Taming Black Swans: Disaster Risk and the Neutral Rate of Interest *

C. Christopher Hyland†

10th November 2026

Abstract

How far should a central bank cut rates when the perceived risk of a severe contraction rises? The answer depends on the *risk-adjusted* neutral rate, which I propose as the benchmark for the stance of monetary policy: the real rate that keeps inflation at target and output at potential once risk is taken into account. When disaster risk rises, asset prices fall, which lowers spending through a wealth effect, so this rate falls even *before* any disaster strikes. Disaster-risk shocks measured from option prices appear in the data as negative demand shocks that lower output and inflation together. In an estimated nonlinear New Keynesian model, higher-order moments matter for the neutral rate: a rise in disaster risk lowers it 1.77 times as much as an equally large rise in Gaussian volatility. Failing to track this rate is costly. Under a standard Taylor rule, households would be willing to give up 23.6 basis points of consumption permanently for the central bank to track it fully instead. To guard against misspecification, I also estimate this rate semi-structurally: it fell well below zero in late 2008, so policy was contractionary even though traditionally used estimates suggest it was accommodative.

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

*I am grateful to Joseph Abadi, Yacine Aït-Sahalia, Andrey Alexandrov, Paul Beaudry, Federico Belo, Harjoat Bhamra, Francesco Bianchi, Javier Bianchi, Adrien Bilal, Luigi Bocola, Lukas Boehnert, Markus Brunnermeier, Vasco Carvalho, Riccardo Cioffi, Louphou Coulibaly, Wei Cui, Sitong Ding, Keshav Dogra, Stéphane Dupraz, Martin Ellison, Sergio de Ferra, Alessandro Ferrari, Andrea Ferrero, Fiorella De Fiore, Jordi Galí, Marco Garofalo, Matthieu Gomez, Goutham Gopalakrishna, Lars Hansen, Kathryn Holston, Miyu Haraguchi, Tobias Kawalec, Rohan Kekre, Paymon Khorrami, Paul Klemperer, Michael Kiley, Alex Kohlhas, Lukas Kremens, Dirk Krueger, Seung Joo Lee, Daniel J Lewis, Simon Lloyd, Wenhao Li, J. Zachary Mazlish, Michael McMahon, Fernando Mendo, Enrique Mendoza, Sebastian Merkel, Thomas Mertens, Silvia Miranda-Agrippino, Ben Moll, Dmitry Mukhin, Stefan Nagel, Bruno Pellegrino, Olaf Posch, Ricardo Reis, Michael Reiter, Giovanni Rosso, Francisco Ruge-Murcia, Andrew Scott, Dejanir Silva, Sanjay R. Singh, Viktor Todorov, Dimitrios Tsomocos, Tim Willems, John C. Williams, 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.dphil@said.oxford.edu.

1 Introduction

> *"Navigating by the stars can sound straightforward. Guiding policy by the stars in practice, however, has been quite challenging of late because our best assessments of the location of the stars have been changing significantly."*
>
> — Jerome H. Powell (2018)

The conduct of monetary policy is often described as "navigating by the stars" (Powell, 2018). In this metaphor, the stars are the "normal" or "desired" values around which output growth, unemployment and interest rates fluctuate. One of these stars is the neutral rate of interest $r_t^\star$: the real interest rate that keeps output at potential and inflation at target, and so the rate that monetary policy aims for (Galí, 2015). The gap between the real policy rate and $r_t^\star$ measures the stance of policy. Policy is contractionary when the real rate is above $r_t^\star$, because households then postpone spending, which lowers output and inflation. This idea goes back to Wicksell (1898) and has long guided both central-bank practice and the field of monetary economics, culminating in the neo-Wicksellian framework outlined in Woodford (2003). However, as $r_t^\star$ cannot be observed, considerable effort has gone into estimating it, the most prominent estimate being that of Laubach and Williams (2003). Figure 1(a) plots the Laubach–Williams estimate, which drifts slowly downwards over the first half of the sample and then rises gently.

In practice, however, navigating by the stars is hard, because the stars move. The real interest rate needed to keep output at potential and inflation at target varies over the business cycle. So, too, does the compensation investors require for bearing risk. A large literature in macro-finance finds that risk premia are countercyclical and that their variation accounts for most of the volatility of asset prices (Cochrane, 2011). Through asset prices, risk premia in turn move household spending and employment (Chodorow-Reich et al., 2021), so the interest rate needed to stabilise the economy should move with them. One influential explanation for these movements in risk premia is time-varying disaster risk, the risk of a rare but severe contraction (Barro, 2006). Figure 1(a) also plots the probability of a large fall in equity prices that option markets assign, which I use to measure disaster risk. It moves sharply with the business cycle, spiking in 1998, in 2008 and in 2020 and falling back within a year or two. There is therefore a tension: canonical estimates of the neutral rate are largely acyclical, whereas risk premia and disaster risk vary sharply over the business cycle.

This paper asks how far disaster risk, the perceived probability of a disaster, moves the neutral rate $r_t^\star$, what this implies for monetary policy and how large the welfare losses are when policy fails to track it. To this end, I propose the *risk-adjusted neutral rate* as the benchmark for the stance of monetary policy: the real rate that keeps inflation at target

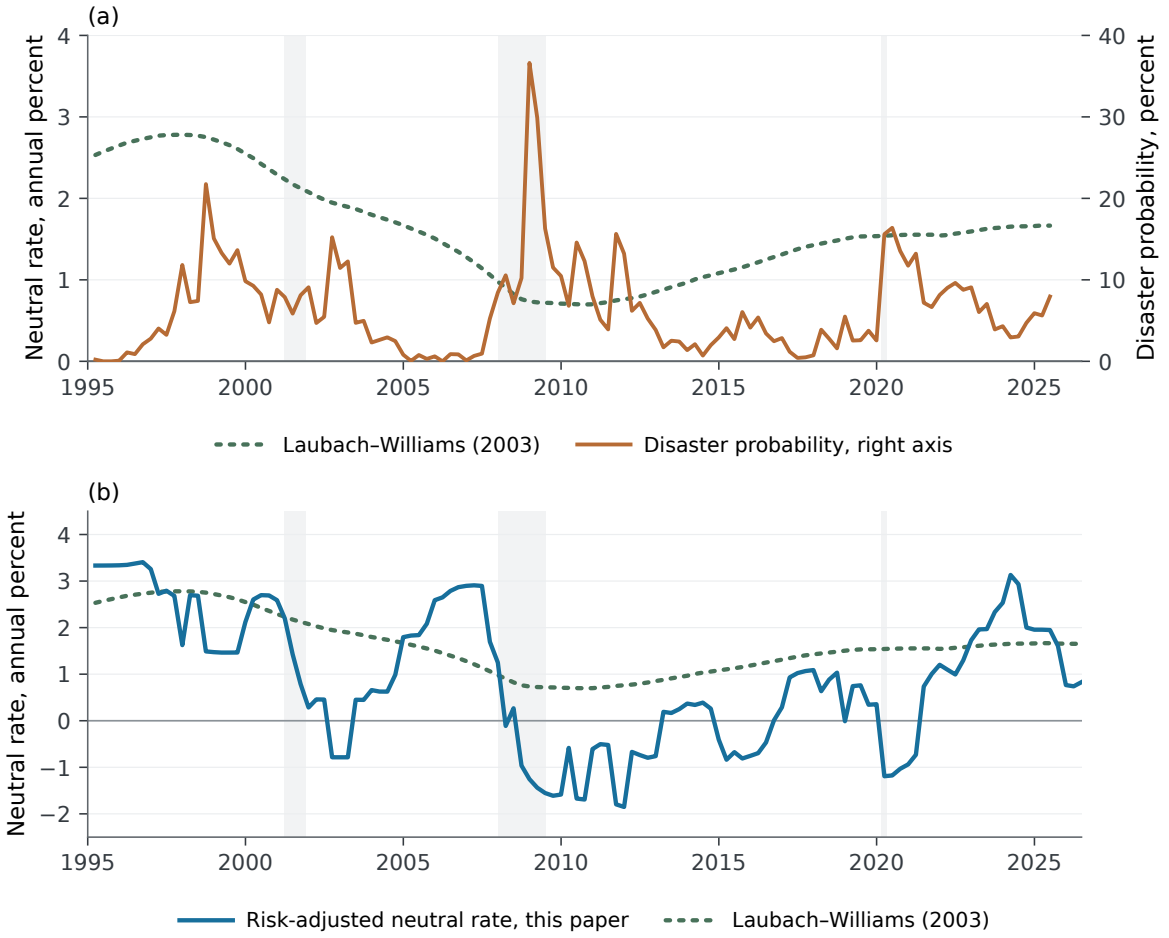


Figure 1: Disaster risk and the neutral rate

Notes: Panel (a): the current-vintage two-sided estimate of the neutral rate of Laubach and Williams (2003), in annual percent (left axis), and the quarterly mean of the annualised probability of a large equity-market decline implied by S&P 500 put options (Barro and Liao, 2021), in percent (right axis), 1995Q1–2025Q2. Panel (b): the smoothed risk-adjusted neutral rate from the estimated model and the Laubach–Williams estimate, in annual percent, 1995Q1–2026Q2. Recessions are shaded.

and output at potential *once risk is taken into account*. It extends the neutral rate of standard New Keynesian models, the real rate that would prevail if prices were fully flexible, to account for risk, which drops out of first-order approximations of those models. When disaster risk rises, asset prices fall, which lowers current spending through a wealth effect. The risk-adjusted neutral rate therefore falls, even *before* any disaster strikes. Because disasters are tail events, the rate that monetary policy aims for then depends not only on the variance of output growth but also on its higher-order moments, including skewness and kurtosis. These moments drop out of first-order approximations, so studying this rate requires a framework that allows for them and for other nonlinearities. This paper

develops such a framework, works out its implications for the risk-adjusted neutral rate and estimates that rate for the United States with a quantitative model fitted to option prices and macroeconomic data. Figure 1(b) plots the estimate: it falls sharply when disaster risk rises, to below zero in 2002–2003, 2008–2012, 2015–2016 and 2020, while the Laubach–Williams estimate barely moves. The main result is that the risk terms left out of first-order approximations are quantitatively large. Disaster risk moves the neutral rate substantially, and faster than the standard estimates do. A central bank that judges its stance against slow-moving estimates can therefore inadvertently tighten policy even as it cuts rates, if the neutral rate falls by more than the cut.

Contributions. The paper brings together two literatures that have developed largely apart: the rare-disaster view of asset prices (Barro, 2006) and the neo-Wicksellian theory of monetary policy (Woodford, 2003). I first develop a two-period New Keynesian model in which output growth can follow any distribution. Using the cumulant-generating function of Martin (2013), I characterise analytically how each moment of the distribution of risk, and not only the variance, affects the economy. I show that changes in the higher-order moments can move the neutral rate substantially. If the safe rate does not adjust, a rise in tail risk lowers output even when the first moment, expected future output, is held fixed. I thus show that movements in tail risk can provide a microfoundation for the "risk-premium" shock in estimated DSGE models such as Smets and Wouters (2007), which likewise lowers demand at a given policy rate. I also show that a central bank that responds only to the second moment of output growth, ignoring its higher-order moments, leaves large welfare gains unrealised, because its policy rate stays above the neutral rate. In fact, the cost of this policy misspecification is about as large as the cost of business-cycle risk itself. I then show that, for a wide range of distributions of output growth, a central bank that tracks the moments up to the fourth, kurtosis, comes close to replicating optimal policy.

I then test this mechanism in the data, using the disaster probability of Figure 1(a), which is inferred from out-of-the-money put options on the S&P 500 (Barro and Liao, 2021). I find that, as the mechanism suggests, disaster-risk shocks appear in the data as negative demand shocks: they lower consumption, investment, output and inflation together. This evidence supports interpreting the risk-premium shock of estimated DSGE models as, at least in part, a shock to perceived disaster risk. Disaster-risk shocks are also an important driver of the business cycle, explaining about 15% of the cyclical variation in output and investment. The joint fall in activity and inflation holds across different identification schemes and on a sample that includes the pandemic.

To measure the size of these effects, I build a quantitative New Keynesian model in continuous time and estimate it on US data. The model lets me quantify how far disaster risk moves the neutral rate, output and inflation over the business cycle, and run policy counterfactuals. The central bank follows a Taylor rule that can track the neutral rate fully, partly or not at all, so I can compare these rules by their effects on output, inflation and welfare.

Quantitative results. First, higher-order moments matter for the neutral rate. I compare two increases in the risk of productivity growth that leave its mean unchanged and raise its variance by the same amount. One makes Gaussian shocks more volatile and the other makes disasters more frequent. Because only disasters also make growth negatively skewed and fat-tailed, the comparison isolates the role of the higher-order moments. The rise in disaster risk lowers the neutral rate by 48.8 basis points, 1.77 times as much as the rise in Gaussian volatility (27.5). A rare, large loss raises households' precautionary saving by more than frequent small fluctuations with the same variance, since marginal utility rises more steeply the further consumption falls. Under log utility the difference is small. It widens considerably with Epstein–Zin preferences (Epstein and Zin, 1989; Duffie and Epstein, 1992), under which households care not only about the immediate loss in a disaster but also about the weaker future it leaves behind.

Second, the estimated model implies that the Federal Reserve tracks the risk-adjusted neutral rate only in part. Its policy rule places a weight of 0.64 on that rate, compared with zero under a standard Taylor rule and one under full tracking. Fitting the model to the onset of the Global Financial Crisis, I find that in late 2007 the Federal Reserve cut rates by more than the original Taylor (1993) rule prescribed, but not by enough to keep pace with the falling neutral rate. Had it followed a standard Taylor rule, which ignores movements in the neutral rate, the same rise in disaster risk would have lowered output by early 2008 about 46 percent more than under its estimated rule, which tracks them in part.

Third, failing to track the risk-adjusted neutral rate is costly. I measure the welfare loss of a policy rule that ignores disaster risk by the share of consumption that households would be willing to give up permanently for the central bank to track the risk-adjusted neutral rate fully instead. Here, full tracking of the risk-adjusted neutral rate achieves the flexible-price allocation, in which output stays at potential and inflation at target despite the same disaster risk. The loss is 0.236% of consumption under a standard Taylor rule. With a Phillips curve 3 times flatter, closer to recent estimates (Hazell et al., 2022; Del Negro et al., 2020), it rises to 0.567%.

Finally, to address concerns about model misspecification, I also estimate the risk-

adjusted neutral rate with a semi-structural model that does not impose the model's structure. It builds on Laubach and Williams (2003) but lets the option-based measure of disaster risk lower the neutral rate. The two approaches agree on how strongly disaster risk moves the neutral rate: a one-percentage-point rise in the disaster probability lowers it by 24.4 basis points in the semi-structural model and by 27.6 in the estimated model. The two estimates move together, with a correlation of 0.73 since 1995, and both fall in every recession. Because option prices respond immediately to news about risk, both drop sharply when risk spikes: in 2008Q4 the semi-structural estimate reaches −5.76% and the model's −1.25%, while the Laubach–Williams estimate stays positive. Published estimates of the short-run neutral rate also fall below zero in the financial crisis, though they differ on how far. Measured against the model's estimate, policy was contractionary at the onset of the crisis, before the large rate cuts of 2008, and again in 2020Q2, when long-run estimates suggest accommodation.

Related literature. This paper builds on several strands of work. First, it takes from the rare-disaster literature the insight that exposure to severe contractions helps explain high equity premia and low safe rates (Barro, 2006; Gabaix, 2012; Gourio, 2012; Wachter, 2013). A related literature shows that the higher-order moments of consumption growth, beyond its variance, can account for a large share of the compensation for risk (Martin, 2013; Backus et al., 2014). Isoré and Szczerbowicz (2017) add time-varying disaster risk to a New Keynesian model and show that, with sticky prices and a low elasticity of intertemporal substitution, a rise in disaster risk lowers consumption, which gives a rationale for discount-factor shocks. I take these insights to monetary policy. In the textbook New Keynesian model, the neutral rate moves with expected potential growth (Galí, 2015); I show that it also moves with the risk to potential growth. I then ask how far the policy rate must move to keep output at potential when that risk rises, and how the answer differs between the volatility of Gaussian shocks and disaster risk.

Second, the demand mechanism relates to work linking asset values, wealth and economic activity (Del Negro et al., 2017a; Chodorow-Reich et al., 2021; Brunnermeier et al., 2024). My two-period model builds on that of Caballero and Simsek (2020), in which higher risk lowers asset prices and, through a wealth effect, aggregate demand. They measure risk by volatility. I instead let potential output growth follow any distribution, so its higher-order moments can describe the bad outcomes that financial markets price. I can then study how these moments move the neutral rate, and what happens to activity and inflation when policy fails to track it. Another literature emphasises the role that financial frictions play, whereby falling asset values tighten credit and deepen downturns

(Bernanke et al., 1999; Gilchrist and Zakrajšek, 2012), which can also move the neutral rate. I abstract from them to focus on how aggregate risk moves the neutral rate. In contrast, Beaudry et al. (2025) offer a different view of how interest rates affect demand. In their life-cycle economy, retirement-saving motives can offset intertemporal-substitution and asset-valuation effects, so activity responds less to persistent changes in interest rates, and a policy rate that fails to track the neutral rate would do less damage. I abstract from such life-cycle forces in this paper.

Third, the paper relates to the literature on uncertainty, tail risk and business cycles (Bloom, 2009; Bloom et al., 2018; Kozeniauskas et al., 2018; Orlik and Veldkamp, 2024). Closest in spirit are Kozlowski et al. (2019, 2020), who argue that persistent revisions to beliefs about tail events after the Great Recession kept the riskless rate low and output depressed, and Pflueger et al. (2020), who find that safe rates fall and investment declines when financial markets perceive more risk. I take such movements in perceived tail risk as given and ask what they imply for the stance of monetary policy. Recent work shows that standard Taylor rules cannot rule out equilibria in which volatility and risk premia are self-fulfilling, and that a rule that tracks a risk-adjusted neutral rate can restore a unique equilibrium (Khorrami and Mendo, 2025; Lee and Dordal i Carreras, 2025); Dordal i Carreras and Lee (2026) study forward guidance at the lower bound when volatility is endogenous. Monetary policy can also move risk premia, by redistributing wealth towards households more willing to bear risk (Kekre and Lenel, 2022). In a heterogeneous-agent New Keynesian model with convenience yields and risk premia, Caramp et al. (2026) find that monetary contractions raise both Treasury convenience yields and the market-implied valuation of downside risk, and that risk and liquidity amplify the transmission of monetary policy. Shifts in the policy regime also help explain persistent swings in asset valuations (Bianchi et al., 2022). Disaster risk is exogenous in my model, so the analysis isolates how it moves the neutral rate, which I also measure in the data.

Fourth, the empirical analysis builds on the extraction of risk-neutral distributions and tail risk from option prices (Breeden and Litzenberger, 1978; Andersen et al., 2015; Martin, 2017; Barro and Liao, 2021) and on the identification of risk shocks in VARs (Uhlig, 2003; Basu and Bundick, 2017; Basu et al., 2021). Rather than volatility or an estimated equity premium, I measure risk with an option-implied disaster probability, show in the two-period model how option prices map into the physical probability of a disaster, and find that shocks to this probability appear in the data as negative demand shocks.

Fifth, the welfare analysis connects to work on the cost of aggregate fluctuations. Normally distributed consumption risk implies small costs under standard preferences (Lucas, 2003), but the costs become much larger with Epstein–Zin preferences or rare

disasters (Tallarini, 2000; Martin, 2008; Barro, 2009). I separate this fundamental cost of risk, which monetary policy cannot remove, from the additional cost of failing to track the neutral rate, which it can.

Finally, the historical analysis complements work on the slow-moving determinants of the neutral rate, such as trend growth and safety premia (Laubach and Williams, 2003; Del Negro et al., 2017b; Brunnermeier et al., 2024), and on the demand for safe assets (Krishnamurthy and Vissing-Jorgensen, 2012; Caballero et al., 2017). The convenience yield on Treasury securities is specific to one asset, whereas disaster risk is a property of the economy as a whole. Relative to these slow-moving determinants, the risk-adjusted neutral rate adds a disaster-risk component that moves with the business cycle, which brings it closer to the short-run neutral rate studied by Barsky et al. (2014). Cúrdia et al. (2015) estimate a similar short-run rate and find that a policy rule that tracks it fits US data better than otherwise identical Taylor rules. My estimated rule nests a standard Taylor rule and full tracking, so I can measure how far the Federal Reserve tracks a neutral rate that moves with disaster risk. I estimate the risk-adjusted neutral rate both from the model and semi-structurally, and compare the two estimates with published estimates of the short-run and long-run neutral rate.

Outline. Section 2 presents the two-period model, and Section 3 the empirical evidence. Section 4 develops the quantitative New Keynesian model. Section 5 estimates it, measures the risk-adjusted neutral rate and the stance of monetary policy, and asks how far disaster risk moves the neutral rate, how monetary rules respond to news of higher disaster risk, including at the onset of the financial crisis, what failing to track the neutral rate costs and how a semi-structural estimate compares. Section 6 concludes.

2 A Two-Period Model of the Neutral Rate

I develop a two-period New Keynesian model in which the growth of potential output, driven by productivity, can follow *any* distribution, so that risk is described by all the moments of that distribution and not only by its variance. The neutral rate nevertheless has a closed form. I use the model to define the risk-adjusted neutral rate, to ask how these higher-order moments affect it, and to study what happens to output and welfare when monetary policy fails to track it. Appendix E.1 sets out the main steps of the derivations, and Online Appendices A–C give the full proofs.

2.1 A two-period economy

Production and prices. The economy has two dates, $t \in \{0, 1\}$, and one consumption good produced with capital, which is fixed and normalised to one. Potential output equals productivity, z_t, with $z_0 = 1$. At date zero, firms face fully rigid prices, so current output is determined by demand: firms meet whatever households wish to consume, up to potential output,

$$y_0 = c_0 \leq z_0 = 1. \tag{1}$$

Date one represents the long run, when prices are flexible and output reaches potential, $y_1 = z_1$. Date-one productivity, z_1, is random and is the only source of aggregate uncertainty in the economy. Since $z_0 = 1$, its logarithm, $G \equiv \log z_1$, is the growth rate of potential output, or potential growth. As in the New Keynesian literature, potential output is the level of output under flexible prices (Woodford, 2003; Galí, 2015); here it moves only with productivity, and a disaster lowers it. For short, I refer to G as output growth.

Monetary policy. With prices fully rigid, the central bank controls the real interest rate through its policy rate. I therefore treat the one-period real safe log rate, r^f, as the policy instrument, and I abstract from an effective lower bound.

Definition 2.1 (Neutral rate)**.** *The neutral rate, $r^\star$, is the real safe rate at which the equilibrium level of current output equals potential, $y_0 = z_0$.*

This rate is the neo-Wicksellian benchmark in this economy (Woodford, 2003): the rate at which monetary policy neither depresses nor stimulates current demand.[1] Equivalently, the neutral rate is the real rate that would prevail if prices were flexible at date zero, so it does not depend on monetary policy. The central bank chooses the policy rate, and the gap between that rate and the neutral rate measures the stance of policy.

[1]The terminology follows Obstfeld (2025), who calls the long-run equilibrium real rate the natural rate and the rate appropriate for monetary policy at a given time the neutral rate. Estimates such as that of Laubach and Williams (2003) target the former; I study the latter. The flexible-price real rate of New Keynesian models, which Woodford (2003) calls the natural rate of interest, is what Obstfeld would call the neutral rate.

Output growth risk. Output growth, G, can follow any distribution, which I summarise by its cumulant-generating function (CGF),[2]

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

where $\kappa_n = c_G^{(n)}(0)$ is the nth cumulant. The first cumulant is mean output growth and the second is its variance, while κ_3 and κ_4 measure (unstandardised) skewness and excess kurtosis. The CGF works as a bookkeeping device for the whole distribution. I write $g \equiv c_G(1) = \log \mathbb{E}[z_1]$ for the log of expected date-one output.

Assets. Households trade two assets: a risk-free bond, in zero net supply, with log return r^f; and a claim to aggregate output in the next period, z_1, which is in unit supply and costs $Q > 0$. I refer to this claim as equity: like the aggregate stock market, it is a risky claim to the economy's future output. Its log return is

$$r^m = \log(z_1/Q) = G - \log Q.$$

Because the price of equity is known at date zero, the equity return inherits the CGF of output growth:

$$c_{r^m}(\theta) \equiv \log \mathbb{E}[e^{\theta r^m}] = c_G(\theta) - \theta \log Q.$$

Because the price of equity shifts only the *mean* return, every cumulant of the equity return beyond the first equals that of output growth G, so skewness and tail risk in output growth pass one for one into the return on equity. I can therefore characterise risk premia in this economy for any distribution of output growth.

Households. The representative household is endowed with current output and with all the equity, and takes income and prices as given. Its preferences are Epstein–Zin (Epstein and Zin, 1989), which separate risk aversion from the elasticity of intertemporal substitution; I set the elasticity to one and relative risk aversion to $\gamma > 1$. At date zero, the household chooses consumption $c_0 > 0$, saving $a_0 > 0$ and the share $\alpha \in [0, 1]$ of saving

[2]The CGF is defined for all θ at which $\mathbb{E}[e^{\theta G}]$ is finite, and satisfies $c_G(0) = 0$. The exact expressions below require $\mathbb{E}[e^{\theta G}]$ to be finite at $\theta = 1$, $-\gamma$ and $1 - \gamma$; the cumulant expansion also requires the Taylor series of the CGF to converge. Knowing every cumulant is equivalent to knowing every moment. Cumulants are convenient for two reasons: the cumulants of a sum of independent random variables are the sums of their individual cumulants, and a deterministic shift in G changes only the first. In asset pricing, Martin (2013) uses the CGF of consumption growth to express safe rates, risk premia and valuation ratios in terms of all the cumulants, and Martin (2008) applies it to the welfare cost of disaster risk. I use it to characterise the neutral rate.

held in equity, with the remainder held in the bond. It solves

$$\begin{aligned} \max_{c_0, a_0, \alpha} \quad & \log c_0 + e^{-\rho} \log U_1, \qquad U_1 = \left(\mathbb{E}[c_1^{1-\gamma}]\right)^{1/(1-\gamma)}, \\ \text{subject to} \quad & c_0 + a_0 = y_0 + Q, \\ & c_1 = a_0 \left[\alpha e^{r^m} + (1-\alpha) e^{r^f}\right], \quad \rho > 0. \end{aligned} \tag{3}$$

Here U_1 is the certainty equivalent of future consumption. The household divides its wealth, $y_0 + Q$, between consumption and saving, and saving earns the portfolio return.

Equilibrium. With a unit elasticity of intertemporal substitution, the household consumes a constant share of its wealth:

$$c_0 = \frac{y_0 + Q}{1 + e^{-\rho}}. \tag{4}$$

Two market-clearing conditions close the model. In the goods market, output is demand-determined at date zero, $y_0 = c_0$, while future consumption equals potential output, $c_1 = z_1$, in the flexible-price long run. In the asset market, the household holds all the equity and no bonds, so $\alpha = 1$ and $a_0 = Q$. Substituting (4) into goods-market clearing gives

$$y_0 = e^{\rho} Q. \tag{5}$$

Equation (5) links current output to the price of equity through a wealth effect. When equity falls in value, households are poorer and consume less. Because output is demand-determined at date zero, firms then produce less. Caballero and Simsek (2020) derive a similar output–asset-price relation, in which risk enters through the volatility of returns. Here, the price of equity depends on the whole distribution of output growth, so its higher-order moments, and in particular the risk of a disaster, also move current demand.

