Factor with Disasters)**.** *Under log utility with consumption jumps $J_C = J_A$, the real SDF evolves as:*

$$\frac{dM_t}{M_{t^-}} = -r_t\,dt - \sigma_{C,t}\,dZ_t + \left(e^{-J_A} - 1\right) d\tilde{N}_t, \tag{29}$$

where $r_t = i_t - \pi_t$ is the real rate, $d\tilde{N}_t \equiv dN_t - \lambda_t\,dt$ is the compensated Poisson increment, and $\sigma_{C,t} \equiv \sigma_A(1 + x'(a_t))$ is the state-dependent consumption diffusion loading. The jump term $e^{-J_A} - 1 > 0$ captures the discrete rise in marginal utility when a disaster arrives and consumption falls.

Proof. See Online Appendix E.2. □

Proposition 4.2 (Equity Premium with Disasters)**.** *Let dR_t^e denote the instantaneous total equity return. The equity premium satisfies:*

$$\boxed{\frac{\mathbb{E}_t[dR_t^e]}{dt} - r_t = \underbrace{\phi_d\,\sigma_{C,t}^2}_{\text{Diffusion Risk}} + \underbrace{\lambda_t\left(e^{-J_A} - 1\right)\left(1 - e^{\phi_d J_A}\right)}_{\text{Disaster Risk}}.} \tag{30}$$

Proof. See Online Appendix E.6. □

The equity premium decomposes into Gaussian and disaster components, each amplified by dividend leverage ϕ_d. More importantly for the paper's mechanism, the same λ_t that raises the disaster component of (30) also enters the neutral rate through the IS curve below. The model therefore makes explicit that rising downside risk is not merely a financial market phenomenon: it is the same state variable that moves required returns, asset valuations, wealth, and aggregate demand.

Proposition 4.3 (IS Curve with Disasters)**.** *The drift of the output gap satisfies:*

$$\boxed{\mu_{x,t} = \underbrace{i_t - \pi_t - r_t^\star}_{\text{Policy Stance}} + \underbrace{\frac{1}{2}\sigma_{C,t}^2 + \lambda_t \mathbb{E}\left[e^{-J_A} - 1\right]}_{\text{Risk Wedge}},} \tag{31}$$

where the time-varying neutral real rate is:

$$\boxed{r_t^\star = \underbrace{\rho + \mu_A(a_t)}_{\text{Ramsey Rate}} - \underbrace{\frac{1}{2}\sigma_{C,t}^2}_{\text{Precautionary}} - \underbrace{\lambda_t \mathbb{E}\left[e^{-J_A} - 1\right]}_{\text{Disaster Wedge}},} \tag{32}$$

with $\lambda_t = \lambda(a_t)$, $\sigma_{C,t} = \sigma_A(1 + x'(a_t))$, *and* $\mu_A(a_t) = \kappa_a(\bar{a} - a_t)$.

Proof. See Online Appendix E.4. □

Equation (32) identifies two channels through which disaster risk depresses the neutral rate below the Ramsey benchmark $\rho + \mu_A(a_t)$. The precautionary term $-\frac{1}{2}\sigma_{C,t}^2$ is the standard Jensen's inequality correction, state-dependent because $\sigma_{C,t}$ varies with the slope of the equilibrium gap function. The disaster wedge $-\lambda_t \mathbb{E}[e^{-J_A} - 1]$ reflects the expected surge in marginal utility at a disaster arrival and is countercyclical because $\lambda_t = \lambda(a_t)$ rises as productivity falls. These are the continuous-time counterparts of the risk-wedge terms in Proposition 2.1 of the stylised model.

The important economic implication is that equation (32) gives the dynamic version of the paper's central mechanism. When disaster risk rises, households value safe claims more highly, risky claims must offer a larger premium, and asset prices fall. That repricing reduces wealth and weakens demand, so the neutral real rate relevant for monetary policy declines. In this sense, the IS curve does not simply add a disaster term to an otherwise standard New Keynesian system; it embeds a risk-premium/wealth-effect channel directly into the determination of $r_t^\star$.

Proposition 4.4 (New Keynesian Phillips Curve)**.** *A local linearisation around the zero-inflation steady state yields:*

$$\boxed{\mu_{\pi,t} = \rho\,\pi_t + \kappa\,x_t - \lambda_t\,\mathbb{E}_t[\kappa_{\pi,t}],} \tag{33}$$

where $\kappa := \eta(\varepsilon - 1)(1 + \varphi)$ *is the slope coefficient and* $\kappa_{\pi,t}$ *is the inflation jump size at a disaster arrival.*

Proof. See Online Appendix E.5. □

The term $-\lambda_t \mathbb{E}_t[\kappa_{\pi,t}]$ captures the expected inflation jump. While disaster arrivals alter the composition of inflation dynamics, they leave the expected instantaneous change unaffected: $\mathbb{E}_t[d\pi_t] = (\rho\,\pi_t + \kappa\,x_t)dt$, which is the same expression as in the standard continuous-time NKPC without disasters. Equations (31), (33), and the Taylor rule (28) form a closed three-equation system for the output gap and inflation given the state a_t. The role of disasters in this system is therefore concentrated where the mechanism suggests it should be: in the movement of the neutral rate and financial valuations, rather than in a fundamentally different Phillips curve.

4.3 Equilibrium Definition and Solution

Definition 4.1 (Competitive Equilibrium)**.** *An equilibrium is a collection of adapted processes* $\{C_t, L_t, Y_t, W_t, P_t, \pi_t, i_t, x_t, M_t\}$ *such that:*

1. *Households maximise* (23) *subject to* (24)*.*
2. *Firms maximise* (25) *subject to demand and technology constraints.*
3. *Markets clear:* $Y_t = C_t$ *and* $Y_t = A_t L_t$*.*
4. *The IS curve* (31)*, NKPC* (33)*, and Taylor rule* (28) *hold.*
5. *Labor does not jump at disaster times:* $\Delta L_t = 0$*.*

The equilibrium is solved numerically using a global nonlinear method. I conjecture Markov equilibria in which the output gap and inflation are functions of log TFP: $x_t = x(a_t)$ and $\pi_t = \pi(a_t)$. The labor continuity restriction implies $x(a + J_A) = x(a)$ for all a, so the output gap is continuous at disaster arrivals and consumption inherits TFP jumps one-for-one.[20]

Under Markov equilibrium, both $\sigma_{C,t}$ and $r_t^\star$ are genuinely state-dependent. Since $x_t = x(a_t)$ and da_t carries diffusion $\sigma_A\,dZ_t$, Itô's lemma yields:

$$\sigma_{x,t} \equiv \sigma_A\,x'(a_t), \qquad \sigma_{C,t} \equiv \sigma_A\big(1 + x'(a_t)\big), \tag{34}$$

so consumption volatility equals $\sigma_A + \sigma_{x,t}$ and inherits the slope of the equilibrium gap function. The neutral rate $r_t^\star = r^\star(a_t)$ is similarly time-varying through both $\lambda_t = \lambda(a_t)$ and $\mu_A(a_t) = \kappa_a(\bar{a} - a_t)$. In Section 5, these objects are computed directly from the numerical solution $x'(a_i)$ at each grid point, and the estimated mapping $r^\star(a)$ is consistent

[20] The solution algorithm and computational details are in Online Appendix E.7.

with the signs implied by (32). In particular, states with higher disaster intensity are states with higher risk premia, lower asset valuations, and a lower policy-relevant neutral rate. That mapping is the quantitative counterpart of the common wedge derived in the stylised model.

5 Quantitative Analysis

This section takes the continuous-time New Keynesian model to the data. The goals are threefold: to calibrate the model so that it jointly matches key asset-pricing and business-cycle moments, to quantify the importance of disaster-driven movements in risk premia for the neutral rate, and to evaluate what happens to output and welfare when the central bank fails to accommodate the resulting change in the policy-relevant $r_t^\star$. The quantitative model extends the analytical framework of Section 4 along three dimensions. First, preferences are generalised from log utility to Epstein–Zin, so that risk aversion and intertemporal substitution can be calibrated independently. Second, the productivity and disaster processes are made state-dependent to generate realistic persistence in valuations and countercyclical risk premia. Third, the policy rule intercept is allowed to vary with the estimated neutral rate, enabling a direct comparison between a central bank that tracks risk-premium-driven movements in $r_t^\star$ and one that does not.

5.1 Model Extensions

Epstein–Zin preferences. Log utility, while analytically convenient, forces risk aversion and the elasticity of intertemporal substitution (EIS) to move in lockstep. Here I replace it with Epstein–Zin preferences, parameterised by risk aversion γ and EIS ψ, which can be set independently. The distinction matters in a disaster-risk environment: a high γ amplifies the effect of downside risk on required returns, asset prices, and the disaster wedge in the neutral rate, while a high ψ governs the household's willingness to shift consumption across time in response to real-rate movements. Under log utility both parameters equal one, and the price of tail risk is understated relative to the data. Full equilibrium conditions under Epstein–Zin are in Online Appendix H.1.

State-dependent processes. Section 4 kept the disaster intensity λ_A and the drift μ_A constant to obtain transparent closed-form expressions. The quantitative model generalises in

two ways. The productivity state a_t is treated as a stationary deviation from trend, with mean-reverting dynamics:

$$da_t = \kappa_a(\bar{a} - a_t)\,dt + \sigma_A\,dZ_t + J_A\,dN_t, \qquad dN_t \sim \text{Poisson}(\lambda(a_t)\,dt), \tag{35}$$

where $\kappa_a > 0$ is the mean-reversion speed and $\bar{a} = 0$ is the unconditional mean. Mean reversion ensures stationarity, so that moments such as the persistence of the log dividend–price ratio are well defined in long simulations. Disaster arrivals are countercyclical: conditional on state a_t, the Poisson intensity is

$$\lambda(a) = \lambda_0 \exp(-\chi_\lambda a), \tag{36}$$

with $\chi_\lambda > 0$, so that bad states carry higher disaster risk. This state-dependence is central for the paper's mechanism: in bad states, disaster intensity rises, risk premia widen, valuations fall, and the risk-adjusted neutral rate declines. The constant-λ_A model of Section 4 is the special case $\chi_\lambda = 0$, $\kappa_a = 0$; all the closed-form results there continue to apply state-by-state upon replacing λ_A with $\lambda(a)$ and $\rho + \mu_A$ with $\rho + \kappa_a(\bar{a} - a)$.

Policy intercept rules. The central question of the quantitative exercise is how the Taylor-rule intercept should respond to time-varying tail risk. The rule takes gap form:

$$i_t - \pi_t = \underbrace{r_t^\star}_{\text{intercept}} + \phi_x x_t + \phi_\pi(\pi_t - \pi^\star),$$

and the intercept can be specified in two ways. Under a *fixed intercept* the central bank sets $r_t^\star = \bar{r} = 2\%$, ignoring time variation in risk premia and therefore ignoring the fall in the true risk-adjusted neutral rate when disaster risk rises. The policy rate then fails to decline sufficiently, opening a persistent wedge between the stance of policy and the true neutral rate. Through the IS curve, that wedge contracts output and inflation. Under a *time-varying intercept* the central bank tracks the full risk-adjusted neutral rate:

$$r_t^\star = \rho + \mu_A - \tfrac{1}{2}\sigma_{C,t}^2 - \lambda_t(e^{-J_A} - 1),$$

accommodating the disaster wedge in real time and keeping the output gap closed. These two regimes bracket the space of possible policy responses and form the basis of the welfare comparison that follows.

5.2 Calibration Strategy

The model is calibrated in two steps. A set of standard New Keynesian parameters governing preferences, nominal rigidities, and the policy rule are fixed at conventional values reported in Table 2. The remaining parameters—those governing disaster risk, technology dynamics, Epstein–Zin preferences, and dividend leverage—are then estimated by simulated method of moments (SMM).

Table 2: Directly Calibrated Parameters

Parameter	Description	Value	Source
Preferences and Technology			
ρ	Discount rate (annual)	0.02	Standard
φ	Inverse Frisch elasticity	1.0	Galí (2015)
ε	Elasticity of substitution	10.0	Markup = 11%
Nominal Rigidities			
η	Rotemberg adjustment cost	0.00556	Targets $\kappa \approx 0.10$
Monetary Policy			
ϕ_π	Taylor coefficient (inflation)	1.5	Taylor (1993)
ϕ_x	Taylor coefficient (output gap)	0.5	Taylor (1993)
$\pi^\star$	Inflation target (annualised)	0.0338	Data mean inflation

Notes: Standard parameters are set to conventional values in the New Keynesian literature. The Rotemberg parameter η is chosen to target a Phillips curve slope $\kappa = \eta(\varepsilon - 1)(1 + \varphi) \approx 0.10$. The inflation target $\pi^\star$ is fixed to match average inflation in the data and is not estimated in SMM.

The SMM procedure targets 13 moments computed from U.S. quarterly data (1954Q1–2024Q4), spanning the mean, standard deviation, and first-order autocorrelation of the log dividend–price ratio, equity returns, the short rate, excess returns, and output-cycle volatility.[21]

State processes in the quantitative model. Section 4 presented the model under constant disaster intensity λ_A and constant drift μ_A to obtain clean closed-form expressions. In the estimation, both are generalised. The productivity state a_t follows mean-reverting dynamics as in equation (35), and disaster arrivals are countercyclical as in equation (36). The estimated parameter vector is

$$\theta = \{\lambda_0, \chi_\lambda, J_A, \sigma_A, \kappa_a, \psi, \gamma, \phi_d\},$$

[21] Complete moment definitions and the SMM algorithm are reported in Online Appendix H.3.

where λ_0 is the baseline disaster intensity, χ_λ governs its countercyclicality, J_A is the disaster jump size, σ_A is diffusion volatility, κ_a is mean reversion, (ψ, γ) are the Epstein–Zin preference parameters, and ϕ_d is dividend leverage. Table 3 reports the SMM estimates.

Table 3: SMM Parameter Estimates

Parameter	Estimate	Description
Disaster Risk		
λ_0	0.0278	Baseline disaster intensity (annual)
χ_λ	3.16	Countercyclical slope
J_A	−0.148	Disaster jump size (log TFP)
Technology Dynamics		
σ_A	0.032	Diffusion volatility
κ_a	0.081	Mean-reversion speed
Preferences		
γ	6.0	Risk aversion
ψ	1.95	EIS
Asset Pricing		
ϕ_d	3.45	Dividend leverage

Notes: Disaster intensity is $\lambda(a) = \lambda_0 \exp(-\chi_\lambda a)$ (clipped in simulation to $[10^{-10}, 0.3]$). The estimated jump size implies a productivity drop of about $1 - e^{J_A} \approx 14\%$ conditional on a disaster.

Several features of the estimates deserve comment. Baseline disaster intensity is about 2.8% per year, higher than in many calibrations of rare-disaster models but within the range estimated by Barro and Liao (2021) from option prices. The countercyclical slope $\chi_\lambda = 3.16$ ensures that disaster risk rises substantially in bad states, generating countercyclical risk premia and time-varying movements in valuations endogenously. The estimated jump size $J_A = -0.148$ implies a roughly 14% decline in productivity conditional on disaster arrival. The preference estimates pair moderate risk aversion ($\gamma = 6$) with a high EIS ($\psi \simeq 1.95$), a combination that amplifies the effect of downside risk on asset prices and the neutral rate. Dividend leverage $\phi_d \simeq 3.45$ translates the underlying productivity risk into realistic equity return volatility.

5.3 Moment Fit: Model vs. Data

Table 4 compares the targeted empirical moments to their model-implied counterparts.

The model matches the equity premium closely: the mean excess return is 1.17% per quarter in the model versus 1.25% in the data, and the volatility of excess returns is 8.23% versus 8.55%. Equity return levels and volatilities are similarly well matched. The mean

Table 4: Targeted Moments: Model vs. Data

Variable	**Moment**	**Data**	**Model**
Equity Premium (quarterly)			
r^x	Mean	1.25%	1.17%
r^x	Std. dev.	8.55%	8.23%
Equity Returns (quarterly)			
r^s	Mean	2.48%	2.57%
r^s	Std. dev.	8.50%	8.23%
Short Rate (quarterly)			
r^b	Mean	1.23%	1.40%
Valuation Ratio			
$\log(D/P)$	Mean	−3.67	−2.58
$\log(D/P)$	Std. dev.	0.40	0.20
$\log(D/P)$	Autocorr.	0.976	0.941

Notes: Moments are computed from quarterly U.S. data (1954Q1–2024Q4). Returns are quarterly log returns. $\log(D/P)$ denotes the log dividend–price ratio. Model moments are computed from simulated economies at quarterly frequency.

short rate is close to its empirical counterpart (1.40% per quarter in the model versus 1.23% in the data). The persistence of equity valuations is captured reasonably well, with the autocorrelation of $\log(D/P)$ at 0.941 in the model versus 0.976 in the data.

These are the margins that matter most for the paper's mechanism. The model generates plausible time variation in expected excess returns and valuations, so it can speak credibly to the way financial repricing feeds into the neutral rate and monetary policy. The main shortcomings concern the valuation ratio itself. The model understates both the mean and the standard deviation of $\log(D/P)$: the mean is -2.58 versus -3.67 in the data, and the standard deviation is 0.20 versus 0.40. The model thus generates dividend yields that are too high and too stable relative to the data. This likely reflects the interaction of dividend leverage with the unconditional disaster probability: with $\phi_d = 3.45$ and $\lambda_0 = 2.78\%$, the discount applied to equity claims is somewhat lower than what the data require to match the level of the price–dividend ratio.

In sum, the updated calibration achieves a substantially improved fit for equity premia, return dynamics, and the short rate, while the valuation ratio remains a margin where additional model ingredients, such as long-run risk or habit formation, could improve the match. For the purposes of this paper, however, the calibration is strong where it needs to be strongest: on the joint behaviour of returns, risk premia, and interest rates.

5.4 Drivers of the Neutral Rate

The neutral-rate formula from Section 4 delivers sharp predictions about how each source of risk affects $r_t^\star$. This subsection quantifies those effects using the calibration in Table 3. The main object of interest is the wedge through which financial repricing feeds into monetary policy.

Evaluating at $a = \bar{a}$, the Ramsey rate is 2.0%, the diffusion wedge subtracts about 5 basis points, and the disaster wedge subtracts about 44 basis points, yielding a neutral rate of approximately 1.5% at the unconditional mean. The disaster wedge is thus nearly nine times larger than the diffusion wedge, despite a baseline disaster probability of only 2.8% per year. This asymmetry arises because disasters are low-probability, high-impact events that generate large spikes in marginal utility and therefore large movements in required returns and asset valuations.

The comparative statics sharpen the point. At a given state the sensitivities are:

$$\frac{\partial r^\star(a)}{\partial \sigma_{C,t}} = -\sigma_{C,t}, \qquad \frac{\partial r^\star(a)}{\partial \lambda(a)} = -(e^{-J_A} - 1), \qquad \frac{\partial r^\star(a)}{\partial J_A} = \lambda(a) e^{-J_A}. \tag{37}$$

A 1 percentage point increase in disaster intensity lowers $r^\star$ by approximately 16 basis points, whereas an equivalent increase in diffusion volatility lowers it by only 3.2 basis points, a ratio of 5 to 1. These magnitudes show that crisis-time movements in the neutral rate are driven primarily by the component of risk that commands a premium in financial markets, not by small changes in Gaussian volatility.

An equivalence exercise makes the same point in more concrete terms. To replicate the 16 basis point decline in $r^\star$ caused by a 1 percentage point increase in $\lambda(a)$ using diffusion volatility alone, consumption volatility would need to rise from 3.2% to about 6.5% annually, roughly doubling, far above any empirical estimate. Even a 10 basis point decline requires about 2.3 percentage points of additional consumption volatility (from 3.2% to 5.5%). The same 10 basis point decline can instead be generated by a 0.6 percentage point increase in disaster intensity, well within the range of option-implied variation documented in Section 3.