2.2 Risk compensation and the neutral rate

To separate the *shape* of output growth risk from its mean, I write output growth as its mean plus a "pure risk shock",

$$G = \kappa_1 + \tilde{\varepsilon}, \qquad \mathbb{E}[\tilde{\varepsilon}] = 0.$$

The pure risk shock isolates the cumulants of G of order two and higher, which describe only the shape of risk, such as its variance, skewness and tail thickness. I measure the compensation households require for bearing this risk by the equity premium, $\mathcal{W}$: the log expected gross return on equity less the safe log return.

Theorem 1 (The equity premium)**.** *The equity premium depends only on the pure risk shock, and it is positive whenever output growth is risky. At a given g and safe rate, a higher premium lowers the price of equity and therefore, by (5), current output:*

$$\begin{aligned} \mathcal{W} &\equiv \log \mathbb{E}[e^{r^m}] - r^f = c_G(1) + c_G(-\gamma) - c_G(1-\gamma), \\ \log \mathcal{Q} &= g - r^f - \mathcal{W}. \end{aligned} \tag{6}$$

Expanding the CGF shows how each feature of the distribution is priced:

$$\mathcal{W} = \sum_{n=2}^{\infty} \Upsilon_n \kappa_n, \qquad \Upsilon_n \equiv \frac{1 + (-\gamma)^n - (1-\gamma)^n}{n!}. \tag{7}$$

The coefficients depend only on risk aversion. The first three are $\Upsilon_2 = \gamma$, $\Upsilon_3 = -\gamma(\gamma-1)/2$ and $\Upsilon_4 = \gamma(2\gamma^2 - 3\gamma + 2)/12$. Variance is therefore priced at γ, and since $\Upsilon_3 < 0$ and $\Upsilon_4 > 0$, negative skewness and positive excess kurtosis add to the premium. Here, all cumulants beyond the fourth enter as well, so the risk premium prices the whole shape of the distribution of output growth.

Because the premium lowers current output at a given safe rate, it also determines the neutral rate: the safe rate that keeps output at potential (Definition 2.1). Following Dordal i Carreras and Lee (2026), I call this rate the risk-adjusted neutral rate.

Proposition 2.1 (The risk-adjusted neutral rate)**.** *Current output equals potential when the safe rate equals the risk-adjusted neutral rate,*

$$r^\star = g + \rho - \mathcal{W}, \tag{8}$$

which falls one for one with the premium at given g and ρ. At any higher safe rate, output falls short of potential by the rate gap: $\log y_0 = r^\star - r^f$ *for* $r^f \geq r^\star$.

The risk-adjusted neutral rate has two parts. The first, $g + \rho$, is the deterministic Ramsey rate: the neutral rate of the textbook New Keynesian model, which moves with expected potential growth (Galí, 2015). It is higher when expected growth is higher or households are less patient, because in either case households want to bring consumption forward, and a higher rate is needed to keep spending at potential. The second is the risk adjustment, $-\mathcal{W}$, which the textbook model leaves out: risk to potential growth raises the precautionary demand for the safe asset, so a lower safe rate is needed to keep current spending at potential. Therefore, a central bank that moves its policy rate one for one with the neutral rate keeps output at potential. As an implication, any rate cut smaller than the fall in the neutral rate leaves output below potential even though policy has *eased*.

Analytical examples. I illustrate Proposition 2.1 with two familiar distributions. The first is Gaussian: when the pure risk shock is $\tilde{\varepsilon} \sim N(0, \sigma^2)$, all cumulants beyond the second vanish, so $\mathcal{W} = \gamma\sigma^2$ and

$$r^\star = g + \rho - \gamma\sigma^2.$$

The risk adjustment is then the familiar precautionary-saving term, which depends on variance alone.

The second is the rare-disaster distribution of Barro (2006) with a single disaster size, a Bernoulli–Gaussian mixture. Output growth has the same Gaussian shock, but with probability p a disaster also scales future output by e^u, where $u < 0$ is the log disaster jump. Output in normal times is raised just enough to keep expected future output at e^g, so a higher disaster probability changes risk without changing expected output.[3] Figure 2 plots the distribution for $u = -0.35$ and three disaster probabilities: a small disaster mode sits well below the main mode and grows with p, while the main mode shifts up to keep expected future output fixed.

The neutral rate is then

$$r^\star = g + \rho - \gamma\sigma^2 - \log \frac{(1 - p + pe^u)(1 - p + pe^{-\gamma u})}{1 - p + pe^{(1-\gamma)u}}. \tag{9}$$

The logarithmic term is the disaster's contribution to the risk adjustment. Its largest factor, $1 - p + pe^{-\gamma u}$, reflects how much households value consumption after a disaster, when marginal utility is $e^{-\gamma u}$ times its level without one. As $p \to 0$ or $u \to 0$, each factor inside the logarithm tends to one, so the term vanishes and the neutral rate returns to its Gaussian value.

2.3 Disaster risk, demand and welfare

I now use the disaster economy to quantify how disaster risk changes the neutral rate, current demand and welfare, using the illustrative calibration in Table 1. The value of g corresponds to expected output growth of about 2.5 percent. I vary the disaster probability p between zero and 5 percent for three disaster sizes, which reduce output by about 18, 30 and 39 percent. Throughout, preferences and the volatility of the Gaussian shock are unchanged, and output in normal times adjusts so that expected future output stays the same. Holding these fixed isolates the role of the moments beyond the second, such as skewness and kurtosis.

[3] Output growth is $G = g - \sigma^2/2 - \log(1 - p + pe^u) + Z + Bu$, where $Z \sim N(0, \sigma^2)$ and $B \sim \text{Bernoulli}(p)$ are independent.

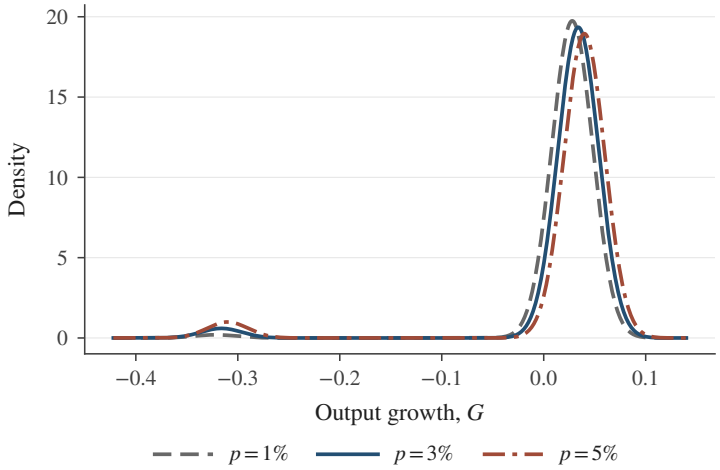


Figure 2: Distribution of output growth with a rare disaster

Notes: The disaster jump is $u = -0.35$, with $g = 0.0252$ and $\sigma = 0.02$ as in Table 1. Each distribution has the same expected long-run output, e^g.

Table 1: Two-period calibration

Parameter	Description	Value
γ	Relative risk aversion	4
ρ	Time preference	0.005
g	Log expected date-one output	0.0252
σ	Gaussian volatility of log output growth	0.02
u	Log disaster jump	$-0.20, -0.35, -0.50$
p	Disaster probability	$[0, 0.05]$

I compare two policies. Under full tracking, the safe rate equals the neutral rate and output stays at potential. Under a fixed rate, the central bank keeps the safe rate at the neutral rate of the Gaussian benchmark, 2.86%, and so ignores disaster risk. The fixed rate mirrors a standard Taylor rule, whose intercept is a constant neutral rate that does not respond to disaster risk directly. Comparing the two policies shows what is lost when the central bank ignores disaster risk.

Figure 3 plots the results against the disaster probability for the three disaster sizes, under the fixed rate. As disaster risk rises, the equity premium rises: equity pays least in a disaster, when consumption is scarce and each unit is most valuable, so households require more compensation to hold it. At $p = 5\%$ and $u = -0.35$, the premium rises from 0.16 to 4.01 percentage points. The risk-adjusted neutral rate falls one for one, from 2.86%

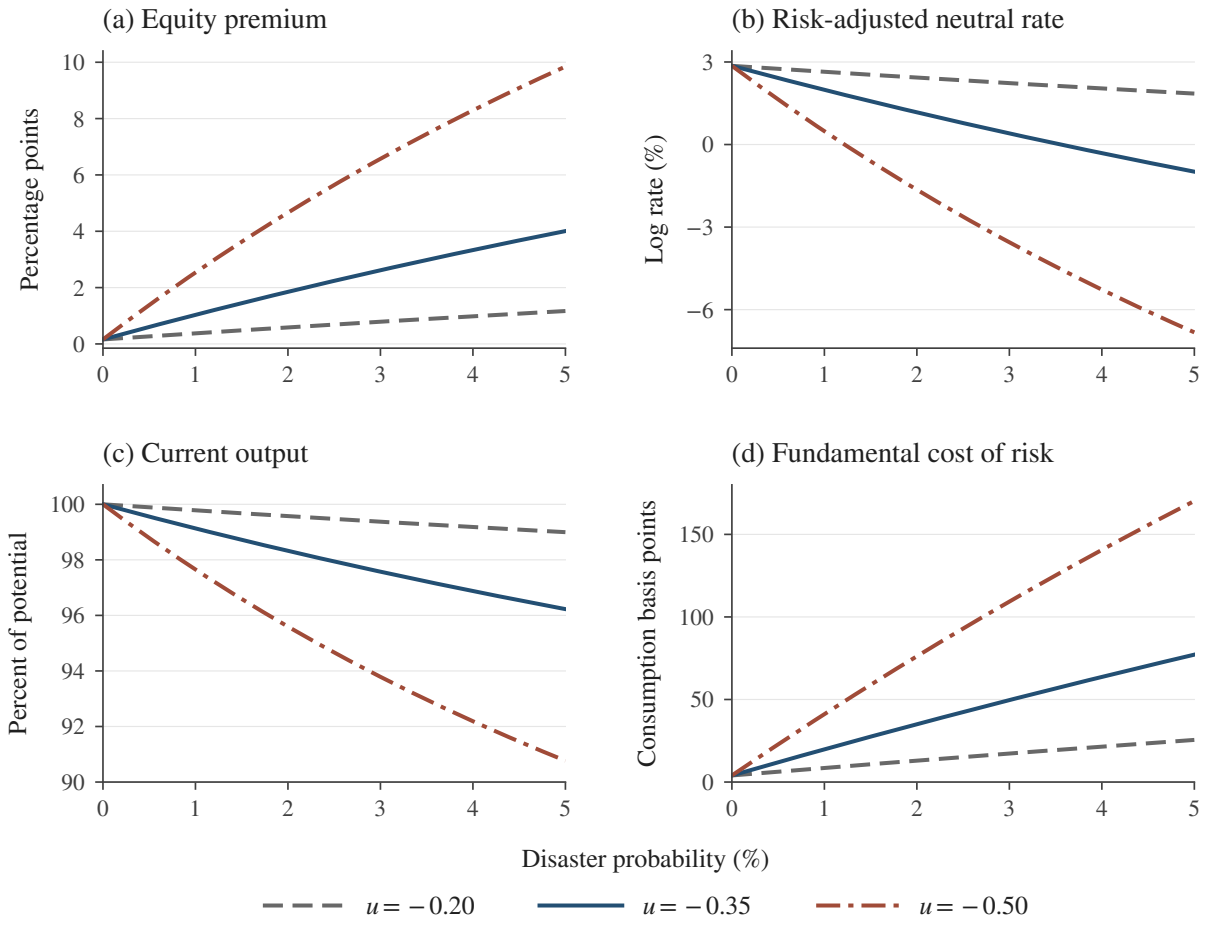


Figure 3: Disaster risk, demand and welfare

Notes: Expected long-run output is fixed. Current output is computed under a fixed rate, with the real safe rate held at the neutral rate without disasters, 2.86%. Panel (d) reports the fundamental cost of risk, CE^{risk}.

to −0.99%, and to −6.83% when $u = -0.50$. Current output falls with it. Facing a riskier future, households want to save more and so spend less today; because prices are rigid, output is determined by demand, and with the safe rate held above the neutral rate it falls to 96.23% of potential, even though expected future output is unchanged.

To check the role of the moments beyond the second, I also construct a Gaussian economy with the same variance and expected long-run output as the disaster economy. At $p = 5\%$ and $u = -0.35$, its premium is only 2.49 percentage points, against 4.01 in the disaster economy. The difference, 1.52 percentage points, is the net contribution of the moments beyond the second. It lowers the neutral rate one for one and, under a fixed rate, deepens the contraction by the same amount. Higher-order moments therefore play a substantial role in movements of the neutral rate.

Disasters also make risk more costly for households. Panel (d) measures this fundamental cost of risk, CE^{risk}, as in Lucas (2003): the share of consumption at both dates that households would give up to replace risky future output with its expected value. Without disasters, the cost is only about 4 basis points of consumption, because Gaussian volatility is small, as in Lucas's calculation. But the equity premium measures how much households dislike risk, and a model that generates a sizeable premium also implies a large cost of risk (Tallarini, 2000). With disasters, the premium rises to 4.01 percentage points and the cost to 77 basis points at $p = 5\%$ and $u = -0.35$, so the gains from eliminating this risk are large. The finding mirrors the large welfare costs of disasters documented by Martin (2008) and Barro (2009), here in a New Keynesian economy.

2.4 A truncated-moment central bank

Suppose the central bank knows expected future output but tracks only the first N cumulants of the pure risk shock. Comparing such rules shows how much the higher-order moments matter for stabilising the business cycle. I define the *Ramsey rule* as the rule that fully tracks the neutral rate, setting the safe rate equal to $r^\star$ in (8). An N-moment rule instead truncates the expansion of the premium in (7) after N terms:

$$r^{f,(N)} \equiv g + \rho - \sum_{n=1}^{N} \Upsilon_n \kappa_n = g + \rho - \mathcal{W}^{(N)}.$$

Since $\Upsilon_1 = 0$, the rule $N = 1$ ignores risk altogether and sets the rate at the deterministic Ramsey rate, $g + \rho$; $N = 2$ accounts for variance, as in the Gaussian case; and $N = 3, 4$ add skewness and excess kurtosis. When the rule understates the premium, it sets the safe rate above the neutral rate, and current output falls below potential by the size of the rate error.

Figure 4 compares the different rules for a disaster size of $u = -0.35$ as the disaster probability rises from zero to 5%. Under the Ramsey rule, the safe rate falls one for one with the neutral rate, and output stays at potential. Monetary policy fully stabilises output against disaster risk.

The $N = 1$ rule ignores risk altogether. As disaster risk rises, it keeps the safe rate at the deterministic Ramsey rate, 3.02%, while the neutral rate falls. Policy therefore becomes increasingly contractionary: at $p = 5\%$, the safe rate is about 401 basis points above the neutral rate, and current output falls to 96.07% of potential through the negative wealth effect.

The $N = 2$ rule now also tracks variance. It lowers the safe rate as disaster risk rises and stabilises much of the fall in output, but at $p = 5\%$ it still leaves the rate about 152

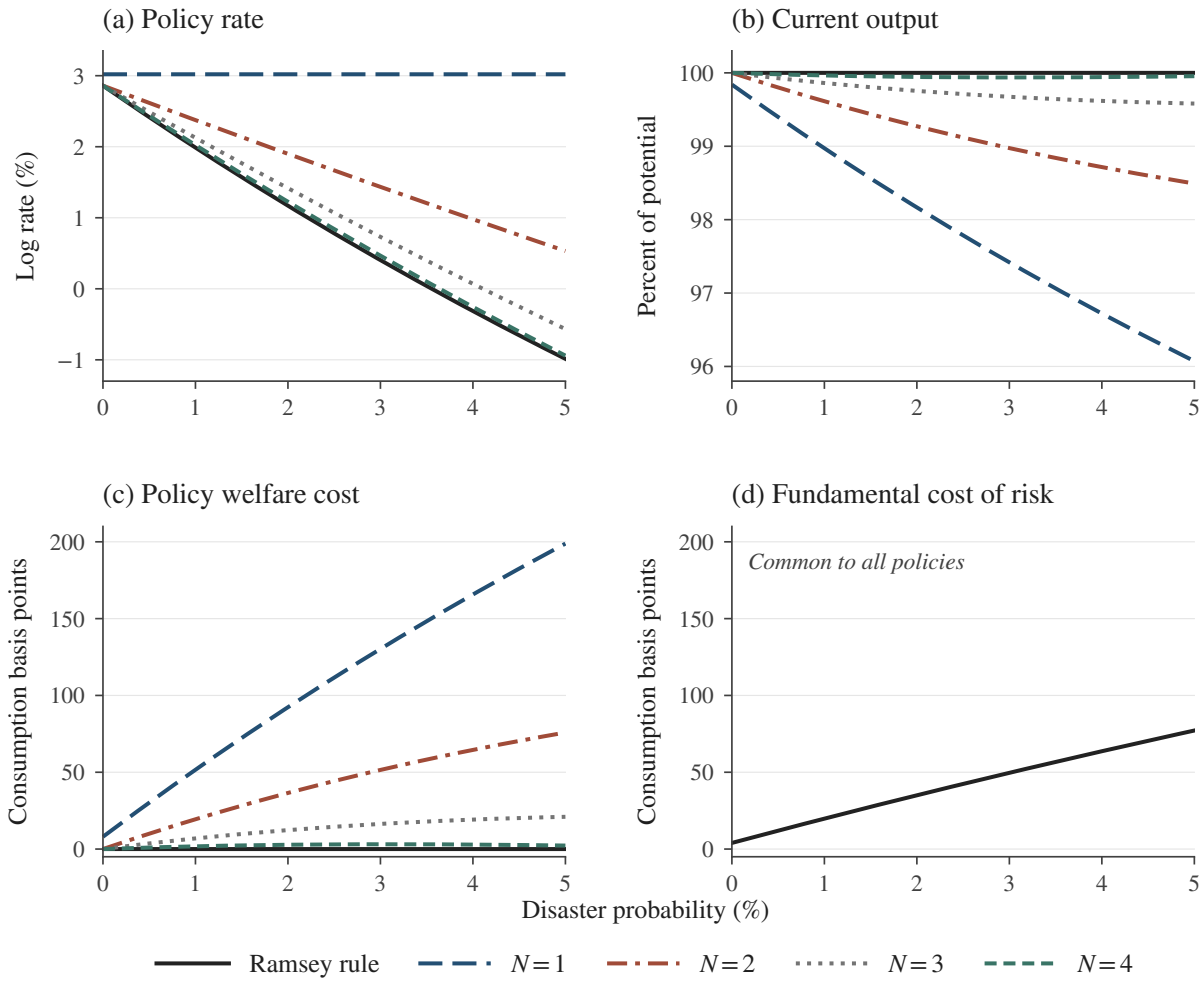


Figure 4: Policy and welfare under truncated-moment rules

Notes: The disaster jump is $u = -0.35$. The Ramsey rule fully tracks the neutral rate and keeps output at potential. Rules $N = 1, \ldots, 4$ know g and approximate only the premium. Panel (c) reports the policy cost, $\mathrm{CE}_N^{\mathrm{pol}}$; panel (d) reports the fundamental cost of risk, $\mathrm{CE}^{\mathrm{risk}}$, which is common to all policies. Under the Ramsey rule, the policy cost is zero at every disaster probability, although households still bear the fundamental cost of risk. Output is expressed as a percentage of potential, and welfare costs are in consumption basis points. All rules leave the safe rate at or above the neutral rate, and their ranking holds for these parameter values.

basis points too high and output at 98.49% of potential. Because this rule matches the first two moments, the policy error comes only from the moments beyond the second, and it equals the 1.52 percentage points found for the variance-matched Gaussian economy in Section 2.3. Tracking skewness ($N = 3$) closes most of the remaining gap, cutting the error to about 42 basis points and raising output to 99.58% of potential. Adding excess kurtosis ($N = 4$) nearly replicates the Ramsey rule, with an error of 4.55 basis points.

Panels (c) and (d) compare the welfare costs. Panel (c) reports the policy cost of each rule, CE_N^{pol}: the share of consumption at both dates that households would give up to have output at potential rather than at the level the rule delivers. Because both allocations face the same risk, this cost comes only from policy. Panel (d) reports the fundamental cost of risk, the Lucas measure of Section 2.3, which is the same under every rule. At $p = 5\%$, the policy cost of the $N = 1$ rule, about 199 basis points of consumption, is larger than the fundamental cost of risk, about 77 basis points, and even the cost of the $N = 2$ rule, about 76 basis points, is close to it. Failing to track the risk-adjusted neutral rate can therefore cost households more than the business-cycle risk itself.

Central banks therefore need to track movements in the higher-order moments of output growth, and not only in its variance: tracking up to the fourth cumulant almost eliminates the policy cost, which remains sizeable when the central bank tracks variance alone.

Generality of fourth-cumulant tracking. In fact, the $N = 4$ policy rule remains accurate well beyond the Bernoulli–Gaussian disaster example. Across eight other distributions of output growth, holding preferences, expected future output and the variance of G fixed (Online Appendix Table 2), tracking up to kurtosis keeps the policy error below seven basis points, whereas tracking only variance misses by up to 83 basis points, so output again falls short of potential by up to that amount. However, some familiar fat-tailed distributions, such as the untruncated Cauchy, Student's t, F and inverse-gamma, lie outside the framework, because their tails are so heavy that the cumulant-generating function, and with it the premium and the neutral rate, is not well defined.

3 Disaster Risk and Business-Cycle Fluctuations

In this section, I ask whether perceived disaster risk affects the macroeconomy. I measure disaster risk directly from option prices and identify shocks to it in a vector autoregression (VAR). I then test whether these shocks behave like negative demand shocks and lower output and inflation together, as the mechanism of Section 2 suggests. Appendices A and B give the data and the specifications, and report all robustness results.

3.1 From put prices to disaster probabilities

I use the two-period model of Section 2 to show how the probability of a disaster can be recovered from option prices. To simplify the exposition, I set $\sigma = 0$ in the Bernoulli–

Gaussian example, so future consumption, $c_1 = z_1$, takes only two values. In the normal state it is $c_N = e^g/(1 - p + pe^u)$, which keeps expected future output at e^g; in a disaster, which occurs with physical probability p, it is $c_D = e^u c_N$, where $e^u < 1$ because the log disaster jump u is negative. Since equity is a claim on future consumption, it pays c_N in the normal state and c_D in a disaster. I introduce a put option on equity with a strike price K between these two payoffs, $c_D < K < c_N$. The put then pays $K - c_D$ in a disaster and nothing otherwise.

Let M_N and M_D denote the stochastic discount factor in the normal and disaster states. The price of the put is the probability of a disaster times the value households place on its payment:

$$P^{\text{put}}(K) = pM_D(K - c_D). \tag{10}$$

A risk-free bond pays one unit in both states and has price $e^{-r^f} = (1 - p)M_N + pM_D$. Under the preferences in (3), $M_D/M_N = (c_D/c_N)^{-\gamma} = e^{-\gamma u} > 1$: households value a unit of consumption more in a disaster, when consumption is low. The bond price therefore pins down the level of the stochastic discount factor, and preferences pin down its relative value across the two states. Substituting both into (10) and solving for p gives

$$p = \frac{e^{r^f} P^{\text{put}}(K)}{e^{r^f} P^{\text{put}}(K) + e^{-\gamma u}\left[K - c_D - e^{r^f} P^{\text{put}}(K)\right]}. \tag{11}$$

Everything on the right-hand side of (11) is either observed or pinned down by the model. The risk-free rate, the put price and its strike are observed in the data, while risk aversion, γ, and the disaster size, u, are model parameters that also determine the disaster payoff, c_D. Given calibrated values of γ and u, (11) therefore recovers the physical disaster probability from option prices. Under rational expectations, the probability investors perceive, and hence the recovered p, equals the true probability of a disaster. If investors instead learn about tail risk from the data, a single rare event can raise the perceived probability persistently, with lasting effects on the economy (Kozlowski et al., 2020).

3.2 Measuring disaster risk

I measure perceived disaster risk with the approach of Barro and Liao (2021), who apply the logic of (11) to short-maturity, far-out-of-the-money S&P 500 puts. They allow disasters to vary in size, with a power law for the ratio of normal to disaster consumption, and account for falls in equity prices when the disaster probability rises before any disaster. Option prices then identify the disaster probability times a pricing factor that depends on risk aversion and the size distribution of disasters; as in (11), calibrating this factor

supplies the level. Their measure is a disaster probability per year, strictly the annual rate at which disasters arrive, rather than the one-period probability p of Section 3.1.

Figure 5 plots the estimated annual probability, monthly from September 1994 to August 2025, together with the Chicago Board Options Exchange Volatility Index (VIX). The VIX is a standard proxy for uncertainty in the literature on uncertainty shocks (Bloom, 2009; Basu and Bundick, 2017). Perceived disaster risk is low in calm periods, below 2 percent for most of 2004–2007, and rises sharply in crises. It reached about 42 percent in November 2008, after the collapse of Lehman Brothers, and 24 percent in autumn 2011, during the euro-area debt crisis. At the onset of the pandemic, in March 2020, it rose to 30 percent.

The two series co-move closely, with a monthly correlation of 0.93. However, they measure different kinds of risk. The VIX summarises expected volatility over the next 30 days using a wide range of put and call options (Martin, 2011), so it is closest to the variance of returns and reflects *upside* as well as downside risk. The disaster probability captures the left tail of the distribution, where the higher-order moments of the model in Section 2 matter. Figure 5 also shows that the VIX is noisier. It moves sharply from day to day, while the disaster probability moves more smoothly.

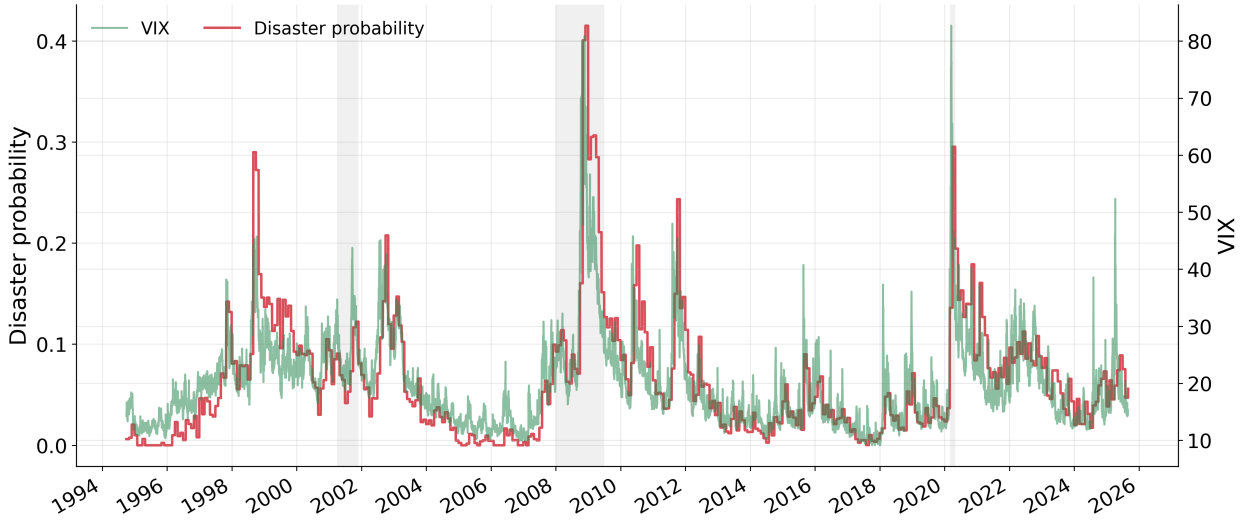


Figure 5: Estimated disaster probability and the VIX

Notes: The red line depicts the annualised disaster probability extracted from S&P 500 put options following Barro and Liao (2021), September 1994 to August 2025. The green line depicts the CBOE's VIX. Shaded areas indicate NBER recessions.