The quantitative lesson is therefore straightforward: the neutral rate is highly sensitive to tail-risk-driven movements in risk premia and only weakly sensitive to diffusion risk. This is why a monetary policy rule that ignores financial repricing can make large errors even if conventional macro indicators move only gradually.

5.5 Narrative Counterfactual Experiment

This section complements the sample-wide Taylor-rule regressions with an episode-level counterfactual in the spirit of Del Negro et al. (2017a), focusing on the window from 2007Q2 to 2008Q1. Option-implied disaster probabilities rose by roughly 10 percentage points over this period, precisely when the Federal Reserve began cutting the federal funds rate faster than any standard Taylor rule would prescribe. The model interprets this episode as a tail-risk shock that lowers the risk-adjusted neutral rate $r_t^\star(a)$ and, in turn, the natural nominal rate $i_t^\star(a) \equiv r_t^\star(a) + \pi^\star$. Crucially, the size and persistence of the shock are disciplined by option-implied disaster probabilities rather than by the macro outcomes used for validation, so the exercise does not reverse-engineer the data.

I parameterise the shock as an unanticipated jump in disaster probability that then mean-reverts deterministically, and I separate the shock path from the policy rule so that the same tail-risk scenario can be run under alternative rules. Three quarterly data series define the targets: output (log real personal consumption expenditures, detrended over 2000Q1–2012Q3 and normalised to zero in the pre-shock quarter), inflation (annualised GDP-deflator inflation minus 2 percent, following Del Negro et al. 2017a), and the quarterly average effective federal funds rate.

To measure how much weight the Fed placed on tail-risk movements in the neutral rate, I introduce a single parameter $\omega \in [0, 1]$ governing the intercept of the real-rate rule:

$$i_t(a) - \pi_t(a) = \Big[\bar{r} + \omega\big(r_t^\star(a) - \bar{r}\big)\Big] + \phi_x x_t(a) + \phi_\pi\big(\pi_t(a) - \pi^\star\big), \tag{38}$$

where $\bar{r}$ is the steady-state real rate. At $\omega = 0$ the intercept is fixed and policy ignores the tail-risk-driven fall in $r_t^\star$; at $\omega = 1$ the intercept tracks $r_t^\star(a)$ one-for-one, fully accommodating the shift in the neutral rate. I choose $\hat{\omega}$ by minimising the mean-squared error between the model's perfect-foresight impulse responses and the data over the episode window.

Figure 9 reports the results. The left column shows the three data series; the right column shows the model-implied paths under the fitted $\hat{\omega} = 0.4$, anchored to observed levels in the pre-shock quarter. The estimate implies that the Fed placed about 40 percent weight on the tail-risk-driven movement in the neutral rate—enough to generate meaningful accommodation, but well short of fully tracking $r_t^\star$.

The model matches the output and interest-rate dynamics well: output declines relative to trend while the funds rate falls rapidly, reproducing the broad co-movement in the data. Inflation is the harder margin. Even at the fitted $\hat{\omega}$, the model's implied inflation path

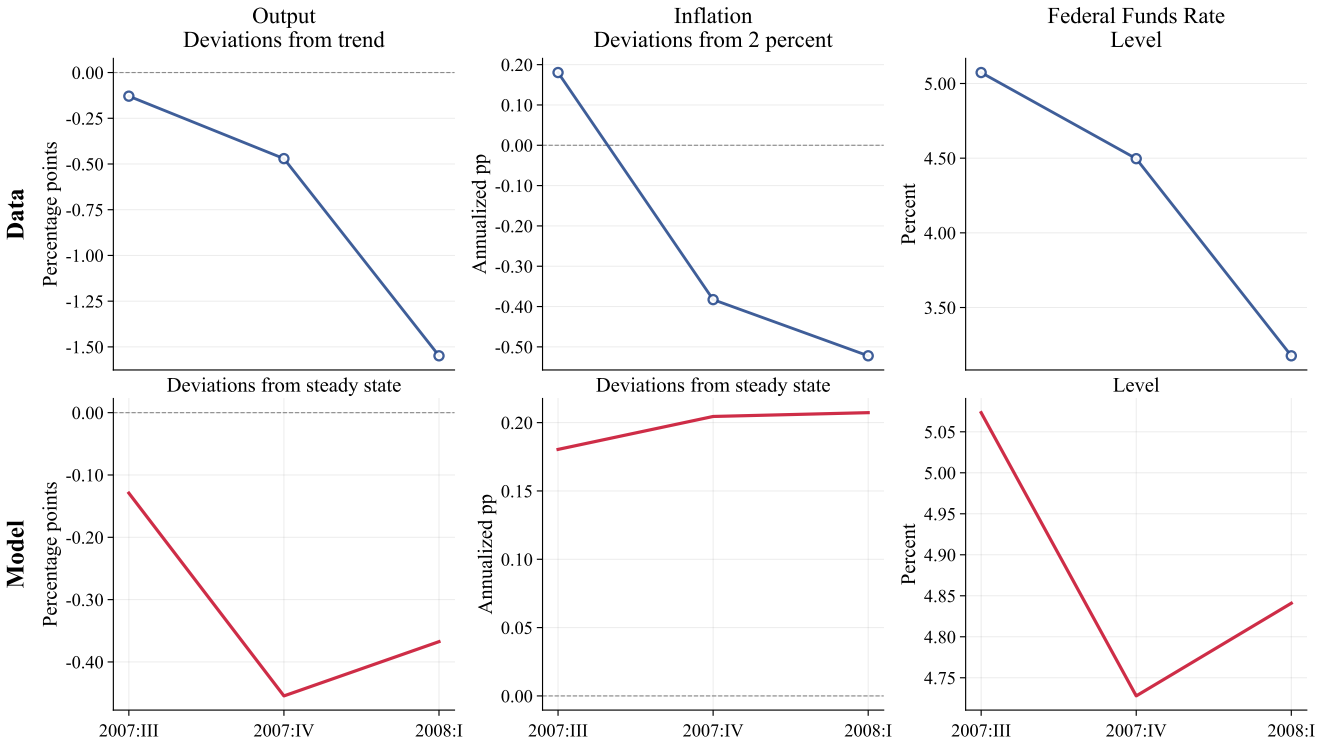


Figure 9: Narrative episode objects in the data (left) and the corresponding model-implied perfect-foresight transition under the fitted accommodation parameter $\hat{\omega} = 0.4$ (right). Output is the deviation of $\log(C + I)$ from a linear trend; inflation is annualised GDP-deflator inflation minus 2 percent; the policy rate is the quarterly average effective federal funds rate. Model paths are anchored to the same pre-shock quarter for comparability.

diverges from the data in level and timing. This is not unexpected. The model features only a single shock and a relatively flat New Keynesian Phillips curve, so generating the observed disinflation would require output-gap movements far larger than those in the data. Moreover, inflation over 2007–2009 reflected forces deliberately excluded from this one-shock experiment—commodity-price swings, cost-push disturbances, and shifts in inflation expectations—which further limit the model's ability to match this margin with a tail-risk shock alone.

The main payoff of the exercise is therefore not a perfect historical reconstruction but a disciplined inference about the policy intercept. The fitted rule suggests that the Fed did respond to the fall in the risk-adjusted neutral rate, but only partially. Under $\omega = 0$, the policy rate remains too high as disaster risk rises, amplifying the contraction: output falls by roughly twice as much as under the fitted rule. Under $\omega = 1$, the intercept tracks the neutral rate in full, largely neutralising the macroeconomic consequences of the shock. Read this way, the episode supports the paper's central message: the relevant policy question is whether the central bank accommodates the fall in $r_t^\star$ induced by financial repricing.

5.6 Counterfactual Analysis: Policy Intercept and Disaster Risk

This section uses perfect-foresight transition dynamics to compare two polar policy responses to a tail-risk surge. The experiment subjects the calibrated model to a transitory disaster-probability shock—a 10 percentage point increase on impact, decaying geometrically with a half-life of eight quarters—and traces the 80-quarter transition under two regimes. Both regimes use the real-rate rule

$$i_t - \pi_t = \left[\bar{r} + \omega\left(r_t^\star - \bar{r}\right)\right] + \phi_x x_t + \phi_\pi(\pi_t - \pi^\star), \tag{39}$$

evaluated at the two extremes: $\omega = 0$ (fixed intercept, tail risk ignored) and $\omega = 1$ (time-varying intercept tracking $r_t^\star$).

Real activity and inflation. Figure 10 reports the impulse responses of the output gap, inflation, and the nominal policy rate.

Under the fixed-intercept regime, the tail-risk shock lowers $r_t^\star$ but the policy rule does not respond, opening a contractionary real-rate gap. The output gap falls by roughly 1.2 percentage points on impact—about three times the 0.4 percentage point decline observed in the data—and inflation drops by about 0.2 annualised percentage points before slowly returning to target. Under the time-varying intercept, the intercept moves with $r_t^\star$, keeping the real-rate gap close to zero and effectively stabilising both the output gap and inflation throughout the transition.

A notable feature of Figure 10 is that the headline policy-rate path looks similar across regimes. The reason is that the source of easing differs. Under $\omega = 0$ the rule cuts rates reactively, after slack and disinflation have already emerged. Under $\omega = 1$ the rule cuts rates pre-emptively through the intercept, because it recognises immediately that the risk-adjusted neutral rate has fallen. This is exactly the distinction between a central bank that responds to realised macro deterioration and one that responds to the underlying wealth-effect shock driving that deterioration.

Financial variables. Figure 11 shows that asset-price responses are nearly identical under both regimes. Equity prices fall by about 40 percent on impact, the model-implied VIX jumps by roughly 20 index points, and the equity risk premium rises by about 12.5 annualised percentage points.