3.3 Disaster risk and aggregate demand

Identification. I estimate a quarterly VAR,

$$Y_t = c + \sum_{\ell=1}^{3} B_\ell Y_{t-\ell} + u_t, \qquad \mathbb{E}[u_t u_t'] = \Sigma_u, \tag{12}$$

where Y_t contains eleven variables: disaster probability, p_t; per-capita logs of real GDP, consumption, investment and employment; GDP-deflator inflation; the 5-year Treasury yield; real stock and Treasury-bill returns; the log dividend–price ratio; and the VIX. The financial variables show how stock prices and broader market volatility move with disaster risk. The baseline sample is 1996Q1 to 2019Q4.[4]

To identify disaster-risk shocks, I use the max-share approach (Uhlig, 2003; Basu et al., 2021), implemented in the frequency domain as in Angeletos et al. (2020). The VAR residuals u_t are combinations of many underlying disturbances. Among all uncorrelated combinations of these residuals, the method picks the one that explains as much of the variation in disaster probability as possible. I call this disturbance the disaster-risk shock and trace its effects through the economy. The method places no restrictions on how activity and inflation respond.

Responses to a disaster-risk shock. Figure 6 plots the impulse responses of the main business-cycle variables to a one-standard-deviation disaster-risk shock. The first panel shows the response of disaster probability itself, the target of the max-share procedure. The shock raises the annual probability of a disaster by about 4 percentage points on impact, and the effect decays within two years. It accounts for over half of the variation in disaster probability at business-cycle frequencies. Stock prices fall at the same time: the real stock return is 6.6 percent lower on impact.

Output, consumption, investment and employment all contract. Output falls by 0.28 percent on impact and by 0.39 percent after three quarters. Consumption reaches its trough after one quarter, with a decline of 0.31 percent. Investment and employment take longer, falling by 1.0 percent after three quarters and 0.31 percent after five, so the contraction in output, investment and employment builds over several quarters. The shock accounts for about 15 percent of the business-cycle variation in output and investment.

The joint fall in consumption and investment is an important pattern, because it distinguishes between explanations of how risk affects the economy. In a neoclassical

[4]The start is the first quarter with the full set of VAR variables, including the option-based volatility series. The end precedes the COVID-19 pandemic, whose quarters dominate the reduced-form residual covariance and widen the bootstrap bands considerably.

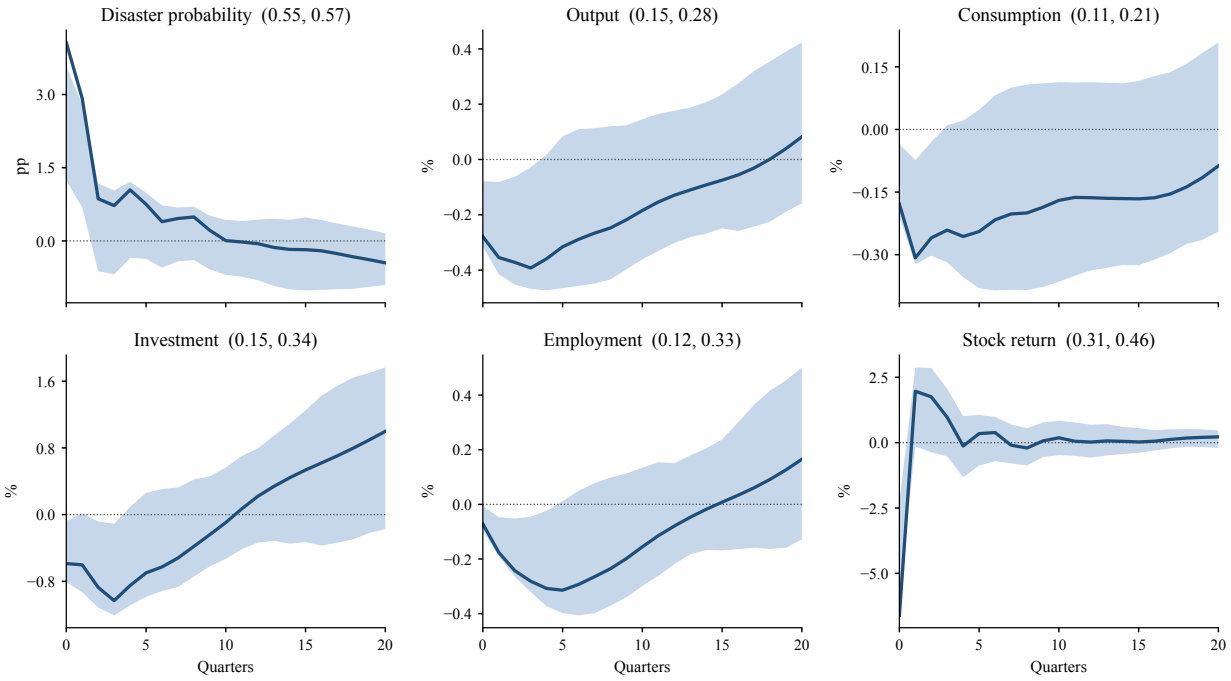


Figure 6: Real effects of a disaster-risk shock

Notes: Impulse responses to the max-share disaster-risk shock, the shock that explains the most variation in disaster probability (1996Q1–2019Q4). Horizons run from impact to 20 quarters. In each panel title, the first number in parentheses is the shock's share of the forecast-error variance at business-cycle frequencies (periods of 6–32 quarters), and the second its share at near-unconditional frequencies (periods of 2–500 quarters). Responses of the disaster probability, yields and inflation are in percentage points; other variables are in percent. Shaded areas are 80% bootstrap bands, from the 10th to the 90th percentile.

economy with flexible prices, output is largely pinned down by productivity, so if precautionary saving lowers consumption, investment must rise. Barro and King (1984) show that, in such an economy, shocks other than those to productivity tend to move consumption and investment in opposite directions. Generating a joint fall in consumption and investment after a rise in risk is a recurring challenge in the uncertainty-shock literature (Bloom, 2009; Basu and Bundick, 2017; Bloom et al., 2018). With sticky prices, by contrast, output is determined by demand: when households cut consumption, firms produce less, and investment falls with output. A model in which demand determines output therefore predicts the joint fall. The response of inflation then shows whether the disaster-risk shock is a demand shock or a supply shock.

Inflation and interest rates. Inflation falls with activity (Figure 7). One quarter after the shock, inflation is 0.4 percentage points lower at an annual rate, and the 5-year Treasury yield is 22 basis points lower, after a 11-basis-point decline on impact. Both responses are

distinguishable from zero. The VIX also rises, by more than 4 index points on impact.

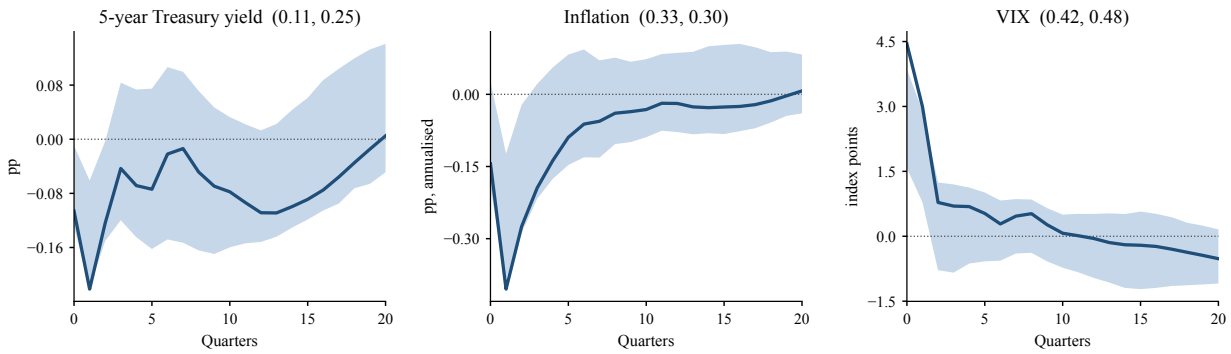


Figure 7: Financial and nominal responses to a disaster-risk shock

Notes: Impulse responses of the 5-year Treasury yield, inflation (GDP deflator, annualised), and the VIX to the max-share disaster-risk shock (1996Q1–2019Q4). Horizons run from impact to 20 quarters. In each panel title, the first number in parentheses is the shock's share of the forecast-error variance at business-cycle frequencies (periods of 6–32 quarters), and the second its share at near-unconditional frequencies (periods of 2–500 quarters). Shaded areas are 80% bootstrap bands, from the 10th to the 90th percentile.

Disaster-risk shocks therefore appear in the data as negative demand shocks: they lower output, consumption, investment, employment, inflation and Treasury yields together, whereas an adverse supply shock would raise inflation as output falls. The fall in activity is consistent with the mechanism of Section 2. There, higher disaster risk raises the equity premium and lowers the price of equity and household wealth, so it depresses demand when the policy rate does not track the fall in the neutral rate. This evidence supports interpreting the risk-premium shock of estimated New Keynesian models (Smets and Wouters, 2007) as, at least in part, a shock to perceived disaster risk, in line with the risk-centric view of recessions of Caballero and Simsek (2020).

These results differ from those for the risk-premium shock of Basu et al. (2021), who also use a max-share procedure but target the 5-year equity risk premium rather than disaster probability. Their shock raises inflation and Treasury yields, so it behaves more like a supply shock, while targeting disaster probability produces falling inflation and yields. The equity premium mixes downside risk with other disturbances, whereas disaster probability measures downside risk more directly.

3.4 Alternative identification schemes and controls

The max-share shock is a disaster-risk shock only if surprises in p_t mostly reflect news about disaster risk, rather than about growth, credit conditions or monetary policy. When other shocks also move p_t, the max-share shock can mix them in (Dou et al., 2025). I therefore re-estimate the responses under two alternative identification schemes from the literature, each resting on different identifying assumptions, and then remove productivity shocks before identifying the disaster-risk shock. A final check controls for lagged risk-neutral variance.

Identification by heteroskedasticity. I adapt the approach of Brunnermeier et al. (2021) and identify the shocks from shifts in their relative volatility across six historical regimes. The identifying assumption is that the effects of each shock on the economy stay the same from one regime to the next. If disaster-risk shocks were far more volatile during the financial crisis than in the calm years before it, disaster probability and output would move together more closely during the crisis. With the effects unchanged, that extra comovement must come from disaster-risk shocks, and its size measures how strongly output responds to them. This approach imposes no timing or sign restrictions. Relative to the baseline, I move to monthly data over 1994M9–2025M8 and allow for fat-tailed shocks. The system has seven variables: disaster probability, real activity, the price level, the policy rate (the Wu and Xia (2016) shadow rate while it is available), the term spread, the excess bond premium of Gilchrist and Zakrajšek (2012) and real equity wealth.

Recursive identification. I also estimate quarterly VARs over 1996Q1–2025Q2, a sample that includes the pandemic, under two standard timing assumptions. In the first, disaster probability is ordered after the slow-moving macroeconomic variables (Gilchrist and Zakrajšek, 2012), so the shock is the part of the change in disaster risk that same-quarter macroeconomic news does not explain. In the second, disaster probability is ordered first (Basu and Bundick, 2017), so that it can affect all other variables within the quarter.

Under both alternative schemes, the baseline result holds: disaster-risk shocks appear as negative demand shocks, lowering economic activity and inflation together. Under identification by heteroskedasticity, the fall in activity shows up in employment, while monthly GDP is roughly flat. Re-estimating the baseline max-share VAR through 2025Q2 leaves the impact pattern unchanged, although the bands widen because the pandemic quarters dominate the estimates.

Controlling for productivity shocks. The max-share shock could also pick up news about productivity, which moves equity prices and with them the disaster probability. I check this with the procedure of Basu et al. (2021): I first identify a productivity shock and then choose the disaster-risk shock among the shocks orthogonal to it. I use two productivity shocks: the surprise in the utilisation-adjusted productivity series of Fernald (2014), and a supply shock, the combination of VAR innovations that moves productivity and inflation most strongly in opposite directions. Removing either leaves the disaster-risk shock almost unchanged. Its correlation with the baseline shock is 0.994 and 0.976, it explains about the same share of fluctuations in disaster probability, and output, consumption, investment, employment and inflation fall much as in the baseline (Appendix A.2.2). The shock behind the joint fall in activity and inflation is therefore not a productivity shock in disguise.

Controlling for lagged risk-neutral variance. A further check changes the variables in the baseline VAR rather than its identification. I replace the VIX with SVIX^2, the risk-neutral variance of the market return of Martin (2017), so that the VAR controls for lagged risk-neutral variance. Output still falls after a disaster-risk shock, and its 80 percent pointwise interval excludes zero on impact and one quarter later. The result is consistent with the mechanism of Section 2, in which left-tail risk lowers demand. When the VAR includes both the VIX and SVIX^2, the interval for output includes zero at every horizon, and stronger restrictions that keep SVIX^2 unchanged on impact or for a year yield no output response distinguishable from zero. Appendix A.2.3 gives the details.

These empirical results motivate my quantitative model. In the model, higher disaster risk lowers the neutral rate, and output falls when policy does not track it.

4 A Quantitative New Keynesian Model with Disaster Risk

To measure how far disaster risk moves the neutral rate, and what happens when monetary policy fails to track it, I extend the two-period model of Section 2 to a quantitative New Keynesian model in continuous time (Woodford, 2003; Galí, 2015), in which disasters arrive as sudden jumps. The quantitative model adds four features. Prices, which were fixed in the two-period model, now adjust gradually, so inflation as well as output responds to demand. A disaster is followed by a gradual recovery, and the risk of disasters moves over time. Finally, monetary policy follows a Taylor rule that can range from ignoring the risk-adjusted neutral rate to tracking it fully. Appendix E.2 sets out the main steps of the derivations, and Online Appendix D gives the full proofs.

4.1 Households

Households choose consumption C_t, labour L_t and holdings of a safe bond. They have Epstein–Zin preferences, which separate risk aversion, $\gamma > 1$, from the elasticity of intertemporal substitution, set to one as in Wachter (2013). As risk aversion exceeds the inverse of that elasticity, households care not only about their current consumption but also about how shocks change their future prospects. A rise in disaster risk therefore matters to them *before* any disaster has occurred. In continuous time (Duffie and Epstein, 1992), continuation utility U_t satisfies, for any $t < T$,

$$U_t = \mathbb{E}_t\left[\int_t^T f(C_u, L_u, U_u)\, du + U_T\right],$$

with aggregator

$$f(C, L, U) = \rho(1-\gamma)U\left[\log C - \frac{L^{1+\varphi}}{1+\varphi}\right] - \rho U \log\Big((1-\gamma)\rho U\Big),$$

where $\rho > 0$ is the rate of time preference and $1/\varphi$, with $\varphi \geq 0$, is the Frisch elasticity of labour supply.

Households earn wages, receive the profits of the firms they own, pay lump-sum taxes, and can save or borrow through a short-term safe bond. Their flow budget constraint is

$$db_t = \left(r_t b_t + \frac{W_t L_t}{P_t} + \frac{\Pi_t}{P_t} - T_t - C_t\right) dt, \tag{13}$$

where b_t denotes real bond holdings, r_t the real safe interest rate, W_t the nominal wage, P_t the price level, Π_t nominal profits and T_t real taxes. Because bonds are in zero net supply, $b_t = 0$ in equilibrium. Taking wages, profits and taxes as given, households can adjust labour at every date, including after a disaster, and optimal labour supply equates the real wage to the marginal rate of substitution between leisure and consumption:

$$\frac{W_t}{P_t} = C_t L_t^{\varphi}. \tag{14}$$

4.2 Firms and productivity

Firms. A competitive final-good sector aggregates a continuum of differentiated intermediate goods, $j \in [0, 1]$, with elasticity of substitution $\varepsilon > 1$. Each intermediate firm faces the

usual Dixit–Stiglitz demand and produces with a linear technology in labour,

$$Y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\varepsilon} Y_t, \qquad Y_t(j) = A_t L_t(j), \tag{15}$$

where $P_t(j)$ is firm j's price, P_t the aggregate price index, Y_t aggregate output, $L_t(j)$ the firm's labour input and A_t aggregate productivity, which is common to all firms. A constant employment subsidy, financed by lump-sum taxes, offsets the monopoly markup.

Price setting. Intermediate firms set prices facing the demand in (15) and quadratic costs of adjusting them (Rotemberg, 1982). They choose price paths to maximise households' valuation of profits, net of labour and price-adjustment costs. Changing firm j's price at rate $\pi_t(j) = \dot{P}_t(j)/P_t(j)$ uses $Y_t[\pi_t(j) - \pi^\star]^2/(2\eta)$ units of the final good, where $\pi^\star$ is target inflation and $\eta > 0$, so a smaller η means stickier prices.[5] The price level cannot jump, but inflation can, including when a disaster occurs. Because adjustment is costly, firms set prices taking into account expected future marginal costs as well as current ones, which yields a forward-looking Phillips curve.

Productivity process. Aggregate productivity is exogenous. Its log combines a growing permanent component with a mean-reverting component:

$$\log A_t = a_t^P + a_t^r, \qquad da_t^P = (\bar{g} + g_t)\,dt, \qquad dg_t = -\kappa_g g_t\,dt + \sigma_g\,dZ_t^g.$$

The permanent component grows at rate $\bar{g}$ plus a deviation g_t that slowly returns to zero, so expected productivity growth, and with it the neutral rate, can change over time. The component a_t^r, log productivity relative to the permanent component, is mean-reverting, with small continuous shocks and rare, sudden falls:

$$da_t^r = -\kappa_a a_{t-}^r\,dt + \sigma_A\,dZ_t^a + J_A\,dN_t, \qquad \kappa_a > 0, \quad \sigma_A \geq 0, \quad J_A < 0,$$

where Z_t^g and Z_t^a are standard Brownian motions, N_t counts the disasters that have occurred, and a_{t-}^r is productivity just before any disaster at date t. The first two terms describe productivity in "normal times" and are the continuous-time counterpart of the autoregressive process in the textbook New Keynesian model (Galí, 2015). Productivity is

[5] I treat Rotemberg adjustment costs as real resources used up in changing prices, not as fees rebated to households. Rebating them would leave firms' pricing unchanged but imply $C_t = Y_t$, so inflation would have no direct welfare cost. In Calvo models that cost comes from the dispersion of relative prices, which is absent here because all firms charge the same price (Woodford, 2003; Galí, 2015).

pulled back towards the permanent component at rate κ_a, and σ_A scales the small, normally distributed (Gaussian) shocks that drive ordinary fluctuations.

The third term is where disasters enter. Each time a disaster occurs, productivity, and with it output at a given level of employment, falls by the fraction $1 - e^{J_A}$. This term models the rare but very large contractions that Barro (2006) and Barro and Ursúa (2008) document. In Barro (2006), such losses are permanent. Here, a disaster leaves the permanent component unchanged and mean reversion gradually unwinds the loss, so a disaster is a deep contraction followed by a gradual recovery (Gourio, 2008; Nakamura et al., 2013).

Disaster risk. Disasters arrive with intensity

$$\lambda_t = \lambda_0 e^{-\chi_a a^r_{t-} + \chi_\xi \xi_t}, \qquad d\xi_t = -\kappa_\xi \xi_t \, dt + \sigma_\xi \, dZ^\xi_t, \tag{16}$$

where $\lambda_0 > 0$ and $\chi_a, \chi_\xi \geq 0$, so the probability of a disaster over a short interval dt is approximately $\lambda_t \, dt$. Disaster risk has two sources. When $\chi_a > 0$, it is *countercyclical*: disasters become more likely when productivity is already low. The second source is a disaster-risk state ξ_t, which mean-reverts with its own shock, Z^ξ_t, independent of Z^a_t and Z^g_t. Disaster risk can therefore rise before productivity has fallen, as the option-based measure of disaster risk does in crises (Figure 5). Without this state, disaster risk would rise only once productivity had fallen, so the neutral rate could not fall ahead of a downturn. Time-varying disaster probability of this kind is standard in the rare-disaster literature (Gabaix, 2012; Gourio, 2012; Wachter, 2013), and its shocks are the model counterpart of the disaster-risk shocks identified in Section 3. Setting $\chi_a = \chi_\xi = 0$ recovers the constant intensity of Barro (2006).

Beyond the growth of the permanent component, expected productivity growth has two parts: the recovery towards that component, less the losses that disasters are expected to inflict. A realised disaster lowers a^r_t but leaves a^P_t, ξ_t and g_t unchanged, so it pulls these parts in opposite directions: when $\chi_a > 0$, it makes further disasters more likely, but it also speeds up the expected recovery, since productivity now lies further below its permanent component.

Unlike the two-period model, which leaves the distribution of output growth unrestricted, the quantitative model specialises risk to normal shocks and disasters of a fixed size. Disaster risk depends on productivity and on the disaster-risk state rather than on the output gap, so a downturn caused by weak demand does not itself make disasters more likely.

Market clearing. In a symmetric equilibrium, all firms charge the same price. Let $\pi_t = d\log P_t/dt$ denote inflation and $\widehat{\pi}_t = \pi_t - \pi^\star$ its deviation from target. Production and goods-market clearing then imply

$$Y_t = A_t L_t, \qquad Y_t = \Psi_t C_t, \qquad \Psi_t = \left(1 - \frac{\widehat{\pi}_t^2}{2\eta}\right)^{-1},$$

with $\widehat{\pi}_t^2 < 2\eta$ so that consumption is positive. The resource wedge Ψ_t is the gross output needed to deliver one unit of consumption once price adjustment has been paid for. At target inflation, $\Psi_t = 1$ and consumption equals output. Away from target, price adjustment uses up resources, so $\Psi_t > 1$, and a given inflation deviation uses up more of them when prices are stickier (a smaller η).

4.3 The neutral rate and monetary policy

With flexible prices, the economy reaches its natural allocation: households supply one unit of labour, $L_t^n = 1$, and output and consumption equal productivity, $C_t^n = Y_t^n = A_t$. Because the subsidy offsets firms' markup, the natural allocation is efficient (Rotemberg and Woodford, 1997). It is therefore the benchmark that monetary policy aims to achieve once risk is taken into account. I define the output gap as the log deviation of output from its natural level, $x_t = \log(Y_t/Y_t^n) = \log(Y_t/A_t)$. The neutral rate is the real safe rate that supports the natural allocation. Unlike in log-linear New Keynesian models, it keeps the compensation for risk, including disasters.

The neutral rate also depends on how households value the future. Because productivity grows along a permanent component, I summarise continuation utility by its value relative to that component: $v_t = v(a_t^r, \xi_t, g_t, i_t)$ is the log of households' certainty-equivalent consumption, adjusted for hours worked, relative to the permanent component of productivity. The corresponding value in the natural allocation, $v_t^n = v^n(a_t^r, \xi_t, g_t)$, does not depend on monetary policy.

Proposition 4.1 (Risk and the neutral rate)**.** *With Epstein–Zin preferences and a unit elasticity of intertemporal substitution, the neutral rate is*

$$r_t^\star = \underbrace{\rho + \frac{\mathbb{E}_t[d\log A_t]}{dt}}_{\text{Ramsey rate}} - \underbrace{\mathcal{R}_t^n}_{\text{risk adjustment}}, \tag{17}$$

where $\mathcal{R}_t^n$, the risk adjustment in the natural allocation, is

$$\mathcal{R}_t^n = \tfrac{1}{2}\sigma_A^2 + (\gamma - 1)\sigma_A^2 \partial_{a^r} v_t^n + \lambda_t \left[e^{-(\gamma-1)\Delta_J v_t^n}(e^{-J_A} - 1) + J_A \right], \tag{18}$$

and v^n is the natural value defined above, with $\Delta_J v_t^n$ its change when a disaster occurs.

The neutral rate in (17) has two parts. The first, time preference plus expected productivity growth, is the Ramsey rate of the textbook model: the rate is higher when households are less patient or expect faster growth. Expected growth is $\mathbb{E}_t[d \log A_t]/dt = \bar{g} + g_t - \kappa_a a_t^r + \lambda_t J_A$. The second part is the risk adjustment $\mathcal{R}_t^n$, as in Proposition 2.1. Riskier productivity raises households' demand for the safe bond. Because bonds are in zero net supply, the safe rate must fall to clear the market, and the risk adjustment measures that fall.

The risk adjustment has three terms.[6] The first is the standard precautionary-saving term, which ordinary productivity shocks would generate even under log utility. The second is a continuation-value term, which Epstein–Zin preferences add. When higher productivity raises continuation utility, $\partial_{a^r} v_t^n > 0$, a productivity shock is riskier than its effect on current consumption alone suggests, so households require more compensation for bearing it. The third term is the risk compensation for the disaster jump. A disaster cuts consumption and, with Epstein–Zin preferences, also lowers continuation utility, so households insure against it by holding more of the safe bond.

Monetary policy. The central bank's target for the nominal rate follows a Taylor rule (Taylor, 1993), with the usual responses to inflation and the output gap and an intercept that can track the neutral rate fully, partly or not at all:

$$i_t^T = \pi^\star + \bar{r} + \omega(r_t^\star - \bar{r}) + \phi_\pi \widehat{\pi}_t + \phi_x x_t. \tag{19}$$

Here $\bar{r}$ is the rule's real intercept, a constant estimate of the neutral rate, and ϕ_π and ϕ_x are the responses to inflation and output gaps. The weight $\omega \in [0, 1]$ sets the fraction of the neutral rate's deviation from $\bar{r}$ that passes into the intercept. At $\omega = 0$ the rule is a standard Taylor rule whose intercept does not respond to risk, and at $\omega = 1$ the intercept equals the neutral rate, so the rule tracks it fully.

The policy rate i_t moves gradually towards this target, so the rule is an inertial Taylor

[6] At the natural allocation, consumption does not respond on impact to the shocks to ξ_t or g_t, so their volatilities do not enter (18) as separate variance terms. The volatility of the disaster-risk state can still change the exposures of v^n to productivity and disasters. Because trend growth is independent of the other states, v^n is linear in g_t with slope $1/(\rho + \kappa_g)$: the volatility of g_t shifts the level of v^n but not those exposures.

rule:

$$di_t = \kappa_i(i_t^T - i_t)\,dt + \sigma_i\,dZ_t^i. \tag{20}$$

The adjustment speed $\kappa_i > 0$ measures how quickly the policy rate closes its gap to the target. The inertia reflects the well-documented practice of central banks moving their policy rates in steps rather than jumping to the rule's prescription. Policy rules estimated in the empirical literature also typically include interest-rate smoothing terms (Clarida et al., 2000; Coibion and Gorodnichenko, 2012). I also add a monetary policy shock, driven by a standard Brownian motion Z_t^i independent of the other shocks, which captures movements in the policy rate that the rule does not explain. Because the policy rate adjusts gradually, it is itself a state variable.[7]

The real-rate gap, $r_t - r_t^\star$, where $r_t = i_t - \pi_t$, measures the stance of policy. When the gap is positive, policy is contractionary. Because the policy rate cannot jump, a sudden fall in the neutral rate widens the gap on impact under any policy rule. A policy rule that only partly tracks the neutral rate then leaves some of the gap open even as the policy rate adjusts.

The IS relation. Households' choice between consuming and saving in the safe bond, together with the resource constraint, gives the exact IS relation in the output gap:

$$\frac{\mathbb{E}_t[dx_t]}{dt} = (r_t - r_t^\star) + (\mathcal{R}_t - \mathcal{R}_t^n) + \frac{\mathbb{E}_t[d\log\Psi_t]}{dt}, \tag{21}$$

where the expectations include both continuous movements and disasters, and $\mathcal{R}_t$ is the risk adjustment of (18) evaluated at the actual rather than the natural allocation. The first term is the textbook intertemporal-substitution effect. Holding the expected future output gap and the other terms fixed, a real rate above the neutral rate lowers current demand, because households postpone consumption. The equation describes the expected *change* in the output gap, so a high real rate shows up as a low output gap today followed by an expected recovery.

The second term is absent from the textbook model. It is the difference between the risk adjustment under the actual allocation and under the flexible-price allocation. This term is zero when policy stabilises the economy at the flexible-price allocation. If policy lets the output gap fall when productivity falls, or deepen in a disaster, consumption becomes riskier than productivity, and households save more as a precaution. With Epstein–Zin preferences, households also care about how policy changes their future welfare after a

[7] Fiscal policy is Ricardian, or *passive* in the terminology of Woodford (2003). I leave the interaction between disaster risk and fiscal policy to future work.

shock, which can strengthen or weaken this effect.

The last term comes from changes in the resources that price adjustment absorbs. Around the natural allocation, the wedge differs from one only at second order in the inflation gap, so this term drops out of a first-order approximation. The quantitative model keeps the full nonlinear relation.

The Phillips curve. Firms' pricing decisions give the exact Phillips curve:

$$\rho \Psi_t \widehat{\pi}_t = \frac{\mathbb{E}_t^v[d(\Psi_t \widehat{\pi}_t)]}{dt} + \eta(\varepsilon - 1)\left(e^{(1+\varphi)x_t} - \Psi_t\right), \tag{22}$$

where $\mathbb{E}_t^v$ is the expectation under households' stochastic discount factor, with which firms value future price-adjustment costs. Firms therefore put more weight on states in which households' future welfare is low, such as after a disaster. The last term is proportional to the deviation of real marginal cost from its natural level, and the resource wedge Ψ_t scales inflation only away from target.

To first order in the output and inflation gaps around the natural allocation, the curve takes the form of the textbook New Keynesian Phillips curve, except that expectations weight states by continuation utility:

$$\rho \widehat{\pi}_t \simeq \frac{\mathbb{E}_t^{v^n}[d\widehat{\pi}_t]}{dt} + \kappa x_t, \qquad \kappa = \eta(\varepsilon - 1)(1 + \varphi), \tag{23}$$

where $\mathbb{E}_t^{v^n}$ uses continuation utility in the natural allocation and κ is the slope. Stickier prices (a smaller η) lower κ and flatten the curve, so inflation responds less to a given output gap. Solved forward, the curve sets inflation equal to κ times the discounted expected sum of future output gaps, with more weight on states in which continuation utility is low. A persistent negative output gap therefore keeps inflation below target. Higher disaster risk can lower *both* output and inflation when policy fails to track the neutral rate, even when current productivity is unchanged. A joint fall in output and inflation after a rise in disaster risk is what Section 3 finds in the data.

Asset prices. Finally, I model equity as a leveraged claim on consumption (Wachter, 2013), whose dividends respond more strongly than consumption to fluctuations around their common permanent component:

$$D_t = \exp\{a_t^P + \phi_d c_t\}, \qquad \phi_d \geq 1,$$

where $c_t = \log C_t - a_t^P$ is log consumption relative to the permanent component. The parameter ϕ_d plays the role of leverage. When $\phi_d > 1$, a disaster cuts dividends by more than consumption, but in the long run dividends grow with consumption, so a change in trend growth moves both by the same amount.[8] I price equity with households' stochastic discount factor and, following the rare-disaster and long-run-risk literatures (Bansal and Yaron, 2004; Wachter, 2013), treat it as being in zero net supply. Equity therefore adds nothing to households' resources and does not affect the allocation.

The risk adjustment also shows up in asset prices. With a unit elasticity of intertemporal substitution, a claim to aggregate consumption is worth C_t/ρ under the transversality condition, and its expected log return, including dividends, exceeds the real safe rate by $\mathcal{R}_t$. At the natural allocation, this premium equals $\mathcal{R}_t^n$, so the compensation investors require for holding the consumption claim is exactly the amount by which risk lowers the neutral rate.

5 Quantitative Analysis

I now take the model of Section 4 to US data. The estimated model gives a quarterly path of the neutral rate, against which I assess the stance of monetary policy. I then use the model to quantify the effects that moments beyond the variance have on the neutral rate and how news of higher disaster risk affects the economy under different monetary policy rules, including at the onset of the financial crisis. Finally, I measure what failing to track the neutral rate costs and provide an alternative semi-structural estimate of the neutral rate. Appendix D gives the estimation details and additional results.

5.1 Calibration and estimation

Table 2 reports the model's parameters in three groups: those I fix using outside evidence, those I calibrate to asset prices and those I estimate by maximum likelihood.

Policy coefficients and price stickiness. I use the feedback coefficients of Taylor (1993), $\phi_\pi = 1.5$ and $\phi_x = 0.5$, applied to annualised inflation and the model's output gap, and set the inflation target to 2%. Like Nakamura et al. (2025), who vary the intercept of the Taylor rule, I fix these responses rather than estimate them alongside the weight on the neutral rate, because the two are hard to tell apart in the data. When the neutral rate falls,

[8] Equity could instead be valued as a claim to firms' profits. With sticky prices and flexible wages, profits rise relative to consumption when the output gap falls, so a procyclical profit claim would require replacing price rigidity with wage rigidity.

Table 2: Parameters of the estimated model

Parameter	Description	Value	Source (s.e.)
Estimated by maximum likelihood			
ω	Weight on the neutral rate	0.639	[0.59, 0.68]
κ_i	Speed of adjustment	5.36	(0.70)
$\bar{r}$	Real intercept of the target	0.0102	(0.0009)
σ_i	Policy shock	0.00057	(0.0028)
κ_ξ	Disaster-risk state: mean reversion	1.003	(0.31)
σ_ξ	Disaster-risk state: volatility	0.0778	(0.0077)
$\bar{p}$	Mean of the disaster measure	0.0620	(0.0138)
Calibrated to asset prices			
γ	Risk aversion	9.98	Asset returns
ρ	Rate of time preference	0.015	Asset returns
ϕ_d	Dividend loading	9.50	Asset returns
Fixed			
ϕ_π, ϕ_x	Responses to inflation, output gap	1.5, 0.5	Taylor (1993)
η	Price-adjustment parameter	0.153	Nakamura and Steinsson (2008)
ε	Demand elasticity	10	Standard
φ	Inverse Frisch elasticity	1	Standard
$\pi^\star$	Inflation target	0.02	Federal Reserve
J_A	Disaster size	−0.288	Calibration
κ_a	Recovery of productivity	0.0347	Calibration
λ_0	Disaster intensity at $a^r = \xi = 0$	0.030	Calibration
χ_a, χ_ξ	Loadings of the intensity	1.45, 10.74	Calibration
σ_A	Ordinary productivity shocks	0.026	GDP growth volatility
$\bar{g}$	Trend growth	0.025	CBO potential output
κ_g, σ_g	Persistence and volatility of g_t	0.061, 0.0020	Calibration

Notes: Rates are annual decimals; volatilities are per square-root year. Standard errors are in parentheses; the interval for ω is the 95 percent likelihood-ratio interval, holding the other parameters at their estimates. Calibration denotes values chosen by simulated moments (Table 8) in a version of the model with different policy coefficients. Ordinary productivity shocks are close to the 2 percent normal-times volatility of Barro (2006) and Wachter (2013).

inflation and output fall with it, so a rate cut could be a response to either. Fixing the responses also makes a rule with no weight on the neutral rate a standard Taylor rule. I set the price-adjustment parameter to $\eta = 0.153$, so that the Phillips curve has the slope of a Calvo model in which prices remain unchanged for four quarters on average, about as long as regular prices last in US microdata (Nakamura and Steinsson, 2008). With this value of η, the New Keynesian Phillips curve has a slope of about 0.17 in quarterly units. Calibrated New Keynesian models use slopes of a similar size (Galí, 2015).

Productivity and disasters. Disasters are too rare to measure from US data alone, so I fix their size and frequency using cross-country evidence. A disaster lowers productivity by 25% ($J_A = -0.29$), within the range of falls in GDP per capita during the disasters documented by Barro (2006) and Barro and Ursúa (2008). Productivity then recovers with a half-life of 20 years. A half-life this long fits evidence that output losses after financial crises persist for decades (Cerra and Saxena, 2008; Müller et al., 2025). At the steady state, the disaster intensity is 3.0% a year, close to the annual probability of entering a disaster of about 2.9% that Barro and Ursúa (2008) estimate for OECD countries. Ordinary productivity shocks have a volatility of 2.6% a year, which brings the model close to the volatility of quarterly GDP growth. Trend growth averages 2.5% a year, the mean growth of CBO potential output over the sample, and its deviations from this mean are small and persistent, with a half-life of about 11 years.

The disaster-risk state. I estimate the process for the disaster-risk state from the option-based disaster probability of Section 3 alone, following Seo and Wachter (2019) in using option prices to discipline a time-varying disaster probability. The measure is a noisy reading of the state around its mean,

$$p_t = \bar{p} + \xi_t + e_t^p,$$

where e_t^p is measurement noise, and I estimate the mean, $\bar{p} = 6.2\%$, and the persistence and volatility of ξ_t by maximum likelihood. Shocks to disaster risk fade quickly, with a half-life of about 8 months ($\kappa_\xi = 1.00$). The state enters the disaster intensity of (16), $\lambda_t = \lambda_0 e^{-\chi_a a_{t-}^r + \chi_\xi \xi_t}$, so the measure moves the intensity but does not set its level. At the steady state, $\xi_t = 0$ and the intensity is $\lambda_0 = 3.0\%$ a year, whatever the mean of the measure. With the calibrated loading $\chi_\xi = 10.7$, each percentage point by which the measure rises above its mean raises the intensity by about 11% of that level. The measure is higher than the model's probability of a disaster, averaging 6.3 percent over 1994–2025, because it counts equity-market falls of roughly 10 percent or more (Barro and Liao, 2021), which are more frequent than a fall in productivity of 25 percent. The model nonetheless implies that the measure rises and falls with the disaster intensity: equity is a leveraged claim on consumption, so large equity-market falls come from disasters, and the risk of such falls rises with the disaster intensity.

Asset prices. I choose risk aversion, time preference and the dividend loading to match the volatility of equity returns, the equity premium and the mean safe return over 1994Q4–2025Q2 (Table 3). Asset prices discipline these parameters because the compensation

Table 3: Selected moments

Moment	Data	Model
Equity total-return volatility	8.75	8.80
Mean safe return	0.09	0.28
Mean equity excess return	1.81	1.95
Safe-return volatility	0.67	0.19
Safe-return persistence	0.89	0.89
Mean log dividend yield	−4.01	−2.78

Notes: Data cover 1994Q4–2025Q2 (dividend yield: 1995Q3–2025Q2). Means and standard deviations are in percent of quarterly real returns; persistence is the first-order autocorrelation. The safe return in the data is the deflated Treasury-bill return. Model moments average simulated samples of the same length that start from the 1994Q3 state and contain no disasters. The first three moments are calibration targets. The dividend yield is the log of trailing twelve-month real dividends over the ex-dividend price.

households require for bearing risk also determines how far risk lowers the neutral rate. The calibration sets $\gamma = 9.98$, $\rho = 1.5\%$ and $\phi_d = 9.50$, and the model closely matches equity volatility and the equity premium. Risk aversion is close to the value of 10 used by Bansal and Yaron (2004).

The policy rule. The remaining parameters of the policy rule are the weight on the neutral rate, ω; the speed at which the policy rate adjusts towards its target, κ_i; the rule's constant estimate of the neutral rate, $\bar{r}$; and the volatility of policy shocks, σ_i. I estimate them, with the standard deviations of the measurement errors, by maximum likelihood on quarterly GDP growth, the CBO output gap, GDP-deflator inflation, the federal funds rate and the option-based disaster probability over 1986Q1–2026Q2. The model has no lower bound on the policy rate, so I treat the 36 quarters in which the funds rate was at or below 25 basis points as censored observations. Because the model is nonlinear, I use a nonlinear filter: each quarter, it propagates the full distribution of the states with the model's dynamics, including disaster jumps, and updates it by Bayes' rule with the quarter's data. A backward pass then gives the smoothed distribution of the states given the full sample.

The estimated rule places a weight of $\widehat{\omega} = 0.64$ on the neutral rate, with a 95 percent interval of 0.59 to 0.68, which excludes both a standard Taylor rule and full tracking. The policy rate closes about 74% of the gap to its target within a quarter ($\kappa_i = 5.36$), and policy shocks are small. The intercept, $\bar{r} = 1.02\%$, is close to the mean of the estimated neutral rate since 1995, 0.91%.

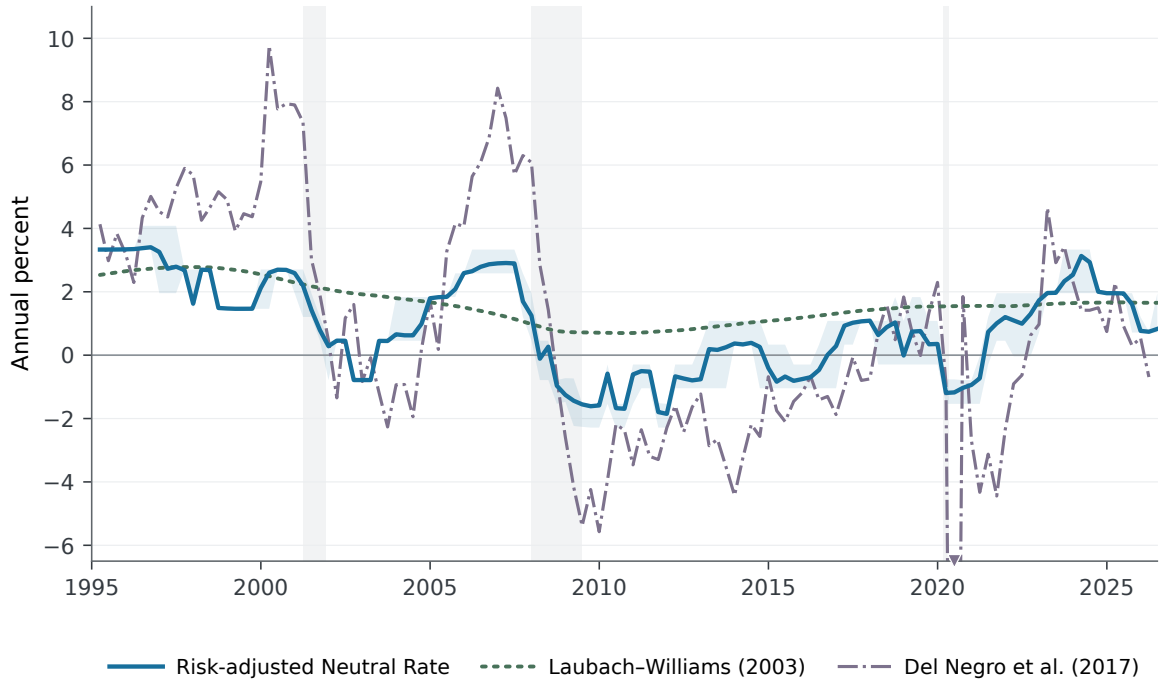


Figure 8: The risk-adjusted neutral rate

Notes: Smoothed estimate of the risk-adjusted neutral rate in the estimated model, the current-vintage two-sided Laubach–Williams estimate, and the May 2026 update of the short-run neutral rate of Del Negro et al. (2017b). The band is the pointwise 90% interval of the smoothed distribution, conditional on the estimated parameters. The Del Negro et al. value of −30.80% in 2020Q2 lies below the plotted range. Recessions are shaded.

5.2 The neutral rate and the stance of policy

The filter recovers the states of the economy from the data, and I evaluate the neutral rate of Proposition 4.1 at the smoothed states. Figure 8 compares this risk-adjusted neutral rate with two benchmarks. The first is the estimate of Laubach and Williams (2003). It infers the neutral rate from the slow co-movement of output, inflation and the real interest rate, so it moves mainly with trend growth and other persistent forces. The second is the short-run neutral rate of Del Negro et al. (2017b), a DSGE estimate that moves at business-cycle frequencies with all the shocks their model needs to fit the macroeconomic data. Like the risk-adjusted rate, it is the real rate in a counterfactual economy without nominal rigidities, here on an asset as safe and liquid as a three-month Treasury bill. In their model, a rise in the convenience yield, the premium investors pay for the safety and liquidity of Treasury securities, lowers this rate, and they attribute most of the decline in interest rates since the late 1990s to a rising convenience yield.

The convenience yield and disaster risk lower the neutral rate for different reasons. The convenience yield is specific to one asset, Treasury securities, whereas disaster risk is a property of the economy as a whole: it raises households' precautionary demand for safe assets in general. In the pandemic, for example, disaster risk rose while strains on dealers' balance sheets turned the convenience yield on Treasuries negative (He et al., 2022). Because the model of Del Negro et al. is linearised, changes in aggregate risk can move its neutral rate only through the convenience yield. The risk-adjusted neutral rate moves with the business cycle. It falls in every recession and stays below the Laubach–Williams estimate for most of the decade after the financial crisis, averaging −0.14% over 2010–2019 against 1.10%. In the crisis, it began to fall before output had fallen much: as the option-based measure of disaster risk rose, the rate fell from 2.90% in 2007Q2 to −0.11% in 2008Q1, when the short-run rate of Del Negro et al. was still 2.85%. By 2008Q4, the risk-adjusted rate had reached −1.25% and the Del Negro et al. rate −2.59%, while the Laubach–Williams estimate was still 0.73%. The risk-adjusted rate reached its lowest value, −1.85%, in 2011Q4. In the pandemic, the Del Negro et al. rate fell to −30.80% in 2020Q2, as their model absorbed the collapse in activity during the lockdown, while the risk-adjusted rate fell to −1.18%.

To assess the stance of policy, I subtract each neutral rate from the same real policy rate: the quarterly mean federal funds rate minus the four-quarter trailing mean of annualised GDP-deflator inflation, which stands in for expected inflation (Figure 9). Because the estimated rule fits the funds rate closely, the gap against the risk-adjusted rate comes mainly from two features of the rule. One is its partial weight on the neutral rate, which leaves policy contractionary when the neutral rate falls below the rule's intercept. The other is its responses to inflation and the output gap.

Measured against the risk-adjusted neutral rate, policy was contractionary at the onset of the financial crisis, by 0.7 to 1.3 percentage points between 2007Q3 and 2008Q1, whereas it was accommodative relative to the Laubach–Williams estimate by mid-2008. The large rate cuts of 2008 then closed the gap: in 2008Q4, the real policy rate of −1.34% was close to the risk-adjusted neutral rate, while the gap relative to the Laubach–Williams estimate was −2.07 percentage points. Policy turned contractionary again in 2009, by 1.0 to 1.7 percentage points, when the funds rate was at zero and inflation fell.

The two gaps diverge when disaster risk rises. In 2020Q2, policy was contractionary relative to the risk-adjusted rate, by 0.51 percentage points, but accommodative relative to the Laubach–Williams estimate, with a gap of −2.21 percentage points. The gaps agree in 2021 and 2022: both show policy as highly accommodative, with inflation rising while the funds rate stayed near zero. In 2021Q4, the gap was −7.2 percentage points against the risk-adjusted rate and −7.5 against the Laubach–Williams estimate. They differ again

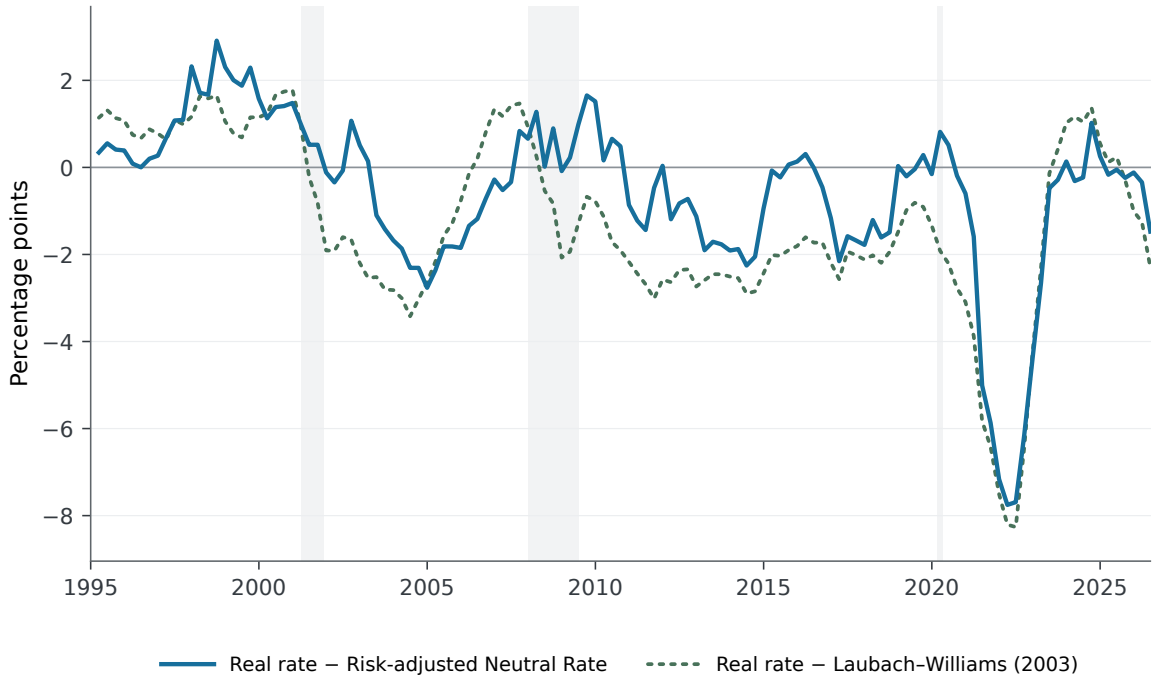


Figure 9: Real-policy-rate gaps

Notes: Each line subtracts the indicated smoothed neutral rate from the same real-rate proxy, the quarterly mean federal funds rate minus four-quarter trailing GDP-deflator inflation. The comparison uses the current-vintage two-sided Laubach–Williams estimate. Positive gaps mean that policy was contractionary relative to the corresponding estimate. Recessions are shaded.

between 2023 and 2025: the gap relative to the risk-adjusted rate was mostly small, while the Laubach–Williams gap shows contractionary policy of up to 1.4 percentage points.

5.3 Volatility versus disaster risk

I now ask whether moments beyond the variance matter for the neutral rate. I compare two permanent increases in risk. The first raises the disaster intensity by one percentage point, holding the size of a disaster fixed. The second raises the volatility of Gaussian productivity shocks from 2.6% to 3.9% a year, an increase of 1.3 percentage points, chosen so that it adds the same variance to productivity growth as the more frequent disasters. Because more frequent disasters also lower expected growth, I also consider a mean-preserving version of the disaster experiment, which raises productivity growth between disasters just enough to keep expected growth unchanged. The mean-preserving increase in disaster risk and the increase in volatility then add the same variance to productivity growth and leave its mean unchanged. Any difference in their effects on the neutral rate therefore

Table 4: The neutral-rate response to higher volatility and disaster intensity

Risk experiment	Log utility	Epstein–Zin
Higher volatility of Gaussian shocks	−4.1	−27.5
Higher disaster intensity	−33.3	−66.9
Higher disaster intensity, expected growth fixed	−4.6	−48.8

Notes: Entries are changes in $r^\star$ in annual basis points at $a^r = \xi = g = 0$; negative values denote declines. The disaster experiments raise the intensity at that state from 3.00% to 4.00% a year, holding the productivity loss at 25.00%. The volatility experiment raises σ_A from 2.60% to 3.88% per square-root year, adding the same instantaneous variance of log productivity. The final row raises productivity growth between disasters to preserve mean log growth. Risk aversion under Epstein–Zin preferences is $\gamma = 9.98$; all other parameters are common across columns.

comes from moments beyond the variance. I report results under Epstein–Zin preferences and under log utility, for which I set risk aversion to one and keep every other parameter, including the policy rule, unchanged.

Even under log utility, the two increases in risk have slightly different effects. Once expected growth is held fixed, disasters lower the neutral rate by 4.6 basis points, against 4.1 for more volatile Gaussian shocks (Table 4). Because households are prudent, the large losses that disasters bring raise their precautionary demand for safe assets by more than a rise in Gaussian volatility that adds the same variance.

Under Epstein–Zin preferences, the gap is much larger. More volatile Gaussian shocks now lower the neutral rate by 27.5 basis points, while more frequent disasters lower it by 66.9 basis points, and by 48.8 basis points once expected growth is held fixed. The mean-preserving increase in disaster risk therefore moves the neutral rate 1.77 times as much as the increase in volatility. Unlike under log utility, households also fear the weaker future that a disaster leaves behind, which raises the weight they place on disaster states.

5.4 News of disaster risk and monetary policy

I now study how news of higher disaster risk affects the economy. When households learn that disasters have become more likely, they want to save more, so the neutral rate falls. I consider an unexpected, one-time disaster-risk shock. It raises the probability of at least one disaster over the next year by one percentage point, from 3.0% at the steady state, and fades with a two-year half-life. Households learn at once how the added risk will evolve, but Gaussian shocks and disasters can still occur. Figure 10 compares the expected responses under three rules: a standard Taylor rule ($\omega = 0$), the estimated rule ($\omega = \widehat{\omega}$) and full tracking ($\omega = 1$).

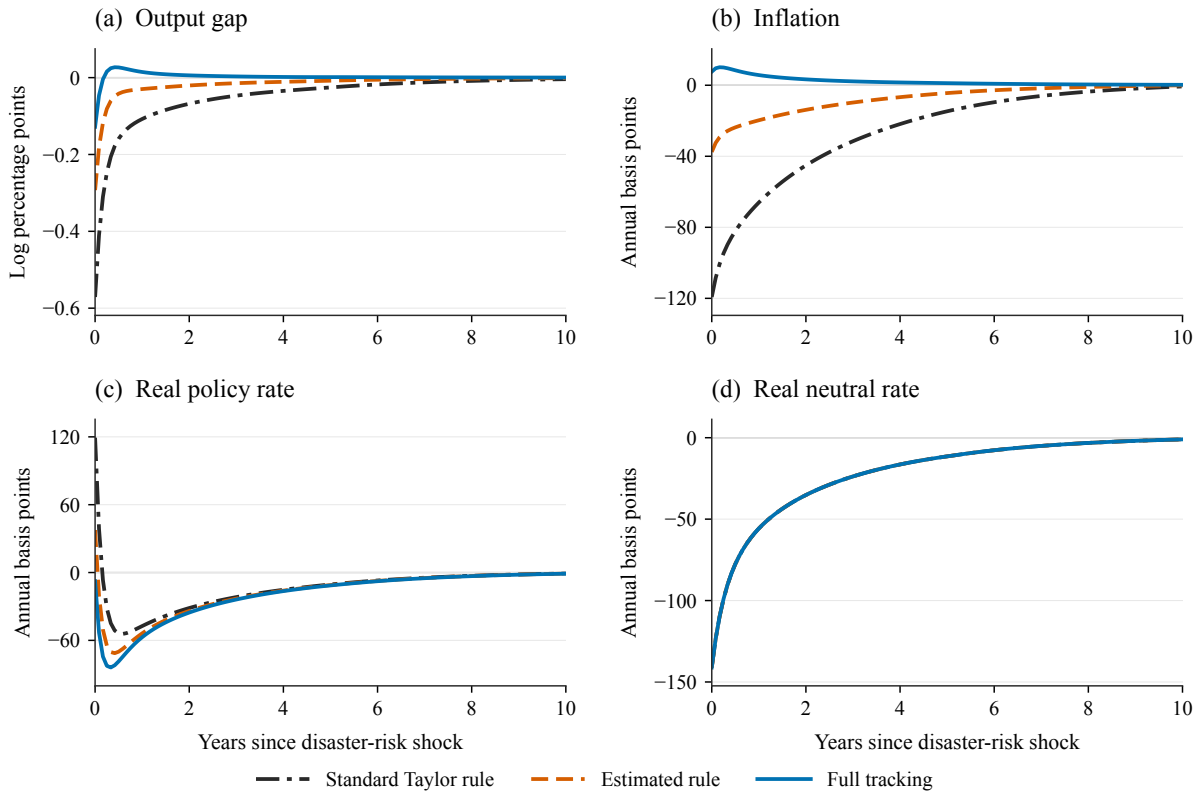


Figure 10: Monetary responses to a disaster-risk shock

Notes: A single unexpected disaster-risk shock raises the one-year disaster probability by one percentage point. The shift in log disaster intensity decays with a two-year half-life; growth between disasters and the volatility of Gaussian shocks remain fixed. All responses average over subsequent productivity and disaster risk and are measured against each rule's own path without the shock. Panel (a) shows changes in expected log output gaps, in log percentage points. Inflation and real rates are in annual basis points. The real neutral-rate responses coincide across rules. Full tracking sets $\omega = 0.99$.