This invariance is important for interpretation. Risk premia and volatility reflect the

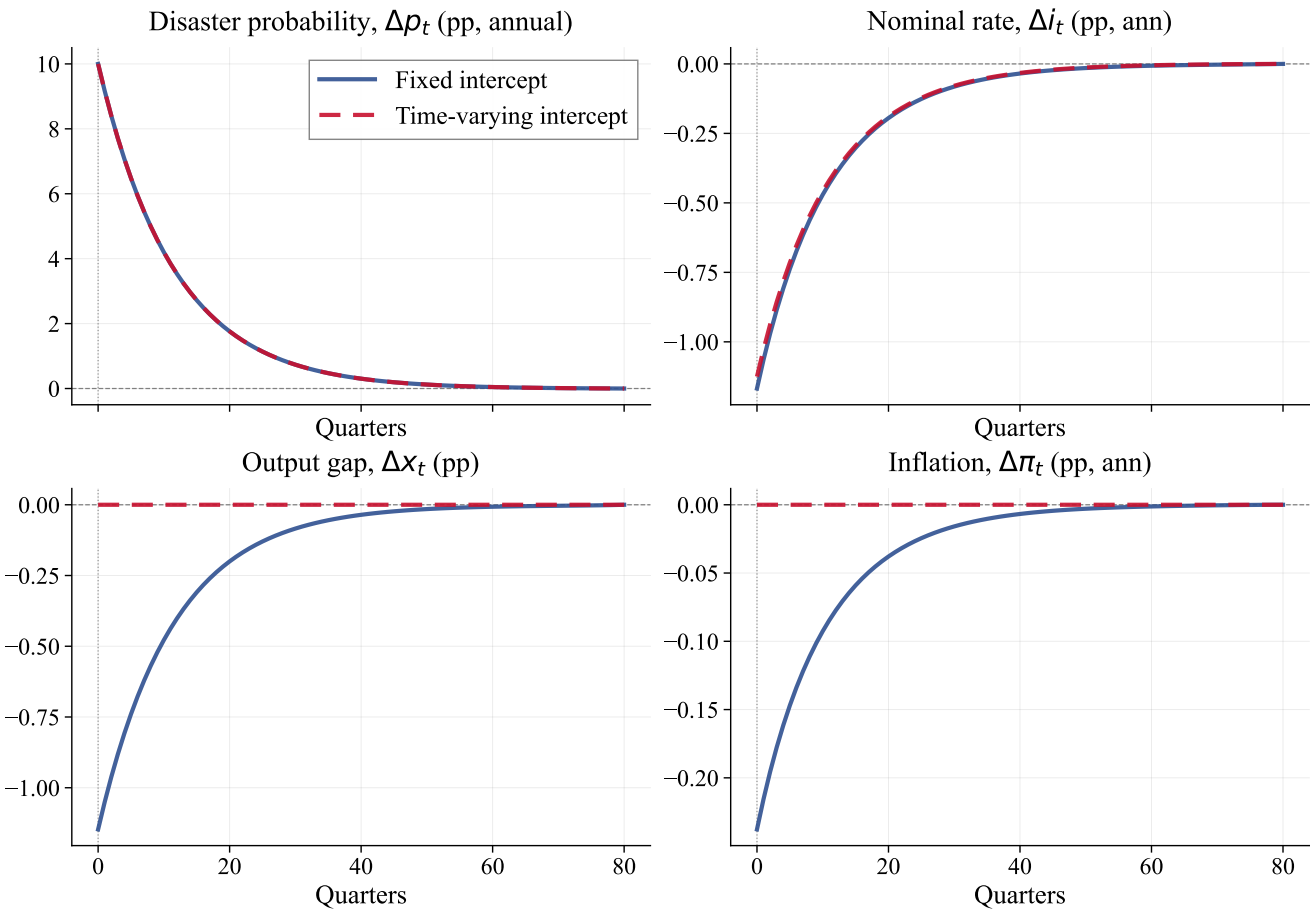


Figure 10: Impulse responses to a transitory 10 percentage point increase in disaster probability under two policy regimes. The fixed-intercept rule ($\omega = 0$) holds the intercept constant; the time-varying-intercept rule ($\omega = 1$) adjusts the intercept to track the risk-adjusted neutral real rate $r_t^\star$.

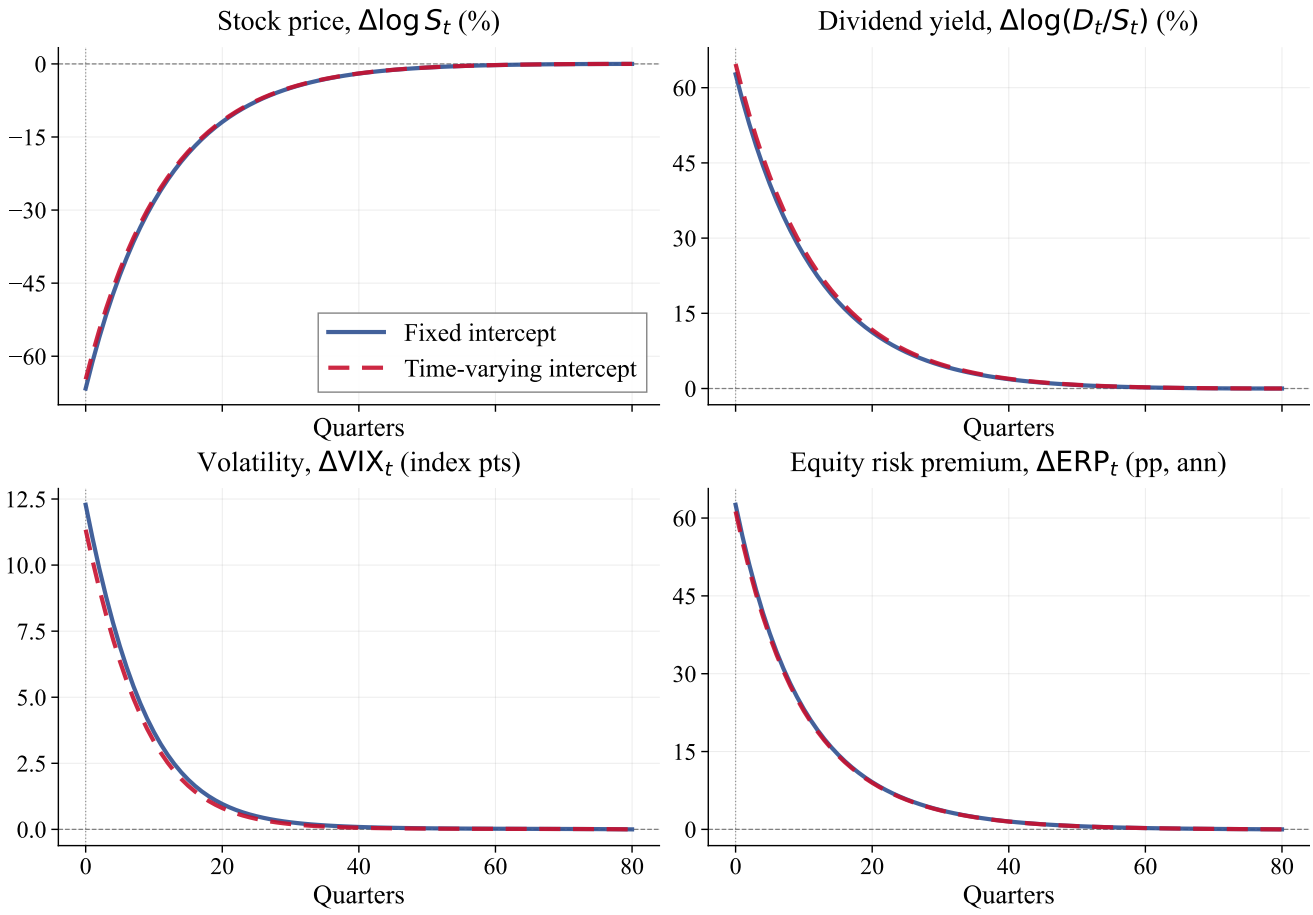


Figure 11: Impulse responses of financial variables under fixed ($\omega = 0$) and time-varying ($\omega = 1$) policy intercepts. Equity valuations, the VIX, and risk premia are driven primarily by the direct effect of disaster risk on the pricing kernel and are largely invariant to the policy regime.

direct effect of tail risk on the pricing kernel, which the policy intercept cannot offset. Monetary policy does not eliminate the financial repricing itself. Instead, it can offset the macroeconomic consequences of that repricing by adjusting to the fall in the neutral rate. In other words, the policy problem is not to compress risk premia, but to prevent risk-premium-driven wealth losses from producing an unnecessarily deep demand contraction.

Real-side adjustments. Figure 12 traces consumption, labour, real wages, and marginal cost under the two regimes.

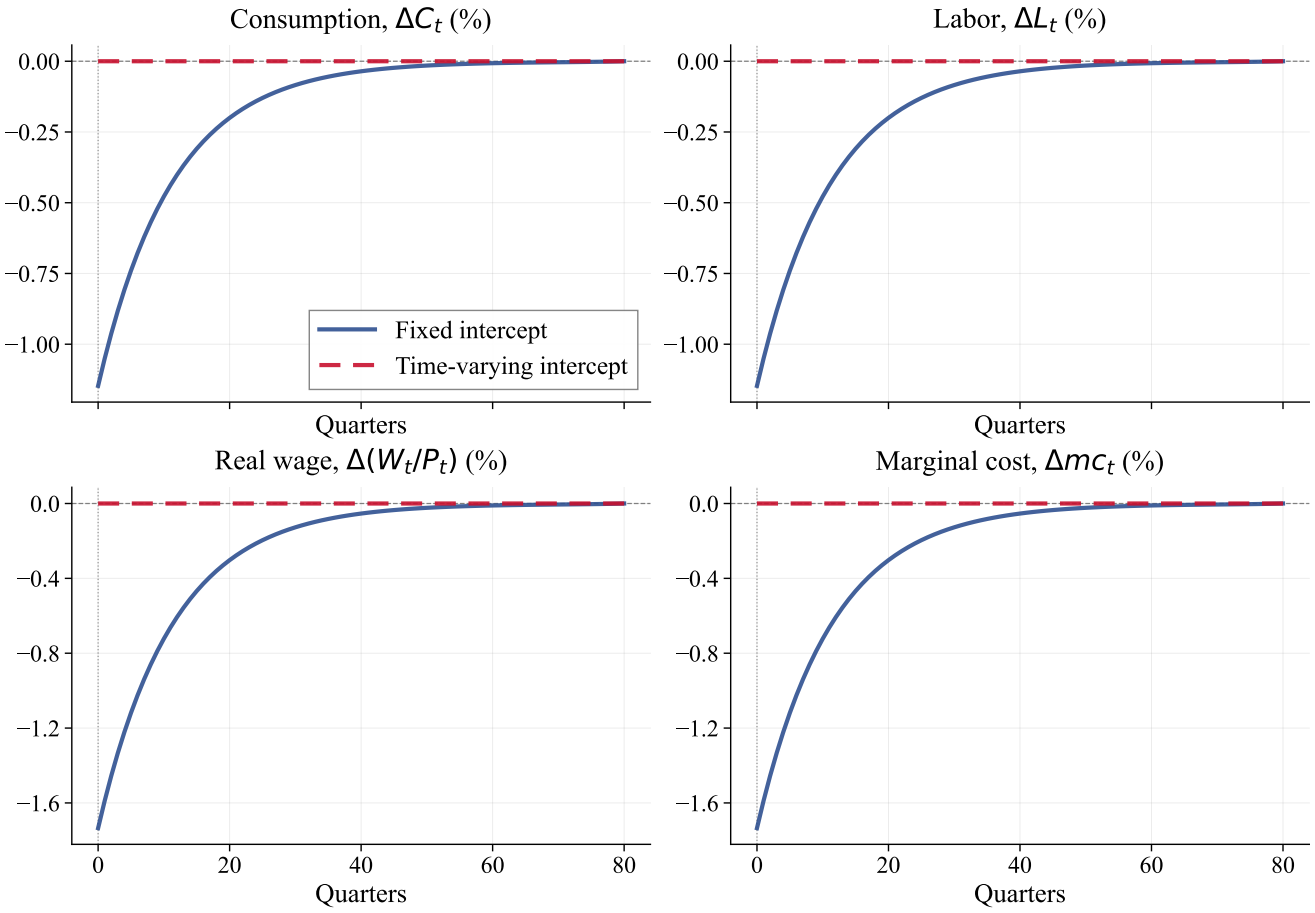


Figure 12: Impulse responses of consumption, labour, real wages, and marginal costs under fixed ($\omega = 0$) and time-varying ($\omega = 1$) policy intercepts.

Under the fixed intercept, consumption and hours fall by about 1.1 percent on impact, and real wages and marginal cost decline by about 1.7 percent. Under full accommodation these movements are shut down almost entirely. This confirms the model's interpretation of the recessionary component of a tail-risk shock: the contraction is not driven mainly by an unavoidable fall in productive capacity, but by a policy-relevant gap between the actual real rate and the lower neutral rate implied by financial repricing. Once policy tracks $r_t^\star$,

nominal rigidities do not by themselves force a deep contraction.

Summary. A tail-risk surge lowers the risk-adjusted neutral real rate. If the policy intercept is fixed, a contractionary real-rate gap opens and the economy contracts sharply. Allowing the intercept to track $r_t^\star$ stabilises output and inflation nearly one-for-one, yet does nothing to prevent the spike in risk premia and volatility, since these reflect the direct repricing of disaster risk. The central bank can therefore insulate the real economy from financial-market turmoil by tracking the risk-adjusted neutral rate, even though it cannot stabilise asset prices themselves. This is the quantitative version of the paper's main claim: movements in tail risk matter for policy because they move risk premia, wealth, and aggregate demand, and hence the neutral rate that a stabilising central bank should track. A practical caveat is that $r_t^\star$ is unobserved in real time; the experiment should therefore be read as a rule-design counterfactual showing the gains available if the central bank can use market-based information to track risk-premium-driven movements in the neutral rate.

6 Estimating the Risk-Adjusted Neutral Rate

The preceding sections established that changes in downside risk move the neutral real rate through a risk-premium and wealth-effect channel, and that ignoring this channel generates welfare-relevant output losses. This section takes that prediction to the data and estimates a historical series for the risk-adjusted neutral rate r_t^* that can be used alongside conventional measures in real time. The goal is to provide a new policy indicator: one that preserves the low-frequency trend-growth component of standard neutral-rate estimates, but also moves at business-cycle frequency with observable changes in financial risk premia, disaster risk, and macroeconomic uncertainty. In that sense, the exercise is the empirical counterpart to the theory. If rising downside risk raises required returns, lowers valuations, and weakens demand through wealth effects, then the neutral rate relevant for monetary policy should fall at exactly those times.

6.1 Definition and Estimation

Recall from the quantitative model that the neutral rate decomposes as

$$r_t^* = \underbrace{\rho + c_g g_t}_{\text{Ramsey component}} - \underbrace{\frac{1}{2}\sigma_{c,t}^2}_{\text{Diffusion wedge}} - \underbrace{\phi_J \lambda_t}_{\text{Disaster wedge}}, \tag{40}$$

where g_t is trend output growth, $\sigma^2_{c,t}$ captures persistent macroeconomic uncertainty (proxied by the VIX), and λ_t is the persistent component of disaster probability from Section 3. Holding risk fixed, faster trend growth raises r^*_t through the standard intertemporal-substitution channel. Holding growth fixed, increases in either diffusion or disaster risk lower r^*_t by strengthening precautionary saving and, in the framework of this paper, by capturing the financial repricing that depresses wealth and aggregate demand. A conventional Laubach–Williams estimate captures only the first force; the risk-adjusted measure captures both.

I estimate the latent components of (40) via a linear-Gaussian state-space model using quarterly U.S. data from 1985Q1 to 2023Q4. The observable vector includes real GDP, inflation, two trend-growth proxies (GDP growth and utilisation-adjusted TFP growth), the option-implied disaster probability, and a VIX-based variance proxy. Five latent states are extracted: potential output, trend growth, persistent disaster risk, persistent diffusion risk, and the output gap. The Kalman filter produces one-sided (real-time) estimates; the Rauch–Tung–Striebel smoother produces two-sided (full-sample) estimates (Hamilton, 1994). Full details of the state-space specification and parameter estimates are in Online Appendix K.1.

6.2 Results

Validation. Figure 13 compares the Ramsey component $\rho + c_g g_t$ against the Laubach–Williams one-sided estimate. The two series co-move closely at low frequencies, confirming that the state-space model recovers the same slow-moving trend-growth variation that drives the Laubach–Williams estimate. This is useful because it shows that the new measure is not discarding the standard benchmark; it nests it. The risk-adjusted estimate departs from Laubach–Williams during crises, and those departures are accounted for entirely by the disaster and diffusion wedges. In other words, the new series differs from conventional estimates precisely when the theory says it should: when financial markets are repricing downside risk and the wealth-effect channel is most relevant for policy.

Full decomposition. Figure 14 presents the complete decomposition. The risk wedges are predominantly negative and widen sharply during episodes of financial stress. During the GFC and COVID-19, the combined disaster and diffusion wedges push r^*_t down by 200–300 basis points below the Ramsey component, consistent with surges in precautionary saving, safe-asset demand, and broader financial repricing. The disaster wedge is the dominant

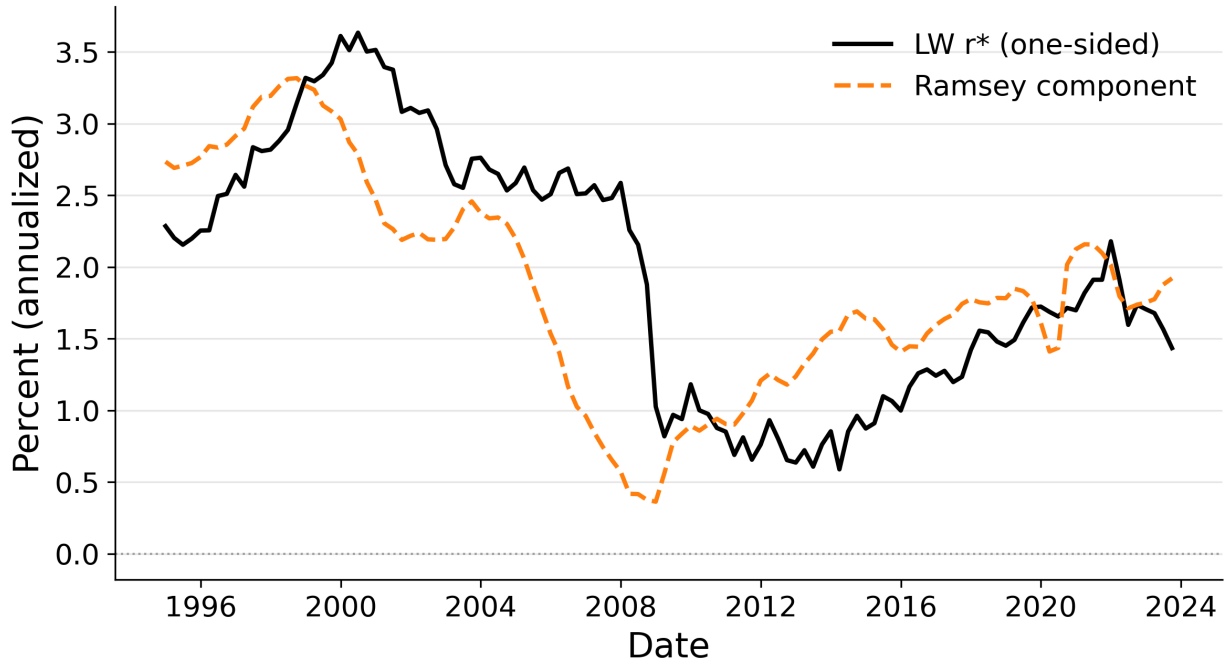


Figure 13: **Ramsey Component vs. Laubach–Williams.** The Ramsey component $\rho + c_g g_t$ tracks low-frequency movements in the Laubach–Williams one-sided estimate.

component in these episodes, which is consistent with the quantitative model's implication that tail-risk variation is more important than diffusion risk for crisis-time movements in the neutral rate.