The neutral rate does not depend on policy: it falls by 141 basis points on impact under every rule and is still 56 basis points below its initial level a year later. The policy rate, by contrast, moves only gradually towards the rule's target, so under every rule it is unchanged at the moment the news arrives. Under the standard Taylor rule, inflation falls by 119 basis points and the output gap by 0.57 log percentage points on impact, so the real policy rate rises by 119 basis points just as the neutral rate falls, and the real-rate gap widens by 260 basis points. The rule then cuts the policy rate quickly, through its responses to falling inflation and output: the nominal rate reaches its lowest point, 137 basis points below its initial level, after 5 months. A standard Taylor rule thus responds to the downturn but does not keep pace with the falling neutral rate.

The estimated rule limits the contraction while cutting the policy rate by less. On

impact, its output gap falls by 0.29 log percentage points and inflation by 38 basis points, and its nominal rate falls by at most 96 basis points. The direct response to the neutral rate largely *replaces* the rule's responses to inflation and output rather than adding to them. By limiting the contraction, it reduces the falls in inflation and output that drive those responses.

Finally, full tracking keeps inflation close to target: inflation rises by 8 basis points on impact, and the output gap falls by 0.13 log percentage points before closing within a few months. The output gap still falls on impact because the policy rate cannot jump, but once the rate has adjusted, tracking the neutral rate keeps the output and inflation gaps close to zero, a form of the divine coincidence of Blanchard and Galí (2007). Full tracking achieves this with the smallest cut in the policy rate of the three rules, 74 basis points: by keeping inflation at target, it turns that cut into a larger fall in the real rate. Tracking the neutral rate stabilises the natural allocation in New Keynesian models (King and Wolman, 1999; Woodford, 2003), and the gap between the real and neutral rates is a useful indicator of inflationary pressure (Neiss and Nelson, 2003). A mean-preserving version of the shock, which raises productivity growth between disasters to keep expected growth unchanged, lowers the neutral rate by 96 basis points on impact and produces responses of the same shape.

The onset of the financial crisis. The onset of the Global Financial Crisis in late 2007, almost a year before the collapse of Lehman Brothers, is a period in which disaster risk rose before any disaster had occurred, so it is a natural test of the model's news shock. Figure 11 shows that, from 2007Q4 through 2008Q3, the federal funds rate stayed below the prescription of the original Taylor (1993) rule, which was elevated during 2007 because oil shocks had raised inflation. The option-based measure of disaster risk from Section 3 rose from about 1% in early 2007 to around 10% in late 2007 and early 2008, before output had fallen much.[9] One reading of the Federal Reserve's deviation from the Taylor rule in this period is that it was responding to tail risk, which the model captures through the weight its policy rule places on the neutral rate. The FOMC's statement of 17 August 2007 is consistent with this reading: "although recent data suggest that the economy has continued to expand at a moderate pace, the Federal Open Market Committee judges that the downside risks to growth have increased appreciably" (Federal Open Market Committee, 2007).

The cuts that took the funds rate below the Taylor prescription tracked only part of

[9] Appendix C shows that the pattern holds with alternative measures of financial-market stress, the VIX and the excess bond premium, and Appendix C.3 reports alternative policy rules.

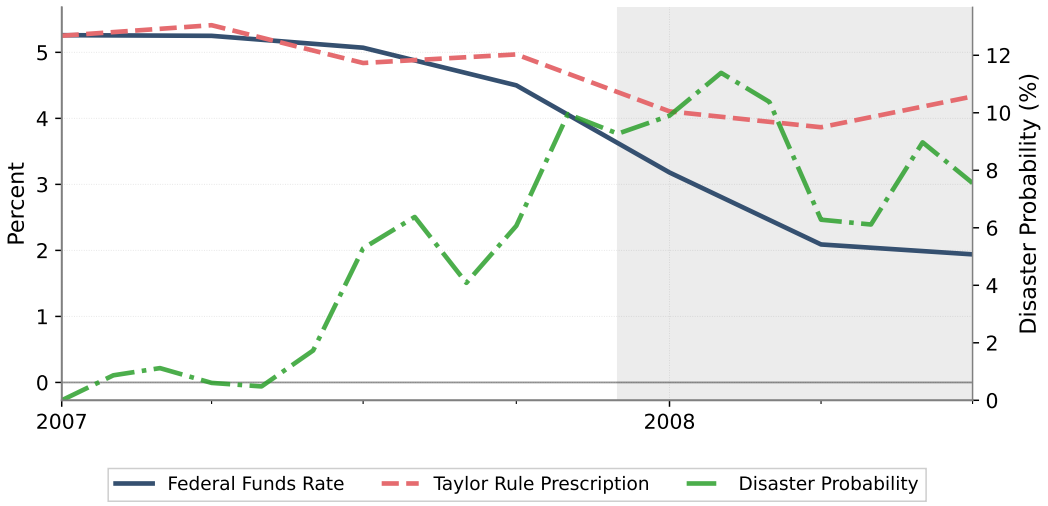


Figure 11: The federal funds rate and the Taylor rule during the financial crisis

Notes: The figure plots the effective federal funds rate, the prescription from the original Taylor (1993) rule, $i_t^{\mathrm{TR}} = 2 + \pi_t^{(4)} + 0.5(\pi_t^{(4)} - 2) + 0.5x_t^{\mathrm{RT}}$, with four-quarter GDP-deflator inflation $\pi_t^{(4)}$ and the real-time output gap x_t^{RT} of Nakamura et al. (2025), and the disaster probability implied by out-of-the-money S&P 500 put options following Barro and Liao (2021). The rate and prescription observations are quarterly, cover 2007Q1–2008Q3, and are dated at the start of each quarter. Disaster probability is monthly, from January 2007 through July 2008. The July-dated rate and prescription observations therefore cover the third quarter. Shading marks the part of the NBER-dated recession (December 2007–June 2009) within the plotted window.

the fall in the risk-adjusted neutral rate, so policy remained contractionary relative to that rate until 2008Q1 (Figure 9). The Federal Reserve cut by more than the Taylor (1993) rule prescribed, but by less than the neutral rate fell.

To see how much of the early crisis a disaster-risk shock can account for, I fit the size and date of a single shock like the one above, holding every parameter at its estimate. The fit targets the changes in private output and the federal funds rate from 2007Q2 to each of the next three quarters, which I compare with quarterly averages of the model's expected responses under the estimated rule (Figure 12). The shock raises log disaster intensity from a date in 2007Q3 and then decays with a two-year half-life. The best fit dates the shock about 2 months into 2007Q3, in August 2007, when interbank markets first seized up, and sets its size at a shift of 0.80 in log disaster intensity. The estimated rule matches the timing of the rate cuts: by 2008Q1, the change in the model's funds rate is −1.42 percentage points, against −2.07 in the data. Output falls by less than in the data, −0.29 against −1.55 log percentage points, so a single disaster-risk shock accounts for only part of the early contraction. Financial frictions, which the model leaves out, may have amplified the crisis further (Bernanke et al., 1999).

To measure how much the Federal Reserve's response to disaster risk mattered, I solve

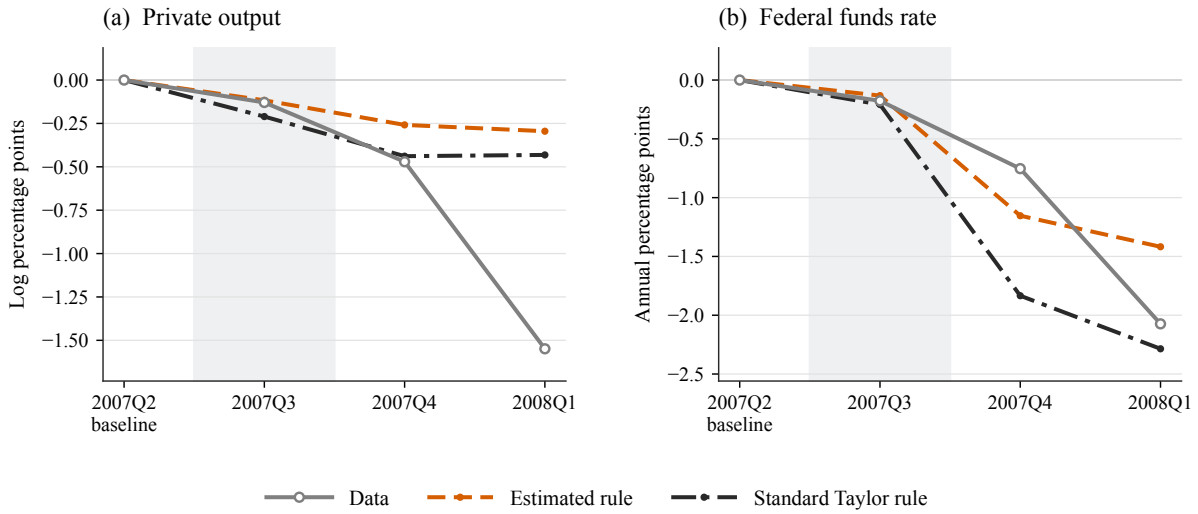


Figure 12: A single disaster-risk shock and the onset of the financial crisis

Notes: Data are changes from 2007Q2 in detrended private output, $\log(C + I)$, in log percentage points, and in the federal funds rate, in annual percentage points. The model shows quarterly averages of the expected changes in log output and the policy rate after a single disaster-risk shock whose size and date are fitted to both series under the estimated rule, and the same shock under a standard Taylor rule ($\omega = 0$). The zero values in 2007Q2 mark the normalisation, not an additional fitting observation. Shading marks the quarter of the shock.

the model again with the same shock, from the same 2007Q2 state, under a standard Taylor rule that ignores the neutral rate ($\omega = 0$). The rule keeps the estimated intercept and response coefficients and drops only the weight on the neutral rate, so the comparison isolates what tracking the neutral rate contributed. The rule is not the real-time prescription of Figure 11. By 2008Q1, the change in output is then −0.43 log percentage points, against −0.29 under the estimated rule, a fall about 46% larger. The standard rule cuts the funds rate by more, with a change of −2.29 against −1.42 percentage points, because it responds to the larger falls in inflation and output that result from ignoring the neutral rate. As after the news shock, a rule that ignores disaster risk ends up cutting by more, in response to a contraction that the estimated rule partly avoids. On this reading, the Federal Reserve's early cuts below the Taylor prescription limited the first stage of the downturn, and the responses to the news shock suggest that fuller tracking would have limited it further.

5.5 The welfare cost of failing to track the neutral rate

I now measure what households lose from the fluctuations in output and inflation that arise when monetary policy does not track the neutral rate. I solve the economy twice, once with flexible prices, which gives the natural allocation with output at potential and

inflation at target, and once with sticky prices under a given monetary rule. In each case, I compute households' lifetime utility, which accounts for consumption and hours worked at every future date and for the risk to productivity, including disasters. Following Lucas (2003), I express the gap in lifetime utility as the share of consumption that households would be willing to give up permanently to move from the allocation under the rule to the natural allocation. Because the natural allocation is efficient, the loss measures the full cost of the rule's departures from it.

For small output and inflation gaps, the loss in each period is approximately

$$\frac{1+\varphi}{2}x_t^2 + \frac{\widehat{\pi}_t^2}{2\eta}.$$

The first term arises because output gaps distort the choice between consumption and leisure, and the second because inflation deviations use up goods in price adjustment. The lifetime loss adds up these flow losses over all future dates and states and, under Epstein–Zin preferences, gives extra weight to bad outcomes.

Table 5 reports the exact loss under the standard Taylor rule, the estimated rule and full tracking. Each row fixes the state of the economy. From that state, I compute households' lifetime utility in the natural allocation and under each rule, and both face the same future shocks. The loss in a row is therefore the cost of the rule from that state. Because the natural allocation also changes with the state, differences between rows do not measure the cost of the shock itself. The first row starts from the steady state with no news, so lifetime utility averages over all the ordinary productivity shocks, movements in disaster risk and disasters that may occur in the future. Under the standard Taylor rule, households would give up 0.236% of their consumption every year to move to the natural allocation; under the estimated rule, they would give up 0.032%. These losses arise because the neutral rate changes over time with productivity and disaster risk, while the standard rule does not respond to these changes directly and the estimated rule responds only partly. Full tracking comes close to the natural allocation, with a loss of only 0.001%, although households still face the same productivity risk. The small remaining loss comes from the gradual adjustment of the policy rate, from policy shocks and from setting ω at 0.99 rather than one.

These losses depend on the slope of the Phillips curve. The calibration implies a slope of 0.17 in quarterly units, while more recent estimates point to flatter curves (Hazell et al., 2022; Del Negro et al., 2020). A flatter curve makes each deviation of inflation from target more costly, so the cost of imperfect tracking rises: with a slope of 0.057, about 3 times flatter, the loss under the standard Taylor rule rises to 0.567% of consumption and under

Table 5: Welfare losses under alternative monetary rules

Evaluation	Standard Taylor rule	Estimated rule	Full tracking
Baseline	0.236	0.032	0.001
Flatter Phillips curve	0.567	0.081	0.001
Disaster-risk shock	0.238	0.032	0.001
Disaster-risk shock, mean-preserving	0.238	0.032	0.001
After a realised disaster	0.265	0.036	0.000

Notes: Losses are percentages of natural consumption under Epstein–Zin preferences, $100[1-\exp(v_\omega - v_n)]$. The estimated rule uses $\widehat{\omega} = 0.639$, and full tracking uses $\omega = 0.99$. Within each row, all rules face the same risks and start from the same state. The second row lowers the slope of the Phillips curve to 0.057 per quarter, about 3 times flatter than the baseline and closer to recent estimates (Hazell et al., 2022). The first four rows start from $a^r = \xi = g = 0$ with the policy rate at $\pi^\star + \bar{r}$, and the last row from $a^r = J_A$.

the estimated rule to 0.081%, while full tracking still costs 0.001% (second row of Table 5). The losses under the calibrated slope are therefore, if anything, on the low side.

The third row also starts from the steady state, but at that moment households learn of the news shock, which raises the probability of a disaster over the next year by one percentage point and fades with a two-year half-life. I solve for the path of the economy after the news, both with flexible prices and with sticky prices under each rule, and evaluate lifetime utility when the news arrives, before the policy rate can respond. Ordinary shocks and disasters can still occur, and once the added risk has faded the economy follows its usual dynamics again. The standard Taylor rule costs 0.238% of consumption, and the estimated rule 0.032%. The fourth row repeats the exercise with the mean-preserving shock. Because that shock also raises productivity growth between disasters to keep expected growth unchanged, it isolates the effect of higher risk. The losses are similar, 0.238% and 0.032%. In both rows, the loss is close to its steady-state value, because the shock fades within a few years and the loss measures the cost of each rule over the whole future.

The fifth row starts immediately after a disaster has struck the economy at its steady state. Productivity is 25% below its permanent component in both the natural and the sticky-price economies, the disaster-risk state and trend growth are unchanged, and the policy rate has not yet moved. Because disasters are more likely when productivity is low, the disaster intensity rises from 3.0% to 4.55% a year. The standard Taylor rule then costs 0.265% of consumption and the estimated rule 0.036%: the neutral rate rises over the following years, as productivity is expected to recover, and a rule that does not track it lets inflation rise. In every row, the loss falls as the rule tracks the neutral rate more closely.

5.6 A semi-structural estimate of the neutral rate

The estimate of Section 5.2 rests on the full structure of the model: its preferences, its disaster process and its policy rule. I also estimate the risk-adjusted neutral rate in a state-space model that builds on Laubach and Williams (2003), as a check that does not impose this structure. As in their model, the neutral rate is identified from an IS curve, in which a real rate above neutral depresses demand, and a Phillips curve, in which inflation falls when output is below potential. I let the neutral rate also respond to the option-based disaster probability, which can rise abruptly when the outlook deteriorates and so lower the neutral rate even when trend growth moves little.

The risk-adjusted neutral rate is

$$r_t^\star = \alpha_L + g_t^\star + z_t - \theta_\lambda \widetilde{\lambda}_t. \tag{24}$$

Trend growth of potential output, $g_t^\star$, passes one for one into the rate, z_t captures other persistent determinants and α_L is a constant. The last term lets disaster risk lower the rate, as in Sections 2 and 4. Here $\widetilde{\lambda}_t$ is the option-based disaster probability of Section 3, a noisy counterpart of the disaster intensity.

The semi-structural model combines slow-moving macroeconomic data with fast-moving option prices. The trend, $\alpha_L + g_t^\star + z_t$, is identified from output, potential output and inflation, as in Laubach and Williams (2003). These series reveal a change in the neutral rate only gradually, as the gap between the real rate and the neutral rate shows up first in output and then in inflation, so the trend moves slowly. The disaster-risk term instead comes from option prices, which respond within days to news about the risk of a severe contraction.

I estimate the loading θ_λ together with the shock and measurement-error variances by maximum likelihood, using output, CBO potential output, changes in inflation and the disaster probability over 1986Q1–2026Q2, and report smoothed estimates from the Kalman filter. The weight on the financial-market signal is therefore estimated rather than imposed, and the loading measures how much disaster risk helps explain output and inflation. Its estimate is 24.4 basis points per percentage point of the disaster probability, with a 90 percent bootstrap interval of 15.4 to 32.3 basis points, and the data reject a zero loading. In the estimated model, the same rise in the option-based probability lowers the neutral rate by 27.6 basis points,[10] which lies inside this interval. Imposing the model's value also barely changes the fit of the semi-structural model, so the two approaches agree

[10]The value is the model's average slope as the option-based probability rises above its mean. At the steady state, the slope is 23.9 basis points.

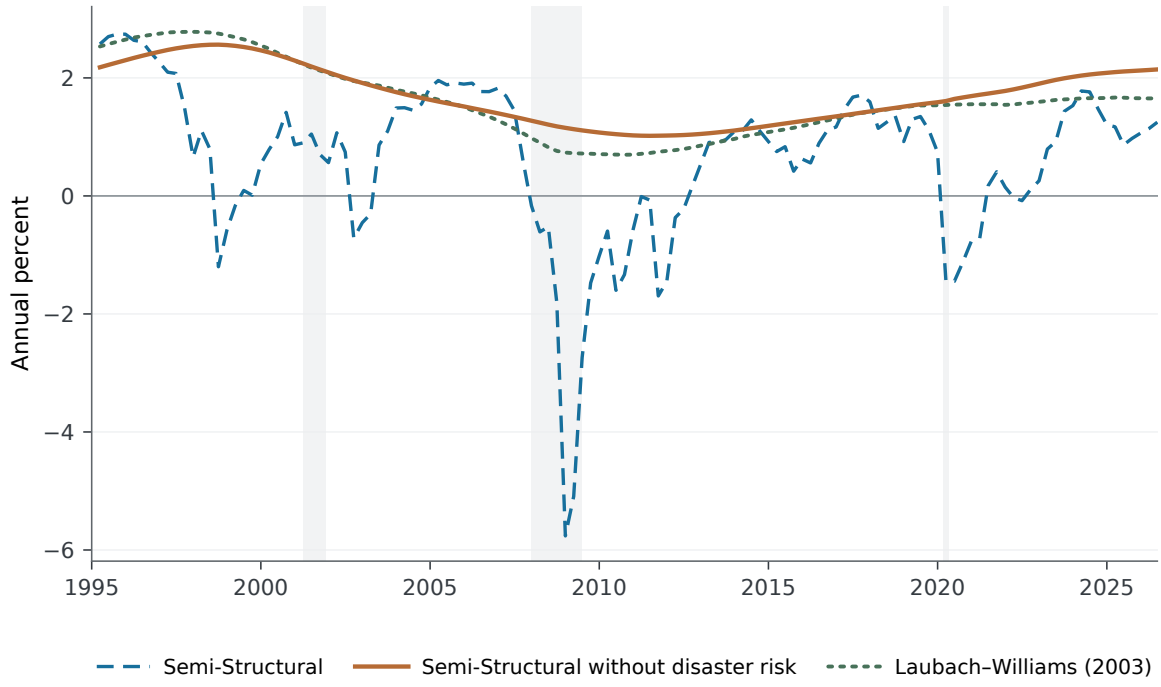


Figure 13: The semi-structural estimate with and without disaster risk

Notes: Neutral rates, in annual percent. The semi-structural estimate is the smoothed estimate of (24); the estimate without disaster risk sets $\theta_\lambda = 0$ in that equation at the same smoothed states, $\alpha_L + g_t^\star + z_t$. The Laubach–Williams estimate is the current-vintage two-sided estimate of Figure 8. Recessions are shaded.

on how strongly disaster risk moves the neutral rate.

Figure 13 plots the two components separately. Without the disaster-risk term, the estimate is the smooth trend, which closely follows the Laubach–Williams estimate, with a correlation of 0.96 since 1995. The disaster-risk term adds the business-cycle movements: it pulls the estimate well below the trend when perceived risk rises, in 1998, around 2001–2002, from 2008 to 2012 and from 2020 to 2023, and the estimate returns towards the trend as risk subsides. In 2008Q4, the trend had changed little, while the full estimate fell to −5.76%. Estimates built from macroeconomic data alone therefore register such falls in the neutral rate late, if at all.

Figure 14 compares the semi-structural estimate with the model's. They move together, with a correlation of 0.73 since 1995, and both fall in every recession. They differ most in the financial crisis: the semi-structural estimate falls to −5.76% in 2008Q4, against −1.25% for the model. In 2020Q2, the two estimates are closer, at −1.45% and −1.18%. Measured against the semi-structural estimate, policy was contractionary by 4.42 percentage points in 2008Q4, while it was close to neutral against the model's estimate. Appendix D.5

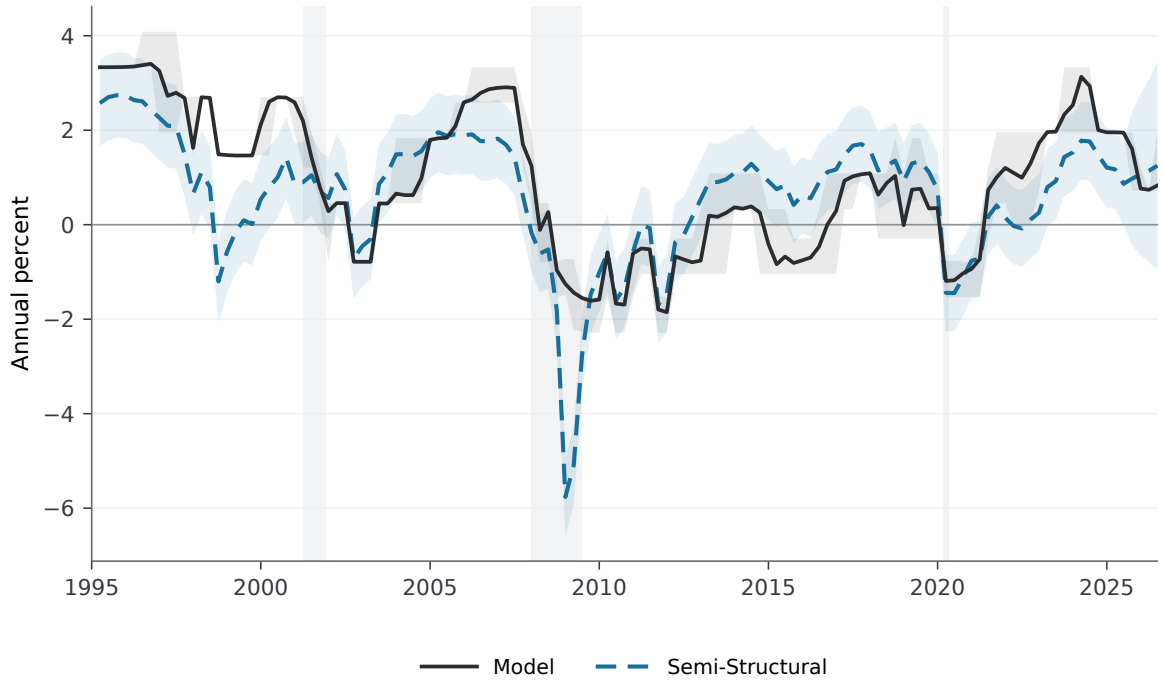


Figure 14: Model-based and semi-structural estimates of the neutral rate

Notes: Neutral rates, in annual percent. The model-based estimate is that of Figure 8; the semi-structural estimate is the smoothed estimate of (24). Bands are pointwise 90% intervals conditional on the estimated parameters. Recessions are shaded.

explains why the two estimates differ in the financial crisis and compares both with other published estimates of the short-run and long-run neutral rate and with measures of the policy setting.

6 Conclusion

Monetary policy is often described as navigating by the stars, whereby the neutral rate of interest $r_t^\star$ is the target for the stance of monetary policy. I have argued that this star moves with the perceived probability of a disaster, so central banks need a benchmark that moves with it. To this end, I have proposed the risk-adjusted neutral rate as the benchmark for the stance of policy: the real rate that keeps inflation at target and output at potential once *risk* is taken into account. When households perceive a higher probability of disaster, asset prices fall, which lowers current spending through a wealth effect. The risk-adjusted neutral rate therefore falls, even before any disaster strikes. Because disasters are tail events, the risk-adjusted neutral rate depends not only on the variance of output growth but also on its higher-order moments, including skewness and kurtosis, which drop out of the first-order approximations used in standard New Keynesian models. I analytically show that a central bank that responds only to the second moment of output growth, ignoring its higher-order moments, leaves large welfare gains unrealised, because its policy rate stays above the risk-adjusted neutral rate. I also show that, for a wide range of distributions, tracking the moments up to the fourth is sufficient to come close to optimal policy and to stabilise the economy against tail risk.

Using option prices to measure disaster risk, I find that, as the mechanism suggests, disaster-risk shocks appear in the data as negative demand shocks: they lower consumption, investment, output and inflation together. They are also an important driver of the business cycle. To measure the size of these effects, I build a quantitative New Keynesian model and estimate it on US data. I find that higher-order moments matter quantitatively: holding the mean and variance of productivity growth fixed, a rise in disaster risk lowers the neutral rate 1.77 times as much as a rise in the volatility of Gaussian shocks. I also estimate how far the Federal Reserve tracks the risk-adjusted neutral rate and find that it does so only in part. Using the model, I then run a policy counterfactual experiment for late 2007. Had the Federal Reserve followed a standard Taylor rule, which ignores movements in the neutral rate, the same rise in disaster risk would have lowered output by early 2008 about 46 percent more than under its estimated rule, which tracks them in part. I also measure the welfare gains from tracking disaster risk: households would be willing to give up 0.236% of their consumption permanently for the central bank to track disaster risk fully, keeping output at potential and inflation at target, rather than follow a standard Taylor rule.