Comparison to benchmark estimates. Figure 15 places the risk-adjusted neutral rate alongside the Laubach–Williams estimate (Laubach and Williams, 2003), the DSGE-based short-run r^* of Del Negro et al. (2017b), and the ex-ante real federal funds rate. The Laubach–Williams estimate captures low-frequency trend-growth variation but misses the sharp crisis-driven movements by construction. The Del Negro et al. measure does exhibit high-frequency variation, though driven by safety and liquidity premia in Treasury markets rather than the option-implied disaster and uncertainty wedges estimated here.

The comparison changes the interpretation of the policy stance. When the real federal funds rate lies above r_t^*, policy is restrictive in a risk-adjusted sense; when it lies below, policy is accommodative. During recessions, r_t^* falls so rapidly that a policy rate that looks accommodative relative to Laubach–Williams may in fact be close to neutral—or even tight—once tail risk is taken into account. This is exactly the empirical implication of the model: a central bank that looks only at slow-moving trend measures can miss a large crisis-time fall in the neutral rate driven by financial repricing and wealth effects.

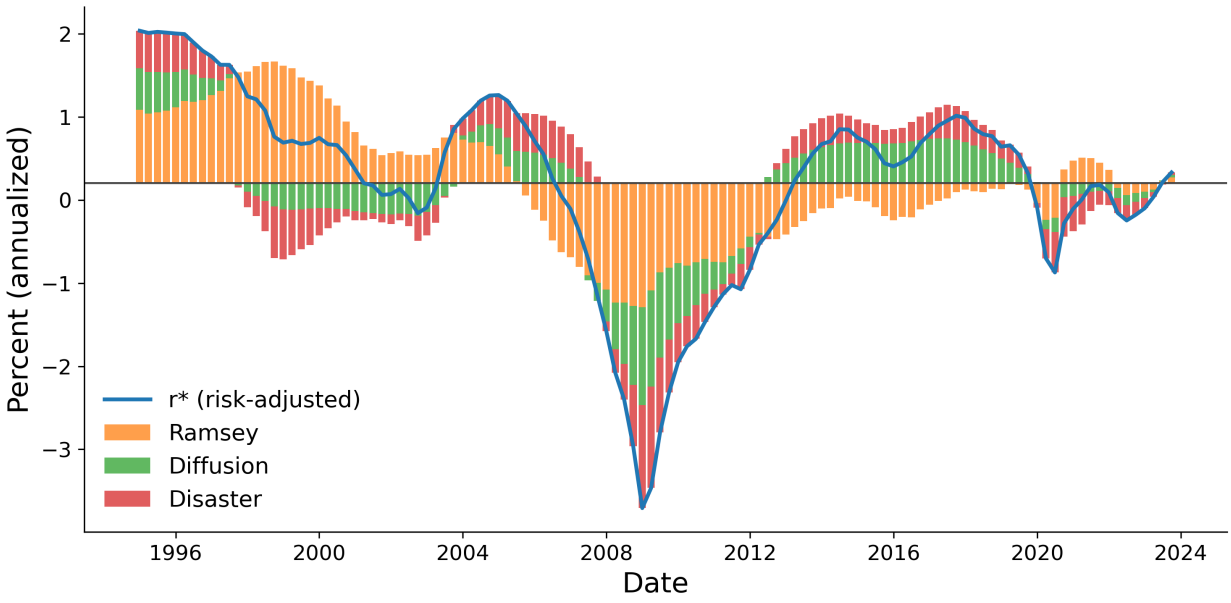


Figure 14: **Risk-Adjusted Neutral Rate Decomposition.** The neutral rate r_t^* decomposed into the Ramsey component and the two risk wedges. The disaster wedge dominates during crises.

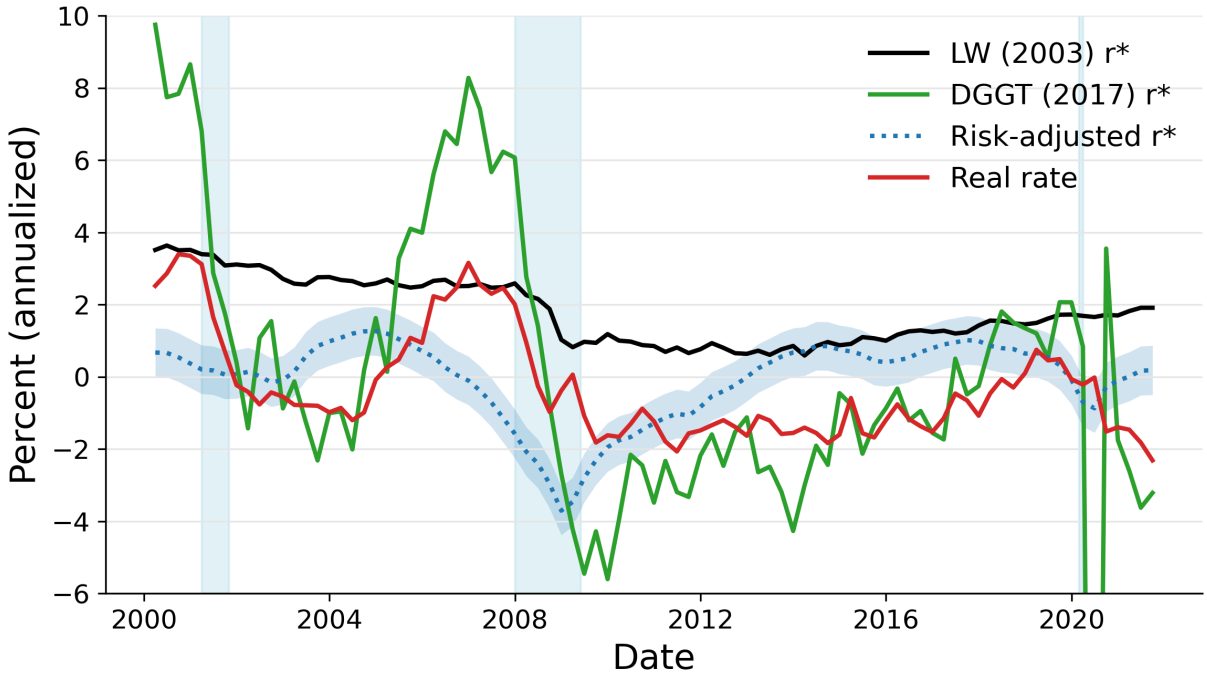


Figure 15: **Risk-Adjusted Neutral Rate and Benchmarks.** The risk-adjusted r_t^* (dotted blue with confidence band) compared to Laubach–Williams (black), Del Negro et al. DSGE-based r^* (green), and the ex-ante real federal funds rate (red). Shaded areas indicate NBER recessions.

6.3 A Risk-Adjusted Taylor Rule

To make the neutral-rate estimate operational, I construct an implied policy-rate path from a balanced-approach rule following Nakamura et al. (2025):

$$i_t^{\text{rule}} = r_t^* + \pi_t + 0.5(\pi_t - 2) + 1.0(y_t - y_t^*), \tag{41}$$

where π_t is 12-month core PCE inflation and $(y_t - y_t^*)$ is the real-time output gap.[22] Figure 16 plots the realised federal funds rate, the Wu–Xia shadow rate, and the rule-implied rate.

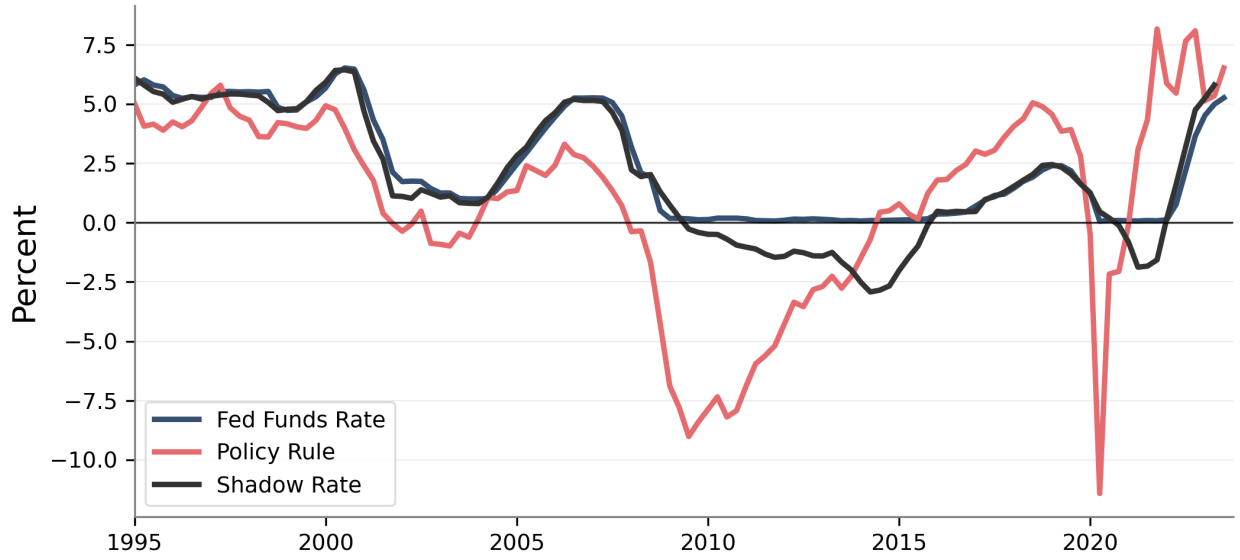


Figure 16: **Federal Funds Rate and Risk-Adjusted Taylor Rule.** The realised federal funds rate (solid black), Wu–Xia shadow rate (dashed), and the balanced-approach rule using risk-adjusted r_t^* (blue).

During normal times the rule-implied rate broadly tracks the federal funds rate. During the GFC and COVID-19, the implied prescription drops sharply as the neutral rate falls with elevated tail risk. Incorporating a risk-adjusted neutral rate therefore alters the recommended policy stance precisely when it matters most: in crisis episodes where conventional rules, anchored to constant or slowly moving neutral rates, would prescribe insufficient easing. Put differently, the policy value of the measure is not that it changes the prescription all the time, but that it changes it in the states of the world where risk premia, wealth effects, and aggregate demand are moving rapidly together. That is the practical payoff of

[22] The specification follows Nakamura et al. (2025) except in the treatment of r_t^*: they use the FOMC's long-run nominal rate from the Summary of Economic Projections minus 2%, whereas I use the risk-adjusted neutral rate estimated here.

the paper's mechanism.[23]

7 Conclusion

Central banks have long practiced risk management, adjusting policy pre-emptively when financial conditions deteriorate and downside risks rise. What has been missing is a formal framework that links those policy decisions to a clear macroeconomic mechanism. This paper provides such a framework. Its central contribution is to show how changes in downside risk move risk premia and asset prices, how those valuation changes transmit to aggregate demand through wealth effects, and how that mechanism in turn shifts the neutral real rate relevant for monetary policy. In the model, the same risk wedge that widens the equity premium also lowers $r_t^\star$. That link between asset pricing and the neutral rate is the paper's core result.

The theory makes this mechanism transparent. When the left tail thickens, the marginal utility of wealth rises in bad states. Investors demand greater compensation for holding risky claims, asset valuations fall, and private wealth declines. With nominal rigidities, that fall in wealth contracts aggregate demand and lowers the equilibrium real rate. A central bank that attends only to inflation, the output gap, and slow-moving trend measures of $r^\star$ will therefore set rates too high precisely when financial repricing is pushing the economy toward contraction. In this sense, the paper develops a risk-premium channel of monetary policy through wealth effects.

The empirical evidence supports this interpretation. Identified increases in disaster probability generate synchronized declines in output, consumption, investment, and employment, accompanied by disinflation and falling safe yields. That pattern is difficult to reconcile with a supply-side interpretation and instead looks like a negative aggregate demand shock under nominal rigidities. The data therefore line up with the mechanism emphasized by the model: when downside risk rises, financial repricing weakens demand and lowers the policy-relevant neutral rate.

The quantitative model shows that this channel is not a second-order refinement. A one percentage-point increase in disaster probability lowers $r_t^\star$ by approximately 16 basis points, much more than an equivalent increase in diffusion volatility. Disaster risk is therefore one of the largest cyclical drivers of the neutral rate in the model. More broadly, the results show

[23] Robustness checks using alternative inflation-expectations measures yield qualitatively similar results; see Online Appendix K.3.

that policy mistakes arise not because central banks fail to stabilize risk premia themselves, but because they fail to accommodate the fall in $r_t^\star$ that those premia imply. A rule that tracks the risk-adjusted neutral rate stabilizes output and inflation far more effectively than one that ignores it, even though risk premia and volatility remain elevated.

The counterfactual exercises sharpen the policy interpretation. Around the onset of the Global Financial Crisis, the calibrated model implies that the Federal Reserve responded to the tail-risk-driven fall in the neutral rate, but only partially. Had policy ignored that fall entirely, the output contraction would have been much more severe. The broader lesson is that monetary policy can insulate the real economy from the demand consequences of a financial repricing shock even if it cannot prevent the repricing itself. This is the sense in which risk premia matter for monetary policy in the framework developed here.

The paper's empirical measure of the risk-adjusted neutral rate provides a practical payoff to this theory. It preserves the low-frequency Ramsey component emphasized by standard estimates such as Laubach–Williams, but it adds the crisis-time variation coming from disaster and diffusion risk. During the GFC and COVID-19, those wedges depressed $r_t^\star$ by 200–300 basis points relative to the Ramsey component. As a result, a policy stance that appeared accommodative relative to conventional benchmarks could in fact be close to neutral or even restrictive once financial repricing and wealth effects were taken into account.

The practical implication is straightforward. Monetary policy should be guided not only by inflation and slack, but also by measures of financial repricing that help track the risk-adjusted neutral rate in real time. Option-implied tail-risk measures are useful in this respect because they capture the component of risk that is most tightly connected to required returns, valuations, and aggregate demand. The case for doing so is strongest in crisis episodes and near the effective lower bound, when the costs of setting rates above the true neutral rate are especially large.

The framework leaves out several potentially important extensions, including financial intermediation, heterogeneity, and the endogeneity of tail risk to policy itself. These could change quantitative magnitudes and add new transmission channels. But they do not alter the core message. When downside risk rises, risk premia widen, wealth falls, aggregate demand weakens, and the neutral rate declines. Optimal monetary policy should respond to that decline. Formalizing this logic helps explain crisis-era policy and offers a more useful way to think about the interaction between risk premia, wealth effects, and monetary policy going forward.

References

Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business-cycle anatomy. *American Economic Review*, 110(10):3030–3070.

Barro, R. J. (2006a). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.

Barro, R. J. (2006b). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.

Barro, R. J. and King, R. G. (1984). Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics*, 99(4):817–839.

Barro, R. J. and Liao, G. Y. (2021). Rare disaster probability and options pricing. *Journal of Financial Economics*, 139(3):750–769.

Basu, S. and Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958.

Basu, S., Candian, G., Chahrour, R., and Valchev, R. (2021). Risky business cycles. Boston College Working Papers in Economics 1029, Boston College Department of Economics. Revised 17 Sep 2024. RePEc handle: RePEc:boc:bocoec:1029.

Billingsley, P. (1995). *Probability and Measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, New York, 3 edition.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, 86(3):1031–1065.

Brunnermeier, M. K., Merkel, S., and Sannikov, Y. (2024). Safe assets. *Journal of Political Economy*, 132(11):3603–3657.

Caballero, R. J. and Simsek, A. (2020). A risk-centric model of demand recessions and speculation. *The Quarterly Journal of Economics*, 135(3):1493–1566.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Del Negro, M., Eggertsson, G., Ferrero, A., and Kiyotaki, N. (2017a). The great escape? a quantitative evaluation of the fed's liquidity facilities. *American Economic Review*, 107(3):824–857.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2017b). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*. Spring 2017.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press, Princeton, NJ, 2 edition.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton, NJ.

Kozeniauskas, N., Orlik, A., and Veldkamp, L. (2018). What are uncertainty shocks? *Journal of Monetary Economics*, 100:1–15.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2019). The tail that keeps the riskless rate low. *NBER Macroeconomics Annual*, 33(1):253–283.

Laubach, T. and Williams, J. C. (2003). Measuring the natural rate of interest. *Review of Economics and Statistics*, 85(4):1063–1070.

Lucas, R. E. (1983). On the welfare cost of inflation and business cycles. *Journal of Monetary Economics*, 12(1):49–69.

Lustig, H. and Verdelhan, A. (2012). Business cycle variation in the risk-return trade-off. *Journal of Monetary Economics*, 59(S):S35–S49.

Martin, I. (2011). Simple Variance Swaps. NBER Working Papers 16884, National Bureau of Economic Research, Inc.

Martin, I. (2013). Consumption-based asset pricing with higher cumulants. *Review of Economic Studies*, 80(2):745–773.

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

Martin, I. W. (2008). Disasters and the welfare cost of uncertainty. *American Economic Review*, 98(2):74–78.

Nakamura, E., Riblier, V., and Steinsson, J. (2025). Beyond the taylor rule. Manuscript.

Rotemberg, J. J. (1982). Monopolistic price adjustment and aggregate output. *The Review of Economic Studies*, 49(4):517–531.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39(1):195–214.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

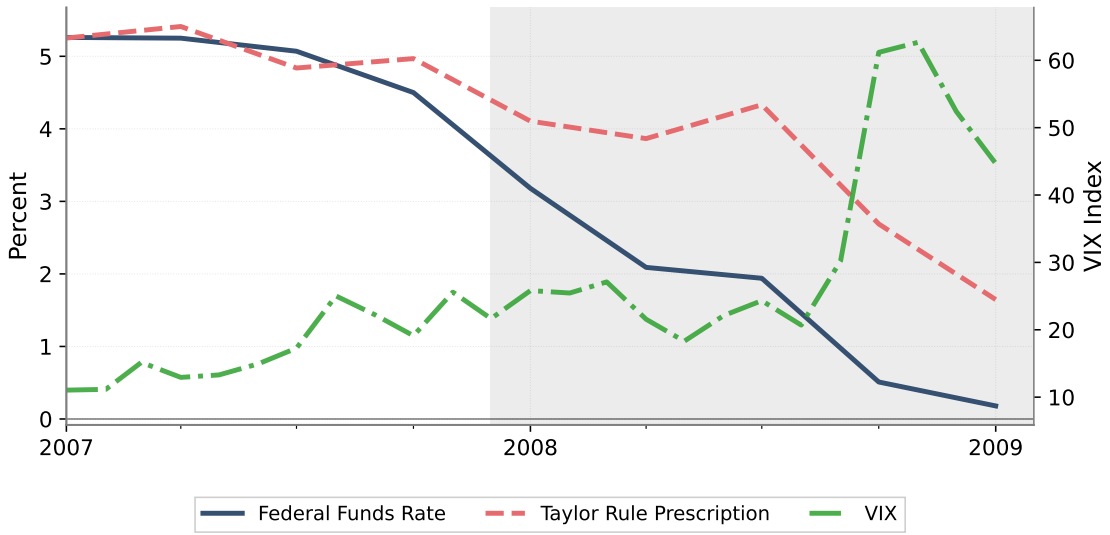
A Alternative Measures of Financial Market Stress

This appendix validates the measure of tail risk used in the main text—options-implied disaster probabilities from Barro and Liao (2021)—by examining two widely-used alternative indicators of financial market stress: the VIX (CBOE Volatility Index) and the Excess Bond Premium (EBP) of Gilchrist and Zakrajšek (2012). Figures 17 and 18 replicate the analysis from Figure 1 in the main text, replacing the disaster probability measure with the VIX and EBP, respectively. All three measures tell a consistent story: financial markets experienced severe stress beginning in mid-2007, well before the standard Taylor rule components (inflation and output gap) suggested the need for aggressive monetary easing.