Finally, to address concerns about model misspecification, I also estimate the risk-adjusted neutral rate with a semi-structural model. The two approaches agree on how strongly disaster risk moves the neutral rate, and both estimates drop sharply when risk

spikes, falling below zero in the financial crisis while standard estimates stay positive. Measured against the model's estimate, policy was contractionary at the onset of the crisis, before the large rate cuts of 2008. Slow-moving estimates register falls in the neutral rate late, if at all, so a central bank that relies on them can tighten policy without intending to, even while it cuts rates.

Two extensions are left for future work. First, disaster risk is exogenous in my model, but monetary policy may itself move perceived risk and risk premia (Kekre and Lenel, 2022; Caramp et al., 2026). A model in which policy affects the perceived probability of a disaster would show whether tracking the neutral rate also dampens the risk that moves it. Second, the risk-adjusted neutral rate can fall well below zero, as it did in late 2008, when the policy rate cannot follow. How forward guidance and asset purchases can track it at the lower bound (Dordal i Carreras and Lee, 2026) is a natural next question. More broadly, the effects of higher-order moments on monetary economics are a rich area for further research, and the results of this paper suggest that their policy implications are large.

References

Andersen, T. G., N. Fusari, and V. Todorov (2015): "The Risk Premia Embedded in Index Options," *Journal of Financial Economics*, 117, 558–584.

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

Backus, D. K., M. Chernov, and S. E. Zin (2014): "Sources of Entropy in Representative Agent Models," *Journal of Finance*, 69, 51–99.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *The Journal of Finance*, 59, 1481–1509.

Barro, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *The Quarterly Journal of Economics*, 121, 823–866.

——— (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264.

Barro, R. J. and R. G. King (1984): "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *The Quarterly Journal of Economics*, 99, 817–839.

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

Barro, R. J. and J. F. Ursúa (2008): "Macroeconomic Crises since 1870," *Brookings Papers on Economic Activity*, 39, 255–350.

Barsky, R., A. Justiniano, and L. Melosi (2014): "The Natural Rate of Interest and Its Usefulness for Monetary Policy," *American Economic Review*, 104, 37–43.

Basu, S. and B. Bundick (2017): "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 85, 937–958.

Basu, S., G. Candian, R. Chahrour, and R. Valchev (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.

Beaudry, P., P. Cavallino, and T. Willems (2025): "Monetary Policy along the Yield Curve: Why Can Central Banks Affect Long-Term Real Rates?" Staff Working Paper 1117, Bank of England.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Amsterdam: Elsevier, vol. 1C, chap. 21, 1341–1393.

Bianchi, F., M. Lettau, and S. C. Ludvigson (2022): "Monetary Policy and Asset Valuation," *Journal of Finance*, 77, 967–1017.

Blanchard, O. and J. Galí (2007): "Real Wage Rigidities and the New Keynesian Model,"

Journal of Money, Credit and Banking, 39, 35–65.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Breeden, D. T. and R. H. Litzenberger (1978): "Prices of State-contingent Claims Implicit in Option Prices," *The Journal of Business*, 51, 621–51.

Brunnermeier, M., D. Palia, K. A. Sastry, and C. A. Sims (2021): "Feedbacks: Financial Markets and Economic Activity," *American Economic Review*, 111, 1845–79.

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

Caballero, R., E. Farhi, and P.-O. Gourinchas (2017): "The Safe Assets Shortage Conundrum," *Journal of Economic Perspectives*, 31, 29–46.

Caballero, R. J. and A. Simsek (2020): "A Risk-Centric Model of Demand Recessions and Speculation," *The Quarterly Journal of Economics*, 135, 1493–1566.

Caramp, N., A. Sarto, and D. H. Silva (2026): "Managing a Liquidity Crisis: Monetary Policy with Convenience Yields and Risk Premia," Working paper, version dated October 19, 2026.

Cerra, V. and S. C. Saxena (2008): "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, 98, 439–457.

Chodorow-Reich, G., P. T. Nenov, and A. Simsek (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 111, 1613–1657.

Christensen, J. H. E. and G. D. Rudebusch (2019): "A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt," *The Review of Economics and Statistics*, 101, 933–949.

Clarida, R., J. Galí, and M. Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, 115, 147–180.

Cochrane, J. H. (2011): "Discount Rates," *Journal of Finance*, 66, 1047–1108.

Coibion, O. and Y. Gorodnichenko (2012): "Why Are Target Interest Rate Changes So Persistent?" *American Economic Journal: Macroeconomics*, 4, 126–162.

Cúrdia, V., A. Ferrero, G. C. Ng, and A. Tambalotti (2015): "Has U.S. Monetary Policy Tracked the Efficient Interest Rate?" *Journal of Monetary Economics*, 70, 72–83.

Del Negro, M., G. Eggertsson, A. Ferrero, and N. Kiyotaki (2017a): "The Great Escape? A Quantitative Evaluation of the Fed's Liquidity Facilities," *American Economic Review*, 107, 824–857.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017b): "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, spring

2017.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020): "What's Up with the Phillips Curve?" *Brookings Papers on Economic Activity*, 301–357, spring 2020.

Dordal i Carreras, M. and S. J. Lee (2026): "Higher-Order Forward Guidance," *Journal of Economic Theory*, 231, 106108.

Dou, L., P. Ho, and T. A. Lubik (2025): "Max-Share Misidentification," Working Paper 25-02, Federal Reserve Bank of Richmond.

Duffie, D. and L. G. Epstein (1992): "Stochastic Differential Utility," *Econometrica*, 60, 353–394.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Federal Open Market Committee (2007): "FOMC Statement," Press release, August 17, Board of Governors of the Federal Reserve System, https://www.federalreserve.gov/newsevents/pressreleases/monetary20070817b.htm.

Fernald, J. G. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Gabaix, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *The Quarterly Journal of Economics*, 127, 645–700.

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

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

Gourio, F. (2008): "Disasters and Recoveries," *American Economic Review*, 98, 68–73.

——— (2012): "Disaster Risk and Business Cycles," *American Economic Review*, 102, 2734–2766.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022): "The Slope of the Phillips Curve: Evidence from U.S. States," *Quarterly Journal of Economics*, 137, 1299–1344.

He, Z., S. Nagel, and Z. Song (2022): "Treasury Inconvenience Yields during the COVID-19 Crisis," *Journal of Financial Economics*, 143, 57–79.

Holston, K., T. Laubach, and J. C. Williams (2017): "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, S59–S75.

——— (2023): "Measuring the Natural Rate of Interest after COVID-19," Staff Reports 1063, Federal Reserve Bank of New York.

Isoré, M. and U. Szczerbowicz (2017): "Disaster Risk and Preference Shifts in a New

Keynesian Model," *Journal of Economic Dynamics and Control*, 79, 97–125.

Kekre, R. and M. Lenel (2022): "Monetary Policy, Redistribution, and Risk Premia," *Econometrica*, 90, 2249–2282.

Khorrami, P. and F. Mendo (2025): "Fear, Indeterminacy, and Policy Responses," Working paper, version dated May 18, 2025.

King, R. G. and A. L. Wolman (1999): "What Should the Monetary Authority Do When Prices Are Sticky?" in *Monetary Policy Rules*, ed. by J. B. Taylor, Chicago: University of Chicago Press, 349–404.

Kozeniauskas, N., A. Orlik, and L. Veldkamp (2018): "What Are Uncertainty Shocks?" *Journal of Monetary Economics*, 100, 1–15.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2019): "The Tail That Keeps the Riskless Rate Low," *NBER Macroeconomics Annual*, 33, 253–283.

——— (2020): "The Tail That Wags the Economy: Beliefs and Persistent Stagnation," *Journal of Political Economy*, 128, 2839–2879.

Krishnamurthy, A. and A. Vissing-Jorgensen (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120, 233–267.

Laubach, T. and J. C. Williams (2003): "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, 85, 1063–1070.

Lee, S. J. and M. Dordal i Carreras (2025): "Self-fulfilling Volatility, Taylor Rules, and Equilibrium Selection," Conditionally accepted, *Journal of the European Economic Association*.

Lubik, T. A. and C. Matthes (2015): "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches," Economic Brief 15-10, Federal Reserve Bank of Richmond.

——— (2023): "The Stars Our Destination: An Update for Our R* Model," Economic Brief 23-32, Federal Reserve Bank of Richmond.

Lucas, Robert E., J. (2003): "Macroeconomic Priorities," *American Economic Review*, 93, 1–14.

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

——— (2013): "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80, 745–773.

——— (2017): "What is the Expected Return on the Market?" *The Quarterly Journal of Economics*, 132, 367–433.

Martin, I. W. (2008): "Disasters and the Welfare Cost of Uncertainty," *American Economic Review*, 98, 74–78.

Müller, K., C. Xu, M. Lehbib, and Z. Chen (2025): "The Global Macro Database: A New International Macroeconomic Dataset," Working Paper 33714, National Bureau of Economic Research.

Nakamura, E., V. Riblier, and J. Steinsson (2025): "Beyond the Taylor Rule," Working Paper 34200, National Bureau of Economic Research.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Nakamura, E., J. Steinsson, R. Barro, and J. Ursúa (2013): "Crises and Recoveries in an Empirical Model of Consumption Disasters," *American Economic Journal: Macroeconomics*, 5, 35–74.

Neiss, K. S. and E. Nelson (2003): "The Real-Interest-Rate Gap as an Inflation Indicator," *Macroeconomic Dynamics*, 7, 239–262.

Obstfeld, M. (2025): "Natural and Neutral Real Interest Rates: Past and Future," *IMF Economic Review*, 73, 339–392.

Orlik, A. and L. Veldkamp (2024): "Understanding uncertainty shocks and the role of black swans," *Journal of Economic Theory*, 222, 105905.

Pflueger, C., E. Siriwardane, and A. Sunderam (2020): "Financial Market Risk Perceptions and the Macroeconomy," *The Quarterly Journal of Economics*, 135, 1443–1491.

Powell, J. H. (2018): "Monetary Policy in a Changing Economy," Speech at "Changing Market Structure and Implications for Monetary Policy," Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, Wyoming, August 24, Board of Governors of the Federal Reserve System.

Rotemberg, J. J. (1982): "Monopolistic Price Adjustment and Aggregate Output," *The Review of Economic Studies*, 49, 517–531.

Rotemberg, J. J. and M. Woodford (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," in *NBER Macroeconomics Annual 1997*, ed. by B. S. Bernanke and J. J. Rotemberg, Cambridge, MA: MIT Press, vol. 12, 297–346.

Seo, S. B. and J. A. Wachter (2019): "Option Prices in a Model with Stochastic Disaster Risk," *Management Science*, 65, 3449–3469.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

Tallarini, T. D. (2000): "Risk-sensitive real business cycles," *Journal of Monetary Economics*, 45, 507–532.

Taylor, J. B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Uhlig, H. (2003): "What Moves Real GNP?" Unpublished manuscript / working paper

(mimeo), Humboldt University Berlin. Preliminary draft (first draft dated March 8).

Wachter, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *The Journal of Finance*, 68, 987–1035.

Wicksell, K. (1898): *Geldzins und Güterpreise: Eine Studie über die den Tauschwert des Geldes bestimmenden Ursachen*, Jena: Gustav Fischer, English translation by R. F. Kahn, *Interest and Prices: A Study of the Causes Regulating the Value of Money*, London: Macmillan, 1936.

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

Wu, J. C. and F. D. Xia (2016): "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48, 253–291.

A Empirical Analysis: Data and Robustness

This appendix documents the data and reports robustness checks for the empirical analysis in Section 3. It also reports historical policy-rule comparisons using the risk-adjusted neutral rate estimated in Section 5.6.

A.1 Data construction

A.1.1 The option-based measure

The monthly disaster-probability series retains the published Barro and Liao (2021) estimates through March 2021 and extends them using OptionMetrics S&P 500 put data through August 2025. S&P 500 index options trade in one of the most liquid option markets, and their prices are widely used to infer the risk-neutral distribution of future market returns, including the probability of a market crash (e.g. Breeden and Litzenberger, 1978; Andersen et al., 2015; Martin, 2017). For the local estimate, I select end-of-month contracts nearest to moneyness targets of 0.5, 0.6, 0.7, 0.8 and 0.9, retaining actual moneyness no greater than 0.9. Target maturities are 30, 60, 90 and 180 days, with a seven-day tolerance and admissible maturities of 16–183 days. I reconstruct put prices from implied volatilities using the Black–Scholes formula with interest and dividend rates set to zero.

Under power-law disaster sizes with tail parameter α and CRRA utility with coefficient γ, the relative option price $\Omega = P/(S \cdot T)$ satisfies $\Omega \approx \eta_1 \, p_t \, \epsilon^{1+\alpha-\gamma}$, where P is the put price, S the index level, T the maturity, ϵ the moneyness (the strike relative to S) and η_1 a collection of structural constants. The corresponding theoretical log representation is

$$\log \Omega_{i,t} = \beta_0 + \beta_1 \log \epsilon_i + \mu_t + e_{i,t},$$

where $\mu_t = \log p_t$ is the log of the time-varying disaster probability and $\beta_1 = 1 + \alpha - \gamma$. The monthly effects μ_t therefore trace out movements in p_t, while the moneyness slope β_1 is the same in every month. In practice, I fit relative option-price levels by least squares, allowing for maturity and moneyness elasticities, nonnegative monthly effects and an additional common moneyness term. The local annual disaster probability is the estimated monthly effect divided by η_1. Risk aversion is fixed at $\gamma = 3$ and the disaster threshold at 1.1, so a disaster is a fall in the index to at most 1/1.1 of its level, a decline of at least about 9 percent. The estimated tail parameter is $\alpha = 7.1$.

To splice the published and locally estimated series, I forward-fill both to daily frequency and take monthly means. I then calibrate the local series with an OLS regression of the published series' logit on a constant and the local series' logit over their 301 overlapping

months. The slope is 0.37 and the R^2 is 0.46. The spliced series keeps the published monthly observations from September 1994 to March 2021. From April 2021 to August 2025, it uses the calibrated local estimates, transformed back into probabilities. For the VAR analysis, I average the monthly disaster probabilities within each quarter.

A.1.2 Macroeconomic variables

Real GDP, consumption and investment are from the Bureau of Economic Analysis national accounts (seasonally adjusted; nominal consumption and investment aggregates deflated by the GDP deflator) and enter the max-share VAR as per-capita log levels. The population series is an HP-filtered trend of the civilian noninstitutional population. For employment, I use total nonfarm payrolls from the Bureau of Labor Statistics, again in per-capita logs. Real stock and Treasury-bill returns are the CRSP value-weighted market return and the CRSP 90-day Treasury bill return, both deflated by the GDP deflator. I build the log dividend–price ratio from the CRSP returns with and without distributions. The 5-year Treasury yield averages the daily constant-maturity yield from the Federal Reserve H.15 release over each quarter. The VIX is a one-month volatility index, integrated from S&P 500 option quotes following Martin (2011) and computed from the same option panel as the disaster probability. It tracks the CBOE index closely.

A.2 VAR specification and identification

A.2.1 Baseline max-share identification

The baseline VAR is equation (12), with the eleven variables listed in Section 3.3, a constant and three lags, estimated on quarterly US data over 1996Q1–2019Q4. Appendix A.2.4 extends the sample through 2025Q2.

The structural shocks ε_t relate to reduced-form residuals via $u_t = A\varepsilon_t$, where A is the impact matrix. Identification is in the frequency domain, following Angeletos et al. (2020): the disaster-risk shock is the unit-norm impact vector a that maximises the share of the target's forecast-error variance attributable to a frequency band Ω,

$$a^\star = \arg\max_{\|a\|=1} \frac{\int_\Omega \left| e_1' B(e^{-i\omega})\, a \right|^2 d\omega}{\int_\Omega e_1' B(e^{-i\omega}) B(e^{i\omega})' e_1 \, d\omega},$$

where e_1 selects the disaster probability and $B(e^{-i\omega})$ is the moving-average lag polynomial of the VAR, post-multiplied by a square root of Σ_u and evaluated on the unit circle. I integrate the spectral density numerically over the *near-unconditional* band, $\Omega = [2\pi/500, 2\pi/2]$,

which covers cycles with periods between 2 and 500 quarters. I normalise the sign of the shock so that it raises the disaster probability and lowers stock returns on impact. I report its forecast-error variance shares at this band and at the business-cycle band (periods 6–32 quarters).

The confidence intervals are 80% bands, from the 10th to the 90th percentile of 600 residual-based bootstrap replications. Each replication re-estimates the VAR on a bootstrapped sample, re-applies the frequency-domain max-share identification and computes the impulse responses.

A.2.2 Controlling for productivity shocks

To check that the max-share shock does not combine disaster risk with news about productivity, I adapt the procedure of Basu et al. (2021). I first identify a productivity shock in the baseline VAR and then maximise the share of disaster-probability fluctuations explained by a shock orthogonal to it. The VAR keeps its eleven variables, three lags and 1996Q1–2019Q4 sample, and productivity enters only as an auxiliary series. I regress quarterly utilisation-adjusted productivity growth (Fernald, 2014) on the VAR's lags and its contemporaneous innovations, and the fitted combination of innovations is the productivity surprise. The first control removes this surprise, and the second a supply shock. The supply shock is the unit-length combination of innovations whose impact effects on productivity and inflation have the most negative product, so it moves the two most strongly in opposite directions. In each case, the disaster-risk shock is the max-share shock within the space orthogonal to the removed direction, with the baseline normalisation. Intervals come from the baseline residual bootstrap with 600 draws, which re-estimate the VAR, the auxiliary regression and both identifications in every draw.

Neither control changes the shock much. The removed directions account for 2.39 and 4.45 percent of the baseline shock's variance. The new shocks have correlations of 0.994 (bootstrap 10th–90th percentiles 0.902 to 0.999) and 0.976 (0.905 to 0.999) with the baseline shock. The shock explains 55.82 and 53.88 percent of the variance of disaster probability at the identification band, against 56.51 percent in the baseline, and its impact effect on disaster probability is 4.003 and 3.923 percentage points, against 4.065. Figure 15 shows the responses. Per percentage point of disaster probability on impact, inflation reaches its trough after one quarter, at −0.099 and −0.099 percentage points, against −0.100 in the baseline, while output reaches a trough of −0.110 and −0.119 percent after 3 quarters, against −0.097 percent. Consumption, investment and employment also fall slightly more than in the baseline, and under the supply-shock control the output trough lies just below the lower edge of the baseline band at that horizon. Removing productivity or supply

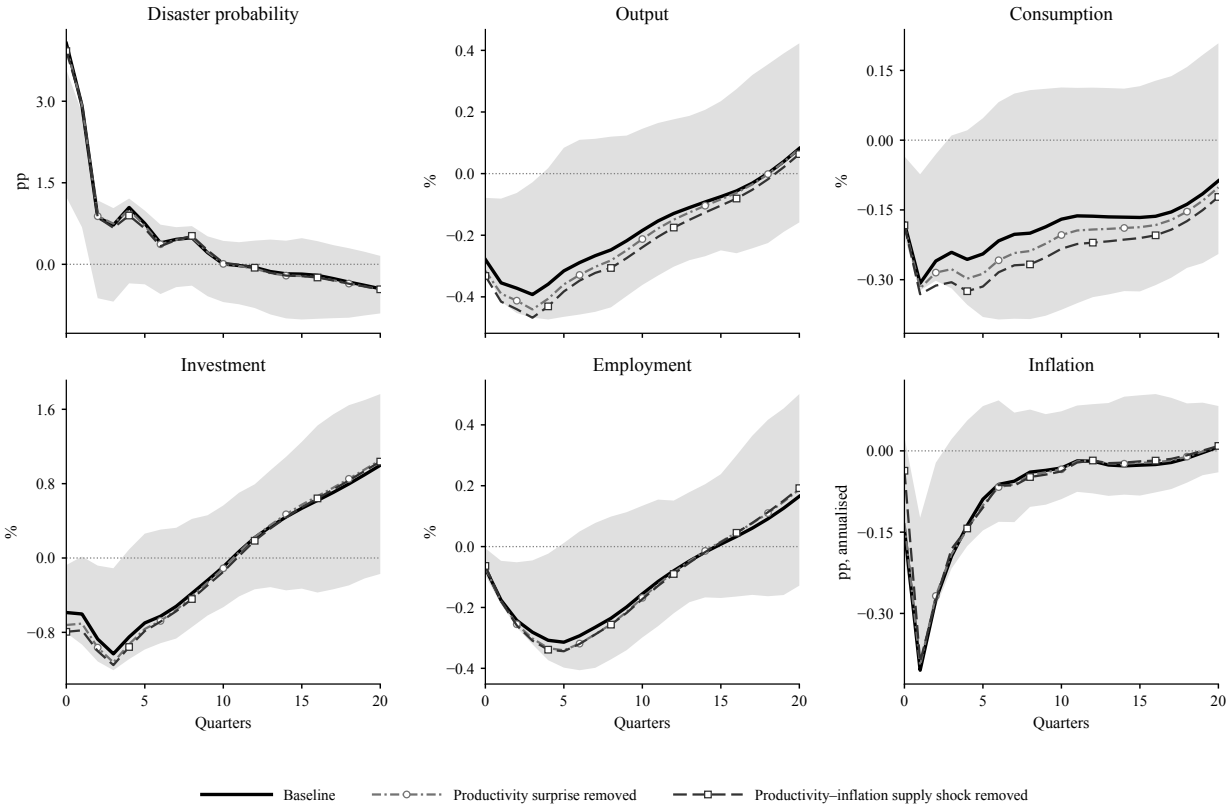


Figure 15: Responses to a disaster-risk shock after removing productivity shocks

Notes: Responses to a one-standard-deviation disaster-risk shock from impact through 20 quarters, 1996Q1–2019Q4. The solid line is the baseline max-share shock and the shaded area its pointwise 80% bootstrap band. The other lines re-identify the shock orthogonal to the fitted surprise in utilisation-adjusted productivity (Fernald, 2014), or to the supply shock that moves productivity and inflation most strongly in opposite directions. Disaster probability is in percentage points, real quantities in percent and inflation in annualised percentage points.

shocks therefore leaves the contraction at least as deep as in the baseline.

A.2.3 Controlling for lagged risk-neutral variance

The baseline VAR measures broad market risk with the VIX. To check whether the output response survives when the VAR conditions on market variance directly, I estimate two further VARs over the baseline sample, with the same constant, three lags and transformations. The first replaces the VIX with SVIX^2, the annualised one-month risk-neutral variance of the market return of Martin (2017); the second keeps the VIX and adds SVIX^2. I square the daily index before averaging within each quarter, using the days on which both the VIX and SVIX are available. Because SVIX^2 enters each VAR as an endogenous variable, three of its lags enter every equation, so each VAR controls for

lagged risk-neutral variance. Neither VAR holds variance fixed, however, because the identified shock remains free to move current and later variance. I identify the disaster-risk shock as in the baseline and orient it to raise the disaster probability on impact, but drop the baseline's additional sign condition on stock returns. Its correlation with the baseline shock is 0.95 in the first VAR and 0.89 in the second.

With SVIX^2 in place of the VIX, a one-standard-deviation shock lowers output by 0.21 percent on impact, with an 80 percent pointwise interval of −0.26 to −0.02 percent. The interval lies below zero on impact and one quarter later and includes zero at every later horizon up to 30 quarters. In the baseline, it stays below zero from impact through three quarters. In this VAR, the interval for the average response from impact through four quarters also includes zero. When the VAR includes both the VIX and SVIX^2, output falls by 0.17 percent on impact, with an interval of −0.24 to 0.04 percent. Here the interval includes zero at every horizon up to 30 quarters.

I also restrict the shock in four ways within each VAR. The first makes it orthogonal to the shock that explains the most variation in SVIX^2. The other three set the response of SVIX^2 to zero on impact, from impact through four quarters, or on impact while also making the shock orthogonal to a fitted productivity surprise. The orthogonality restriction removes one variance-related direction but leaves SVIX^2 free to respond. The zero restrictions fix its measured response only at the stated horizons. The restricted shocks explain less of the variation in the disaster probability and are only weakly correlated with the unrestricted shock. Whatever the restriction, the 80 percent interval for output includes zero at every horizon up to 30 quarters in both VARs. The point estimates rise at first and fall at longer horizons, but neither movement is distinguishable from zero. The restrictions can also remove disaster-risk news that itself moves variance, so the absence of a distinguishable contraction does not show whether disaster risk has an effect with variance held fixed.

Intervals are pointwise 10th–90th percentiles from 600 residual-bootstrap draws, which re-estimate each VAR and re-apply every identification in each draw. Intervals for averages over horizons take the average within each draw. Some draws, 50 in the first VAR and 70 in the second, yield dynamically unstable but finite estimates. I keep them.

A.2.4 Full sample through 2025Q2

I re-estimate the max-share VAR on the full available sample, 1996Q1–2025Q2 ($T = 118$). Relative to the baseline window, it adds the COVID-19 collapse and recovery, the inflation surge and tightening cycle, and the 2025 tariff episode. Figures 16 and 17 report the resulting impulse responses, which I construct as in the baseline.

The impact pattern is unchanged and demand-like. A shock that raises disaster risk is

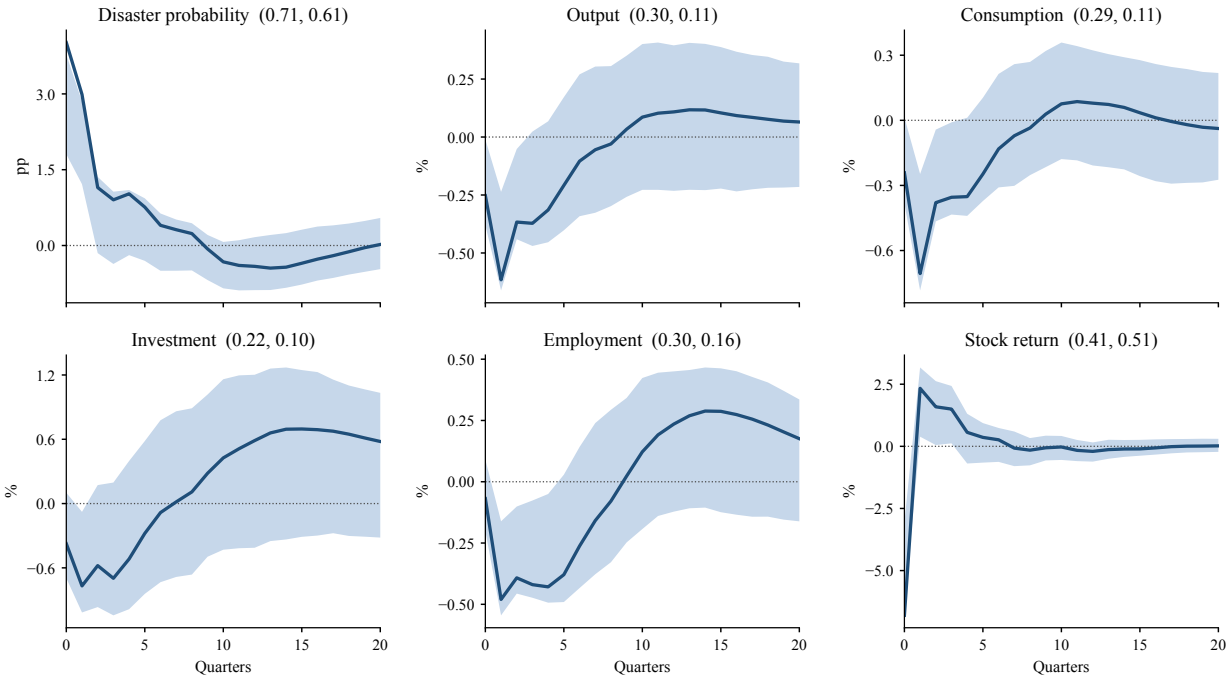


Figure 16: Real effects of a disaster-risk shock: full sample (1996Q1–2025Q2)

Notes: Impulse responses to the max-share disaster-risk shock estimated on the full sample including COVID-19. Horizons run from impact to 20 quarters. In each panel title, the first number in parentheses is the shock's share of the forecast-error variance at business-cycle frequencies (periods of 6–32 quarters), and the second its share at near-unconditional frequencies (periods of 2–500 quarters). Shaded areas are 80% bootstrap bands, from the 10th to the 90th percentile.

followed by a synchronised decline in output, consumption, investment and employment, a fall in the 5-year Treasury yield and inflation, and a rise in the VIX. The declines in real activity are larger than in the baseline but less persistent: output and consumption trough at roughly -0.6 and -0.7 percent one quarter after the shock and turn positive after 8 quarters. Because the pandemic quarters dominate the reduced-form residual covariance, the bootstrap bands also widen substantially. The output response is distinguishable from zero for the first three quarters, while the consumption and investment bands include zero at most horizons.