The VIX as a Measure of Tail Risk

Figure 17 plots the federal funds rate, the Taylor rule prescription, and the VIX index over the 2007–2009 period. The VIX, often referred to as the "fear gauge," measures the 30-day implied volatility of S&P 500 index options and is widely interpreted as a forward-looking indicator of market uncertainty and risk aversion (Bloom, 2009).

Figure 17: Federal Funds Rate vs. Taylor Rule with VIX



Notes: The figure plots the effective federal funds rate, the prescription from the original Taylor (1993) rule using real-time data, and the VIX index (monthly average). The VIX measures the implied volatility of S&P 500 index options and is obtained from the CBOE. Shaded region indicates NBER-dated recession (December 2007–June 2009). Higher VIX values indicate greater expected market volatility and increased risk aversion.

The VIX remained relatively stable, fluctuating between 10 and 15, through mid-2007. However, starting in August 2007—coinciding with the onset of subprime mortgage market turmoil—the VIX began to rise sharply. By early 2008, it had climbed to the 20–30 range, and following the collapse of Lehman Brothers in September 2008, it spiked above 60, reaching a peak near 80 in November 2008. This dramatic increase in implied volatility reflects a fundamental repricing of risk in financial markets and heightened uncertainty about future economic outcomes.

Critically, the VIX began its sustained elevation in mid-2007, precisely when the Federal Reserve began cutting rates despite relatively benign inflation and output gap readings. The parallel movement of the VIX and the deviation between the actual federal funds rate and the Taylor rule prescription provides strong prima facie evidence that the Federal Reserve was responding to financial market stress rather than to traditional macroeconomic fundamentals alone.

The Excess Bond Premium as a Credit Spread Measure

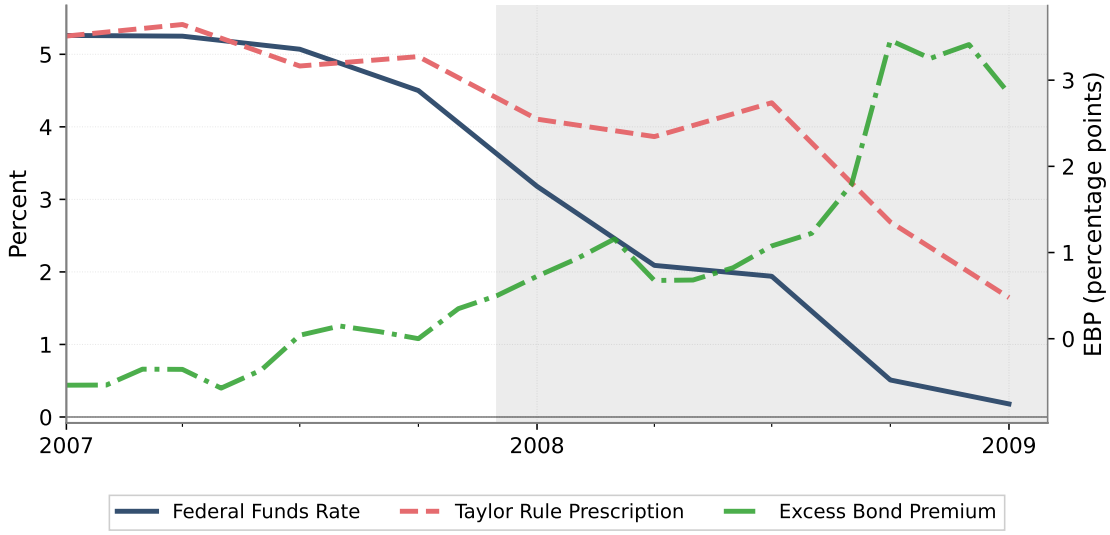
Figure 18 plots the federal funds rate, the Taylor rule prescription, and the Excess Bond Premium (EBP) from Gilchrist and Zakrajšek (2012). The EBP is constructed as the component of corporate bond spreads that cannot be explained by expected default risk, effectively isolating the pure risk premium demanded by investors for holding corporate credit.

Gilchrist and Zakrajšek (2012) demonstrate that the EBP is a powerful predictor of economic activity: elevated credit spreads that cannot be justified by fundamentals (i.e., high EBP) forecast sharp contractions in output, investment, and employment. Moreover, they show that the EBP contains information about economic activity beyond that contained in standard financial indicators such as stock prices or the slope of the yield curve.

The pattern in Figure 18 closely mirrors that in Figures 1 and 17. The EBP remained near zero through mid-2007 but began rising sharply in August 2007, reaching approximately 1 percentage point by early 2008. Following the Lehman Brothers collapse, the EBP spiked to nearly 3 percentage points—its highest level in the post-war period. This dramatic widening of credit spreads, orthogonal to default risk, reflects a surge in the risk premium demanded for holding risky assets, consistent with an increase in tail risk and risk aversion.

As with the VIX, the timing of the EBP's rise coincides precisely with the Federal Reserve's departure from the standard Taylor rule. This suggests that policymakers were responding not only to equity market volatility but also to credit market stress—a distinct

Figure 18: Federal Funds Rate vs. Taylor Rule with Excess Bond Premium

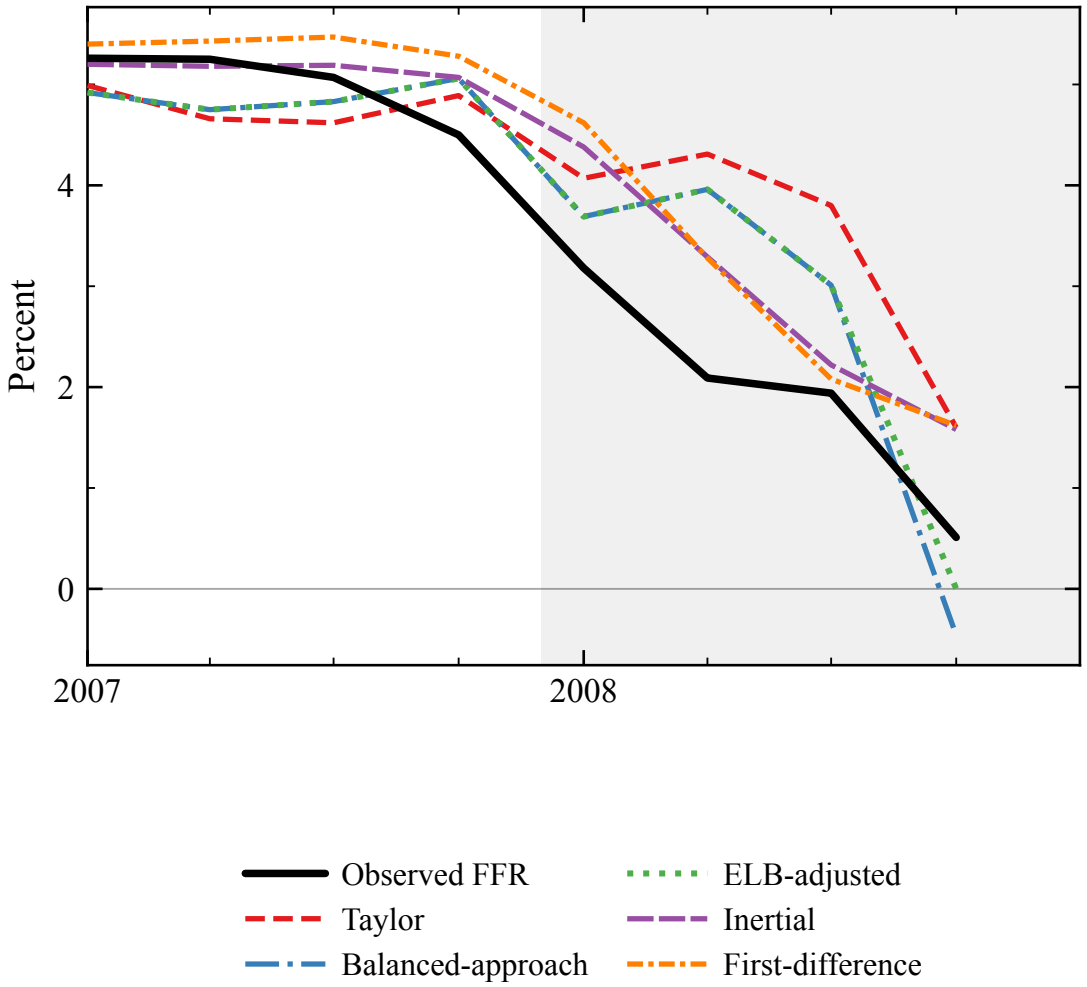


Notes: The figure plots the effective federal funds rate, the prescription from the original Taylor (1993) rule using real-time data, and the Excess Bond Premium (EBP) from Gilchrist and Zakrajšek (2012). The EBP measures the component of corporate bond spreads orthogonal to expected default risk, capturing pure risk premia in credit markets. Shaded region indicates NBER-dated recession (December 2007–June 2009). Higher EBP values indicate greater credit market stress and elevated risk premia.

dimension of financial market tail risk that affects the corporate sector's access to funding and, ultimately, real investment and employment.

A.1 Alternative Taylor Rules

Figure 19: Federal Funds Rate vs. Alternative Taylor Rule Variants



Notes: The figure plots the observed federal funds rate against five standard Taylor rule variants: the original Taylor (1993) rule, the balanced-approach rule, the effective-lower-bound-adjusted rule, the inertial rule, and the first-difference rule. All rules use real-time data. Shaded region indicates NBER-dated recession (December 2007–June 2009). Across every specification, the prescribed rate lies materially above the observed federal funds rate for much of 2007 and into 2008, confirming that the Fed's aggressive easing cannot be attributed to any standard rule variant.