The shock's business-cycle variance shares rise on the full sample. It accounts for roughly 71 percent of the business-cycle forecast-error variance of the disaster probability, 41 percent of stock returns and 59 percent of the VIX. For output, consumption, investment and employment, the shares are 22–30 percent. The shock's near-unconditional shares of real activity (10–16 percent) are, however, below the baseline values. The 2020–2025 quarters therefore load heavily on the identified shock at cyclical frequencies but make its

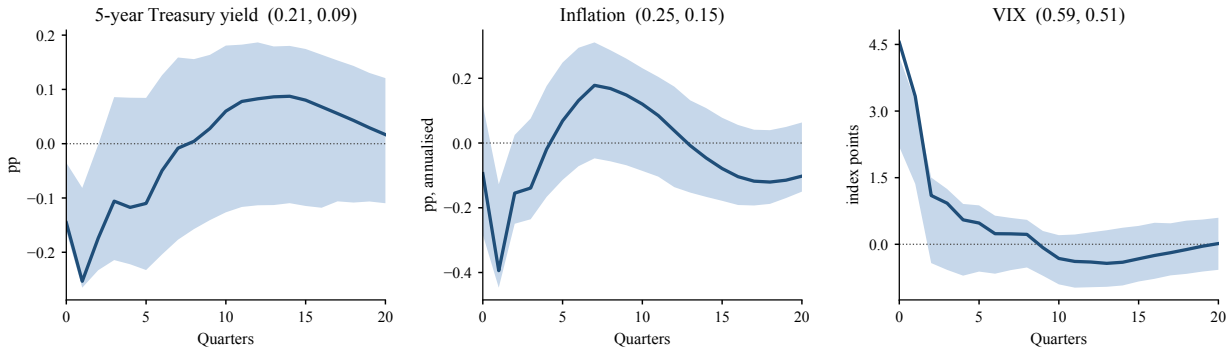


Figure 17: Financial and nominal responses to a disaster-risk shock: full sample (1996Q1–2025Q2)

Notes: Impulse responses of the 5-year Treasury yield, inflation (GDP deflator, annualised), and the VIX to the max-share disaster-risk shock estimated on the full sample including COVID-19. Horizons run from impact to 20 quarters. In each panel title, the first number in parentheses is the shock's share of the forecast-error variance at business-cycle frequencies (periods of 6–32 quarters), and the second its share at near-unconditional frequencies (periods of 2–500 quarters). Shaded areas are 80% bootstrap bands, from the 10th to the 90th percentile.

effects less precisely estimated.

A.2.5 Recursive identification

I also estimate recursive (Cholesky) VARs under two standard timing assumptions from the uncertainty and credit-spread literatures. Recursive identification makes its assumptions explicit, because the disaster-risk shock is the part of the change in disaster risk that the variables ordered before it do not explain. The results can depend on the ordering, so I report two. The reduced-form VAR is

$$y_t = c + \sum_{\ell=1}^{p} B_\ell y_{t-\ell} + u_t, \qquad E[u_t u_t'] = \Sigma_u,$$

and recursive identification sets $\Sigma_u = PP'$, where P is lower triangular and the structural shocks satisfy $u_t = P\varepsilon_t$. All the identifying assumptions therefore lie in the ordering of y_t. I estimate both VARs on quarterly US data over 1996Q1–2025Q2 ($T = 118$). The policy-rate measure is the Wu–Xia shadow federal funds rate (Wu and Xia, 2016) through June 2023, when the published series ends, and the effective federal funds rate thereafter.

The first, preferred ordering follows the logic of Gilchrist and Zakrajšek (2012). It places the slow macro block before option-implied disaster risk and the financial and policy block

after it:

$$y_t^{GZ} = \left[\Delta c_t,\ \Delta i_t,\ \Delta y_t,\ \pi_t,\ p_t,\ r_t^e,\ y_t^{5Y},\ i_t^{pol}\right]',$$

with $p = 2$ lags. Innovations to consumption, investment, output growth and inflation can then move disaster-risk pricing within the quarter, and the disaster-risk shock, the fifth orthogonalised shock, can move equity returns, the safe yield and policy on impact. The ordering is the more conservative one here, because it does not assume that option-implied disaster risk is predetermined with respect to same-quarter macro news. The second ordering, from the uncertainty-shock approach of Basu and Bundick (2017), puts disaster risk first:

$$y_t^{BB} = \left[p_t,\ y_t,\ c_t,\ i_t,\ h_t,\ P_t,\ m_t,\ i_t^{pol}\right]',$$

with $p = 4$ lags. This risk-first ordering assumes that disaster-risk news is predetermined within the quarter and can affect every other variable contemporaneously. In both specifications, I orient the disaster-risk shock so that p_t rises on impact. I then scale all responses by $s = 0.01/IRF_p(0)$, with disaster probability in decimal units, so the plotted responses correspond to a one-percentage-point increase in annual option-implied disaster probability. Growth-rate and inflation responses in the Gilchrist–Zakrajšek ordering are cumulated to level responses, and the log-level variables in the Basu–Bundick ordering are reported in percent deviations. I construct confidence bands from a residual bootstrap with 1,000 replications.

The recursive estimates show the same qualitative pattern as the max-share and heteroskedastic designs. After a one-percentage-point rise in disaster risk, real activity weakens, the price level softens, safe yields and risky valuations fall, and the policy rate declines modestly. The pattern is clearest in the Gilchrist–Zakrajšek-style ordering (Figure 18), where the disaster-risk shock is identified after removing contemporaneous slow-moving macro and inflation innovations. Over the first eight quarters, output troughs at about −0.09 percent, consumption at −0.08 percent and investment at −0.20 percent, and the price level falls by about 0.02 percent. Over the same horizon, the 5-year yield and the policy rate fall by about 6 and 6 basis points. Real equity valuations drop by roughly 1.4 percent on impact. At their troughs, the 90 percent bands of the output, consumption, investment, yield and equity responses exclude zero.

In the Basu–Bundick risk-first ordering (Figure 19), output, consumption, investment, hours and the price level all decline as well, and the policy rate troughs at about −0.11 percentage points. The result therefore holds under both recursive timing conventions. The disaster-risk shock accounts for a sizeable share of the eight-quarter forecast-error variance of risky valuations (about 35 percent of the real stock return in the Gilchrist–Zakrajšek

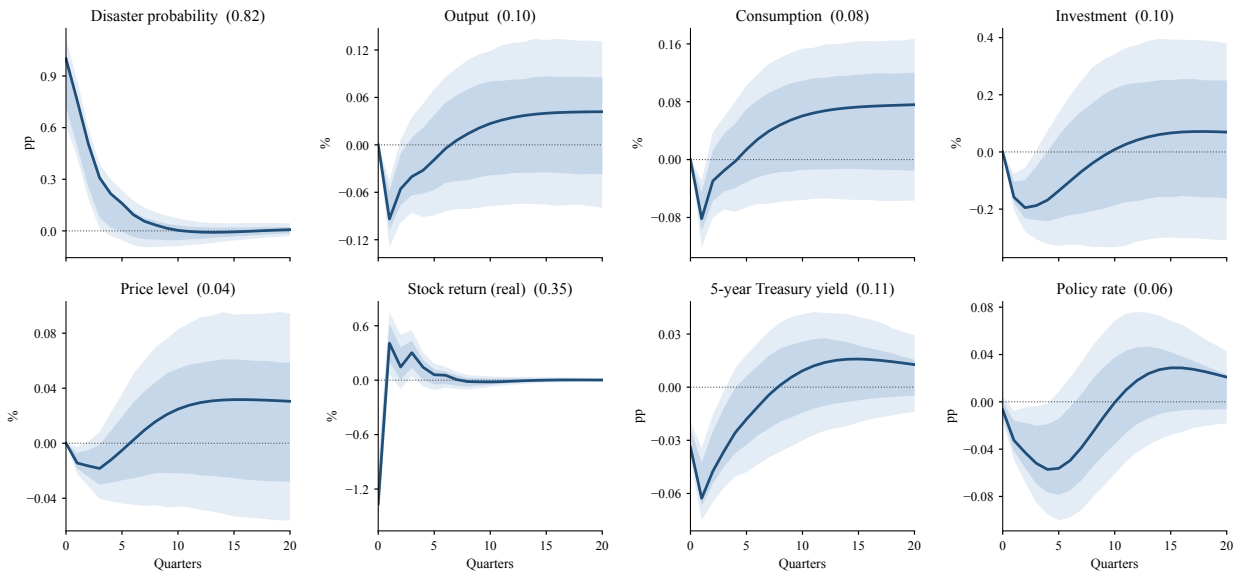


Figure 18: Recursive VAR responses to a disaster-risk shock: Gilchrist–Zakrajšek-style ordering

Notes: Orthogonalised impulse responses from a quarterly recursive VAR in which consumption growth, investment growth, output growth, and inflation are ordered before option-implied disaster risk, while the real equity return, the 5-year Treasury yield, and the policy rate are ordered after it (1996Q1–2025Q2, $T = 118$, 2 lags). The shock is signed and scaled so that annual option-implied disaster probability rises by one percentage point on impact. Horizons run from impact to 20 quarters. Growth and inflation responses are cumulated to level responses. Dark and light shaded regions are 68% and 90% residual-bootstrap bands. The number in parentheses in each panel title is the disaster-risk shock's share of the forecast-error variance of that equation at the eight-quarter horizon.

ordering) and of the policy rate (about 28 percent in the Basu–Bundick ordering). Its shares of real activity are smaller, as in the other designs. Table 6 lists the signs of the responses under all four identification schemes.

A.3 Historical evaluations of a risk-adjusted rule

I now ask how the risk-adjusted neutral rate changes the policy rate prescribed by a Taylor rule, which also responds to inflation and the output gap. Figure 20 compares the realised federal funds rate with the standard Taylor rule and with a version that replaces its fixed neutral rate with my risk-adjusted estimate. Both rules use the real-time macroeconomic inputs of Nakamura et al. (2025). In annual percentage units, the two rules are

$$\begin{aligned} i_t^{\mathrm{TR}} &= 2 + \pi_t^{(4)} + 0.5(\pi_t^{(4)} - 2) + 0.5x_t^{\mathrm{RT}}, \\ i_t^{\mathrm{RA}} &= r_t^{\star} + \pi_t^{(4)} + 0.5(\pi_t^{(4)} - 2) + 0.5x_t^{\mathrm{RT}}. \end{aligned} \tag{25}$$

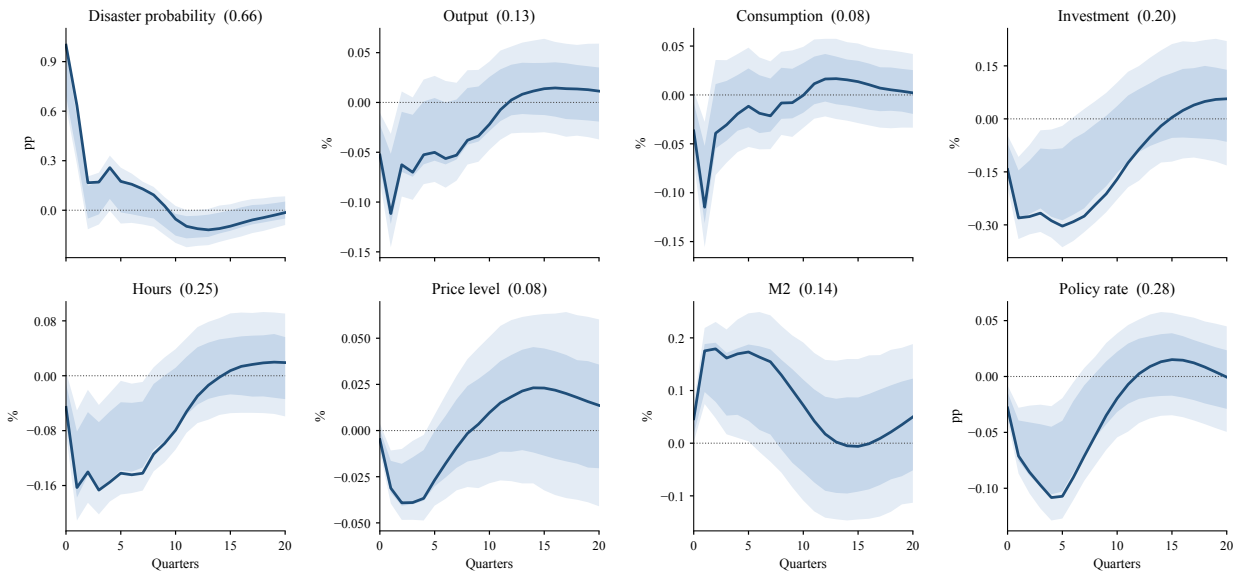


Figure 19: Recursive VAR responses to a disaster-risk shock: Basu–Bundick risk-first ordering

Notes: Orthogonalised impulse responses from a quarterly recursive VAR in which option-implied disaster risk is ordered first, followed by output, consumption, investment, hours, the price level, money, and the policy rate (1996Q1–2025Q2, $T = 118$, 4 lags). The shock is signed and scaled so that annual option-implied disaster probability rises by one percentage point on impact. Horizons run from impact to 20 quarters. Log-level variables are reported in percent deviations; rates and disaster probability are in percentage points. Dark and light shaded regions are 68% and 90% residual-bootstrap bands. The number in parentheses in each panel title is the disaster-risk shock's share of the forecast-error variance at the eight-quarter horizon.

Here $\pi_t^{(4)}$ is four-quarter GDP-deflator inflation and x_t^{RT} is the real-time Greenbook output gap, or the real-time CBO gap when the Greenbook measure is unavailable. I use the authors' original data over 1995Q1–2025Q1. The neutral-rate input is the smoothed semi-structural estimate of Section 5.6, which uses observations through 2026Q2. Because the comparison combines real-time macro inputs with a retrospective neutral-rate estimate, it does not test a rule that policymakers could have implemented with the information they had at the time.

Full tracking changes the prescribed rate by exactly $r_t^\star - 2$ percentage points relative to the standard rule. In 2008Q4, the standard rule prescribes 2.69% and the risk-adjusted rule −5.07%, against a realised funds rate of 0.51%. The realised rate therefore lies between the two prescriptions. It is easier than the standard rule but tighter than the risk-adjusted rule, in line with the contractionary stance found in Section 5.6. Both prescriptions are negative in 2020Q2, −3.58% under the standard rule and −7.02% with full tracking, against a funds rate of 0.06%.

Table 6: Responses to a disaster-risk shock across identification schemes

Identification	Output	Cons.	Inv.	Emp./ Hours	Inflation	Safe yield	Policy rate	Equity
Max-share	–	–	–	–	–	–	—	–
Heteroskedastic SVAR	flat	—	—	–	–	–	–	–
Cholesky: GZ-style	–	–	–	—	–	–	–	–
Cholesky: BB-style	–	–	–	–	–	—	–	—

Notes: The table reports the sign of the median response over the first two years to a shock that raises annual option-implied disaster risk. The max-share shock is a unit structural innovation, and the other designs are normalised to a one-percentage-point impact. A dash "—" marks a variable that the specification does not include. In the heteroskedastic SVAR, output is monthly GDP and employment is payroll employment; the safe yield is the 10-year minus 3-month term spread, which falls on impact and reverts within six months.

A deviation from a Taylor-rule prescription and the real-minus-neutral gap in Section 5.6 assess policy against different benchmarks. The rule also responds to inflation and the output gap, and here it uses real-time macroeconomic inputs. In 2020Q2, for example, the real rate was 0.79 percentage points above the risk-adjusted neutral rate, but the funds rate lay far above both rule prescriptions, because the rules respond to the collapse in output.

These calculations change only the neutral-rate intercept and hold the historical macroeconomic path fixed. They do not measure the equilibrium or welfare effects of following either rule, and the negative prescriptions need not have been feasible policy choices. Appendix A.4 repeats the comparison for a balanced-approach rule with current-vintage core PCE inflation, the CBO output gap and a larger response to the gap. It shows that the prescriptions also depend on these conventional inputs and coefficients.

A.4 A balanced-approach rule comparison

To show how the rule comparison of Appendix A.3 depends on its conventional inputs, I evaluate the balanced-approach rule

$$i_t^{\mathrm{BA}}(\bar{r}_t) = \bar{r}_t + \pi_t^{\mathrm{PCE}} + 0.5(\pi_t^{\mathrm{PCE}} - 2) + y_t^{\mathrm{CBO}}, \qquad \bar{r}_t \in \{2, r_t^{\star}\}. \tag{26}$$

Inflation is the four-quarter change in the quarterly mean core PCE index, and the output gap is $100(Y_t/Y_t^{\mathrm{CBO}} - 1)$, the percentage gap between real GDP and CBO potential GDP. Relative to equation (25), the rule keeps the inflation coefficients and the smoothed neutral-rate estimate but replaces the real-time inputs with current-vintage core PCE inflation and the CBO gap, and it doubles the coefficient on the output gap from one half to one.

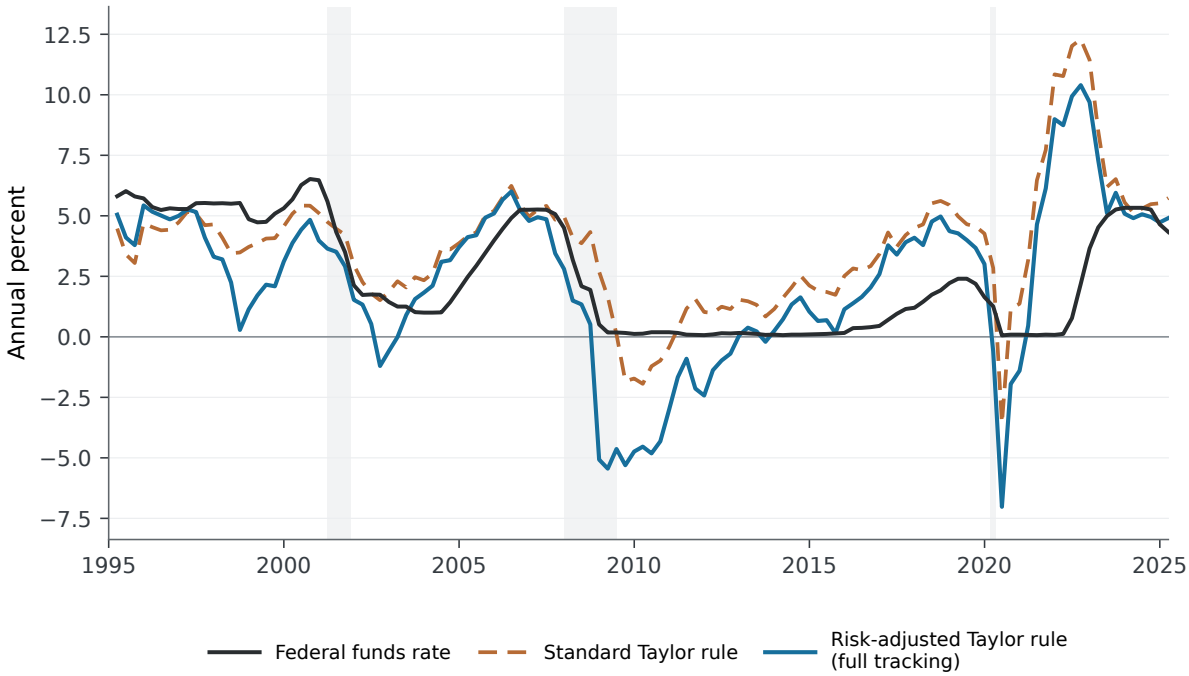


Figure 20: Standard and risk-adjusted Taylor rules

Notes: Realised federal funds rate, the standard Taylor rule, and full tracking of the risk-adjusted neutral rate. Both rules use the original real-time inflation and output-gap inputs of Nakamura et al. (2025) over 1995Q1–2025Q1. Because the risk-adjusted rate is smoothed with observations through 2026Q2, the comparison is retrospective. Neither rule imposes a lower bound or interest-rate smoothing. Recessions are shaded.

Figure 21 plots both evaluations against the funds rate. The conventional inputs matter on their own. With a fixed neutral rate, the balanced-approach rule prescribes 0.05% in 2008Q4, against 2.69% for the standard Taylor rule of Appendix A.3. As in that comparison, the risk-adjusted neutral rate lowers the prescription markedly in the crises, to −7.71% in 2008Q4 and to −9.80% in 2020Q2 (against −6.35% with a fixed neutral rate). At 2026Q2, the funds rate is 3.63%. Core PCE inflation of 3.38% and a CBO gap of 0.83 percentage points imply rule rates of 6.89% with a fixed neutral rate and 6.14% with the estimated neutral rate. Because the two rules differ by $r_t^\star - 2$ in every quarter, the difference comes from the whole neutral-rate estimate, not only from its disaster-risk component.

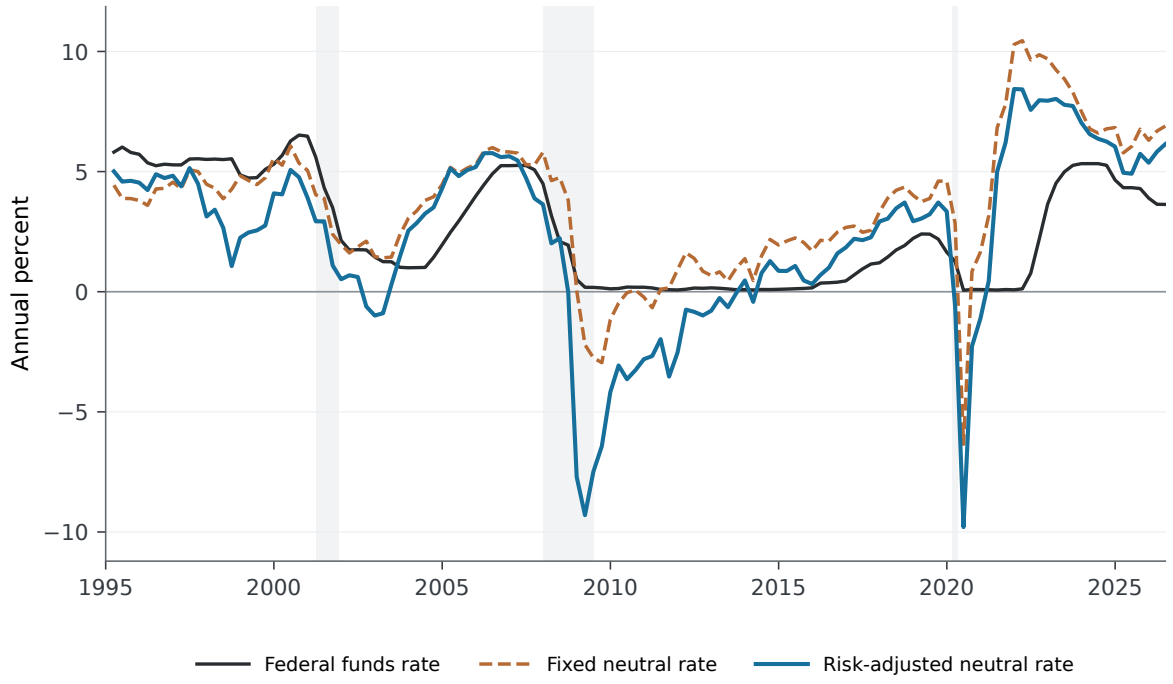


Figure 21: Historical balanced-approach rule comparisons

Notes: Quarterly funds rate and evaluations of equation (26) using a fixed 2% neutral rate or the smoothed risk-adjusted estimate. Inflation and the CBO gap come from the data vintage downloaded on 16 September 2026. Recessions are shaded. Neither rule imposes interest-rate smoothing or a lower bound. The calculations hold the economy's historical path fixed, so they show what each rule would have prescribed, not how the economy would have responded to it.

B Heteroskedastic SVAR for Disaster-Risk Shocks

Section 3.4 summarises a structural VAR (SVAR) identified by heteroskedasticity, and this appendix gives the details. I adapt the framework of Brunnermeier et al. (2021) (henceforth BPSS) to monthly data, with the option-implied disaster probability of Section 3 as the key variable.

B.1 Data and transformations

I estimate the baseline heteroskedastic SVAR at monthly frequency over 1994M9–2025M8. The vector of observables is

$$y_t = [p_t,\ x_t,\ P_t,\ i_t,\ TS_t,\ EBP_t,\ Q_t]',$$

where p_t is the option-implied annual disaster probability, x_t is real activity, P_t is the PCE price index, i_t is the policy-rate measure, TS_t is the 10-year Treasury yield minus the 3-month Treasury bill rate, EBP_t is the excess bond premium of Gilchrist and Zakrajšek (2012), and Q_t is a real equity wealth index. Table 7 lists the variables and their transformations. Quantity and price variables enter in log levels, while rates, spreads and disaster probabilities enter in decimal units, so that a value of 0.01 corresponds to one percentage point. The real equity variable is a CRSP total-return equity wealth index deflated by the PCE price index. For the policy rate, I use the Wu–Xia shadow federal funds rate (Wu and Xia, 2016) until the published series ends in June 2023, and the effective federal funds rate after that.

The baseline activity variable is industrial production, but the figures report the responses of monthly GDP and employment. Each comes from an otherwise identical system in which that variable replaces industrial production, while the financial and nominal panels come from the baseline system.

Table 7: Monthly variables in the heteroskedastic SVAR

Variable	Description	Transformation
p_t	Option-implied annual disaster probability	Level, decimal
x_t	Industrial production / monthly GDP / employment	Log level
P_t	PCE price index	Log level
i_t	Policy-rate measure (shadow rate, spliced)	Level, decimal
TS_t	10-year Treasury minus 3-month Treasury bill	Level, decimal
EBP_t	Excess bond premium	Level, decimal
Q_t	Real equity wealth index	Log level

B.2 Structural VAR

The empirical model is a structural VAR with n variables and p lags:

$$A_0 y_t = c + \sum_{\ell=1}^{p} A_\ell y_{t-\ell} + \varepsilon_t, \tag{27}$$

where A_0 is the contemporaneous structural matrix, A_ℓ are lag coefficient matrices, c is a vector of constants, and ε_t is an $n \times 1$ vector of mutually independent structural shocks. In

the baseline specification, $n = 7$ and $p = 6$. The reduced-form representation is

$$y_t = b + \sum_{\ell=1}^{p} B_\ell y_{t-\ell} + u_t, \qquad u_t = A_0^{-1}\varepsilon_t, \quad B_\ell = A_0^{-1}A_\ell, \quad b = A_0^{-1}c. \tag{28}$$

The identification problem is to recover A_0, which maps reduced-form innovations into orthogonal structural shocks.

B.3 Identification by heteroskedasticity

Following BPSS, identification comes from changes in the relative variances of structural shocks across regimes. Let $m(t) \in \{1, \ldots, M\}$ denote the variance regime at date t. The structural shock covariance matrix in regime m is

$$E[\varepsilon_t \varepsilon_t' \mid m(t) = m] = \Lambda_m, \tag{29}$$

where Λ_m is diagonal. The propagation matrices $A_0, A_1, \ldots, A_p$ are held fixed across regimes, while the diagonal elements of Λ_m vary across regimes. The implied reduced-form covariance matrix in regime m is

$$\Sigma_m = E[u_t u_t' \mid m(t) = m] = A_0^{-1}\Lambda_m (A_0^{-1})'. \tag{30}$$

Changes in Σ_m across regimes can therefore come only from changes in the relative variances of the orthogonal structural shocks. For two regimes i and j,

$$\Sigma_i^{-1}\Sigma_j = A_0' \Lambda_i^{-1} \Lambda_j (A_0^{-1})', \tag{31}$$

so the eigenvalues of $\Sigma_i^{-1}\Sigma_j$ are the ratios of structural variances across regimes, and its eigenvectors recover the rows of A_0, up to sign and permutation, provided that different structural shocks have distinct variance profiles. Scale is fixed by the normalisation

$$\frac{1}{M}\sum_{m=1}^{M} \lambda_{k,m} = 1, \qquad k = 1, \ldots, n, \tag{32}$$

where $\lambda_{k,m}$ is the k-th diagonal element of Λ_m, so that the average structural variance of each shock equals one across regimes.

The baseline model uses $M = 6$ variance regimes:

1. 1994M9–2001M12, late 1990s/dot-com period,
2. 2002M1–2007M7, pre-crisis low-volatility period,
3. 2007M8–2010M12, Global Financial Crisis,
4. 2011M1–2019M12, post-crisis recovery,
5. 2020M1–2021M12, COVID and renewed effective-lower-bound period,
6. 2022M1–2025M8, inflation surge, then tightening and easing.

The dates need not capture every shift in volatility, because identification requires only that the average relative variances of the structural shocks differ across regimes. A misdated boundary would therefore weaken identification and widen the posterior bands, rather than generate the results.

B.4 Fat-tailed structural shocks

In the preferred specification, the structural shocks are Student-t, as in BPSS. Fat tails matter because the crisis observations in disaster probability, credit spreads and equity prices are large, and a Gaussian model can give them too much weight. The Student-t specification can be written as a scale mixture of normals:

$$\varepsilon_{k,t} \mid s_{k,t}, m(t) = m \sim N(0, \lambda_{k,m} s_{k,t}), \qquad s_{k,t} \sim IG(\nu/2, \nu/2), \tag{33}$$

so that marginally $\varepsilon_{k,t} \sim t_\nu$. I use the BPSS calibration, $\nu \approx 5.70$. With fat tails, a large isolated shock can be a draw from the structural distribution itself, so the model need not explain it entirely through persistent regime-level volatility.

B.5 Bayesian estimation

I estimate the model with the BPSS Bayesian procedure. I first compute the Gaussian posterior mode and then use it to start the posterior simulation under Student-t errors. The parameter vector includes A_0 and the lag matrices $A_1, \ldots, A_p$, as well as the regime-specific variance factors Λ_m and the scale-mixture terms $s_{k,t}$.

The priors are those of BPSS. Elements of A_0 have diffuse Gaussian priors centred on a scaled identity matrix, and the lag coefficients a Minnesota-style prior. The variance factors have a Dirichlet prior that enforces the average-variance normalisation in equation (32). I generate posterior draws by Gibbs sampling and compute impulse responses for each draw. The figures plot the posterior median at each horizon, with 68 percent posterior

bands. These bands come from a single chain of the adapted BPSS sampler, without formal multi-chain convergence diagnostics, but the qualitative conclusions are the same at the Gaussian posterior mode and across the activity-variable variants.

B.6 Shock labelling

Identification through heteroskedasticity recovers structural shocks only up to sign and permutation, so I label the shocks ex post by their impact responses. With $IRF_{i,j}(h)$ the response of variable i at horizon h to structural shock j, the disaster-risk shock is

$$j_{DR} = \arg\max_j |IRF_{p,j}(0)|, \qquad IRF_{p,j_{DR}}(0) > 0, \tag{34}$$

with the sign normalised so that disaster probability rises. The signs of the output, price, equity and policy responses to the disaster-risk shock are estimated rather than imposed.

B.7 Scaling

I scale the disaster-risk shock to raise annual disaster probability by one percentage point on impact. Let $IRF_{p,DR}(0)$ be the impact response of disaster probability to the identified disaster-risk shock, where disaster probability is measured in decimal units. Define

$$s_{DR} = \frac{0.01}{IRF_{p,DR}(0)}, \qquad \widetilde{IRF}_{i,DR}(h) = s_{DR}\, IRF_{i,DR}(h). \tag{35}$$

The figures plot $100\,\widetilde{IRF}_{i,DR}(h)$, which is approximately a percent deviation for log-level variables and a percentage-point change for rates, spreads and disaster probability. I construct the inflation response from the price-level response: because P_t enters in logs, the annualised monthly inflation response is

$$\widetilde{IRF}_{\pi,DR}(h) = 1200\left[\widetilde{IRF}_{\log P,DR}(h) - \widetilde{IRF}_{\log P,DR}(h-1)\right], \tag{36}$$

with the impact response defined relative to the pre-shock price level.

B.8 Forecast-error variance shares

Let Θ_h denote the structural moving-average coefficient at horizon h, so that

$$y_{t+h} - E_t y_{t+h} = \sum_{s=0}^{h} \Theta_s \varepsilon_{t+h-s}.$$

The share of the H-month-ahead forecast-error variance of variable i due to the disaster-risk shock is

$$FEV_{i,DR}(H) = \frac{\sum_{h=0}^{H-1} \left(e_i' \Theta_h e_{DR} \right)^2}{\sum_{j=1}^{n} \sum_{h=0}^{H-1} \left(e_i' \Theta_h e_j \right)^2}, \tag{37}$$

where e_i selects variable i and e_{DR} selects the disaster-risk shock. The figures report two summaries of these shares: their average over horizons H from 6 to 32 months, and the share at the 24-month horizon.

B.9 Responses to a disaster-risk shock

Figure 22 reports the responses. The first panel shows a persistent but transitory increase in disaster risk: a one-percentage-point rise in annual disaster probability decays gradually over the 20 months shown. Financial variables respond immediately. Real equity wealth falls by roughly 1.0 percent on impact, and the excess bond premium rises by about 2 basis points and is back at its pre-shock level within about six months. In the model, higher perceived disaster risk raises the compensation investors require for holding risky claims and so lowers their prices, and these responses are consistent with that channel. Credit conditions tighten as well. The contraction in activity is moderate and builds gradually. Employment declines by about 0.06 percent after one year and 0.09 percent after 20 months, with posterior bands that exclude zero, while monthly GDP is approximately flat with wide posterior bands.

The nominal responses point to a demand shock rather than a negative supply shock, which would raise inflation as output falls. The price level instead falls for most of the first year, and the inflation response is negative on impact and partly reverses as the price level recovers. This pattern fits a risk-premium shock that contracts aggregate demand through asset prices, wealth and credit premia. The max-share evidence shows the same pattern, but here it arises without any ordering imposed between disaster risk and financial conditions. The policy rate declines by about 2 basis points and remains below its pre-shock level for roughly the first year, with a 68 percent band that excludes zero over that window. The response is small and grows more uncertain at longer horizons, so I read it as evidence of modest monetary accommodation, not as an estimate of a policy rule. Its sign agrees with the mechanism of the paper, in which higher disaster risk lowers the neutral rate.

The forecast-error variance shares qualify this interpretation. The disaster-risk shock explains about 72 percent of the variation in the disaster-probability series itself, but only a few percent of the variation in activity, equity prices, credit premia, prices and the policy rate. At monthly frequency, the shock therefore moves the economy in the direction the

model predicts but is a mild disturbance, not a dominant source of business-cycle variation.

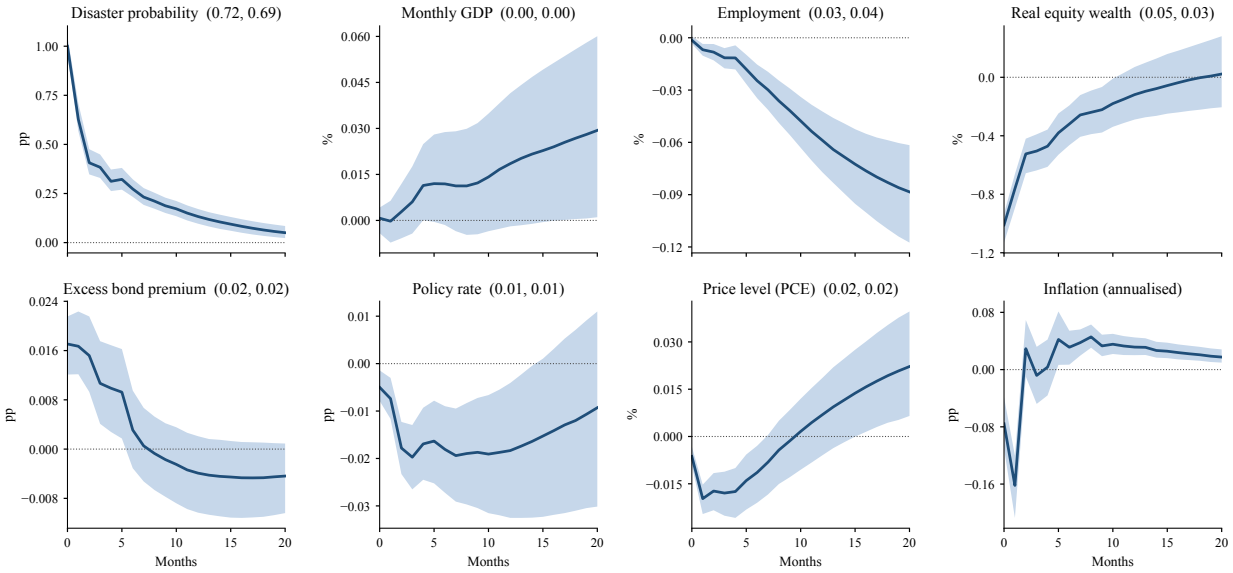


Figure 22: Responses to a disaster-risk shock: heteroskedastic SVAR

Notes: Posterior median responses to the disaster-risk shock identified in the heteroskedastic SVAR (1994M9–2025M8), scaled to raise annual option-implied disaster probability by one percentage point on impact. Horizons run from impact to 20 months, and shaded areas are 68% posterior bands. Quantity and price variables are in percent deviations, and rates, spreads, inflation and disaster probability in percentage points. Only the monthly GDP and employment panels come from the activity-variant systems. In each panel title, the first number in parentheses is the disaster-risk shock's forecast-error variance share averaged over horizons of 6 to 32 months, and the second its share at the 24-month horizon.

B.10 Policy-rate splice

The baseline policy series switches from the Wu–Xia shadow rate to the effective federal funds rate in July 2023, after the published shadow rate ends. By then the funds rate was already above five percent. Because the two rates differ by about one percentage point at the join, the splice introduces a step in the policy series. I therefore re-estimate the baseline and both activity-variant systems with a policy series that uses the shadow rate only in months when the funds rate is below 25 basis points and the funds rate otherwise (Figure 23). The disaster-probability, equity, credit-spread, price-level and inflation responses are unchanged. The policy-rate easing is smaller and shorter-lived (about −1.2 basis points after six months, and back to zero within a year), and the employment decline disappears. The baseline activity and policy responses are therefore sensitive to the splice, but the

financial and nominal evidence is not.

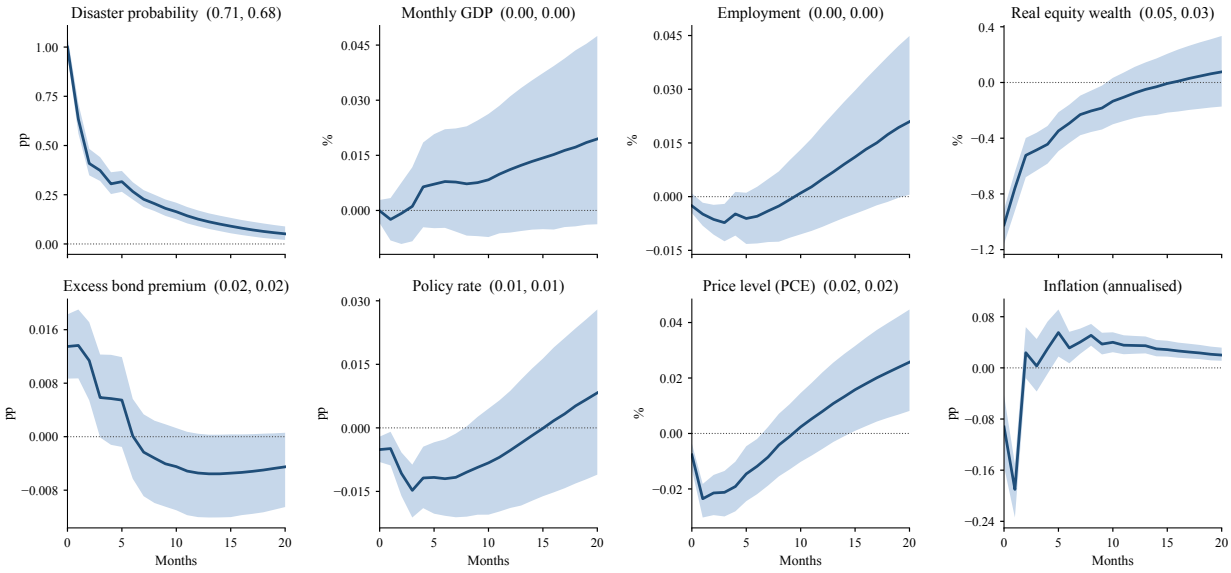


Figure 23: Impulse responses to a disaster-risk shock: lower-bound-only shadow rate

Notes: As in Figure 22, except that the policy-rate measure is the Wu–Xia shadow rate only in months when the effective federal funds rate is below 25 basis points and the effective federal funds rate otherwise (1994M9–2025M8). Horizons run from impact to 20 months.

C Alternative Measures of Financial Market Stress

This appendix extends the comparison in Figure 11 of Section 5.4 with two other indicators of financial-market stress, the VIX and the excess bond premium (EBP) of Gilchrist and Zakrajšek (2012), and with alternative policy rules. The VIX and EBP figures retain their longer plotting windows. As in Figure 11, the Taylor-rule prescription is the standard Taylor rule of equation (25), evaluated with the real-time inflation and output-gap data of Nakamura et al. (2025).

C.1 The VIX

Figure 24 replaces disaster probability with the VIX. The VIX measures the 30-day implied volatility of S&P 500 index options and is often read as a forward-looking indicator of market uncertainty and risk aversion (Bloom, 2009).

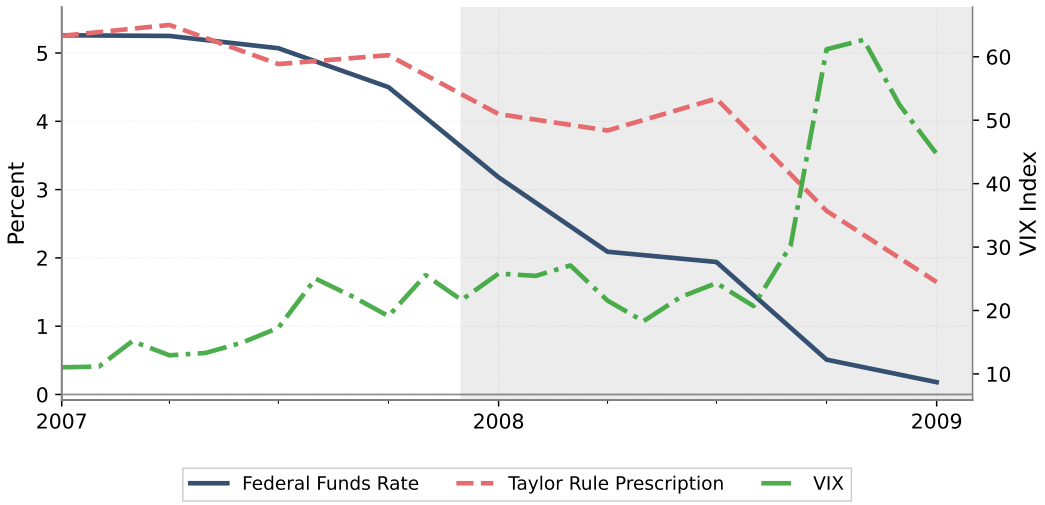


Figure 24: The federal funds rate, the Taylor rule and the VIX

Notes: The figure plots the effective federal funds rate, the prescription from the original Taylor (1993) rule using real-time data, and the VIX (monthly average). The VIX measures the implied volatility of S&P 500 index options and is obtained from the CBOE. The NBER-dated recession (December 2007–June 2009) is shaded.

The VIX fluctuated between 10 and 15 through mid-2007, rose in the summer of 2007 as turmoil emerged in the subprime mortgage market, and moved mostly in the 20–30 range until September 2008. After the collapse of Lehman Brothers, its monthly average rose

above 60, peaking in November.

The sustained rise in the VIX began as the Federal Reserve started cutting rates, although inflation and the output gap called for little easing (Figure 26). Its timing, together with the growing gap between the federal funds rate and the Taylor-rule prescription, is consistent with the Federal Reserve responding to financial-market stress in addition to inflation and the output gap.

C.2 The excess bond premium

Figure 25 replaces disaster probability with the EBP, the component of corporate bond spreads that expected default risk does not explain. The EBP therefore measures the compensation investors demand for bearing corporate credit risk. Gilchrist and Zakrajšek (2012) show that it predicts contractions in output, investment and employment, with information that standard financial indicators such as stock prices and the slope of the yield curve do not contain.

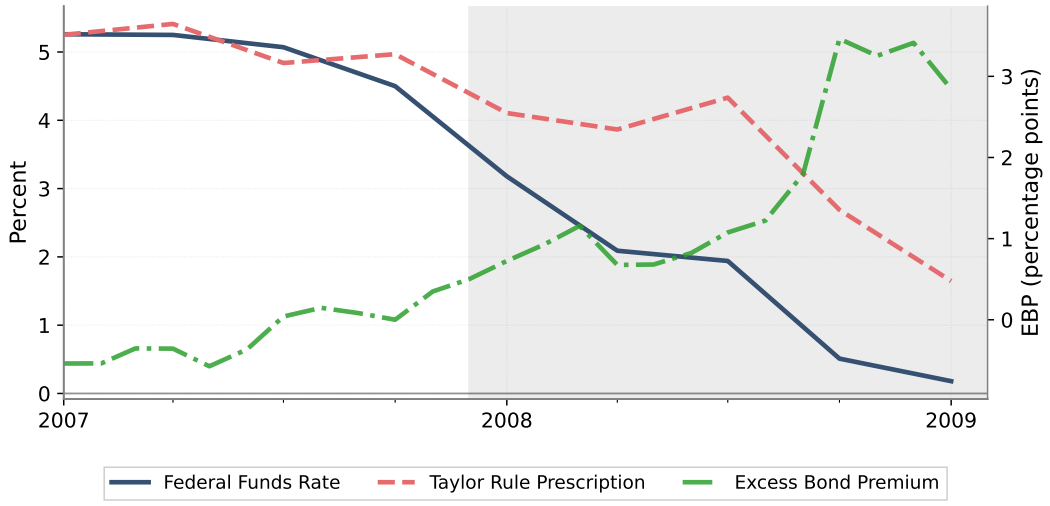


Figure 25: The federal funds rate, the Taylor rule and the excess bond premium

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 is the component of corporate bond spreads orthogonal to expected default risk, and so measures risk premia in credit markets. The NBER-dated recession (December 2007–June 2009) is shaded.

The EBP shows the same timing as disaster probability and the VIX. It was negative in the first half of 2007, began rising in mid-2007 and reached approximately 1 percentage

point by early 2008. After the collapse of Lehman Brothers, it rose to its highest level in the series.

As with the VIX, the EBP rose as the Federal Reserve departed from the standard Taylor rule, so policymakers may have been responding to credit-market stress as well as to equity-market volatility. All three measures rose from mid-2007, before the inputs to the standard Taylor rule called for aggressive monetary easing.

C.3 Alternative Taylor rules

Figure 26 plots the two inputs to the standard Taylor rule, real-time GDP-deflator inflation and the real-time output gap, over the crisis window. Through 2007, inflation was above target and the output gap was positive; the gap turned negative only in 2008, and only modestly until the second half of that year. Neither input therefore called for the aggressive easing that the Federal Reserve delivered from late 2007.

Figure 27 shows that the shortfall of the federal funds rate below the rule is not specific to the original Taylor rule. From late 2007 through the third quarter of 2008, all five standard rule variants prescribe a rate above the federal funds rate.

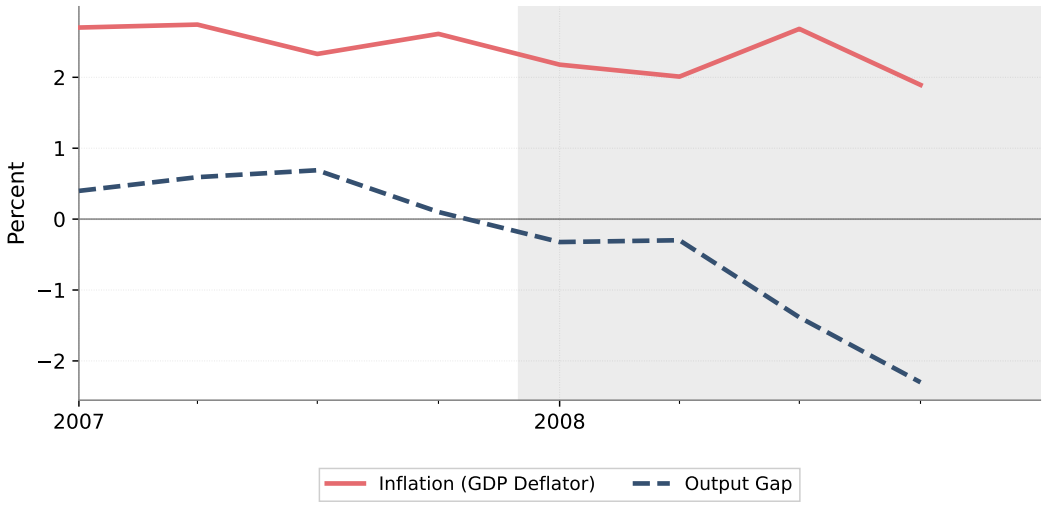


Figure 26: Inputs to the Taylor rule: inflation and the output gap

Notes: The figure plots real-time four-quarter GDP-deflator inflation and the real-time output gap (the Greenbook estimate, or the real-time CBO estimate where the Greenbook estimate is unavailable) from Nakamura et al. (2025). The NBER-dated recession (December 2007–June 2009) is shaded.

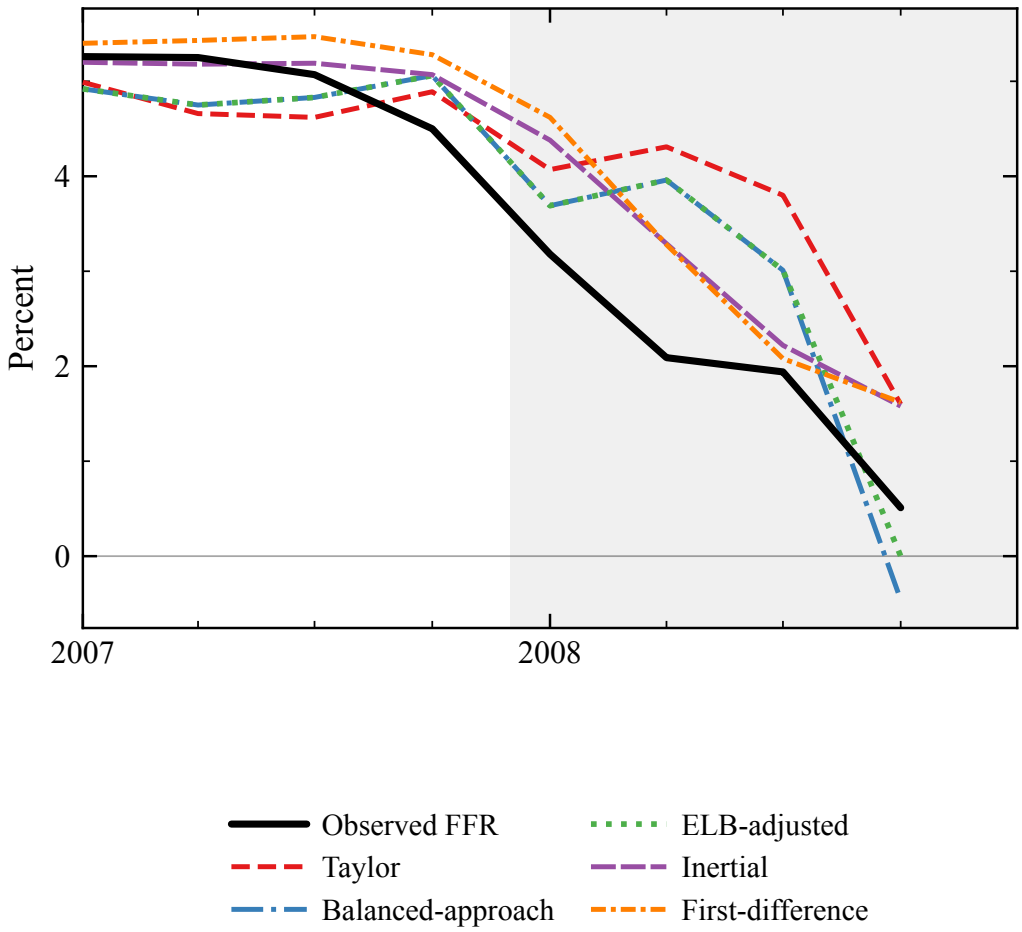


Figure 27: The federal funds rate and alternative Taylor rules

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. The prescriptions are taken from Figure 1 of the Federal Reserve Board's web page "Policy Rules and How Policymakers Use Them". The NBER-dated recession (December 2007–June 2009) is shaded.

D Quantitative Analysis: Estimation and Other Results

This appendix documents the estimation of the model in Section 5.1 and reports additional results for Section 5.

D.1 Measurement and the likelihood

The estimation uses five quarterly series over 1986Q1–2026Q2: GDP, CBO potential output, GDP-deflator inflation, the federal funds rate and the option-based disaster probability, which is available from 1994Q4 to 2025Q2. Model quantities are evaluated at the end of each quarter. The filter observes GDP growth and the gap between CBO potential output and GDP rather than their levels:

$$\begin{aligned}\Delta_q y_t &= \bar{g}/4 + \int_{t-1/4}^{t} g_s\, ds + \Delta_q(a^r + x + d)_t + e^y_t - e^y_{t-1/4},\\ \mathrm{cbo}_t - y_t &= -x_t - d_t + e^c_t - e^y_t,\end{aligned}$$

where $\Delta_q h_t = h_t - h_{t-1/4}$ and d_t is a discrepancy between measured and model output that follows a quarterly first-order autoregression with coefficient 0.5. Inflation and the disaster measure are observed with noise, $\pi^{\mathrm{obs}}_t = \pi^\star + \widehat{\pi}_t + e^\pi_t$ and $p_t = \bar{p} + \xi_t + e^p_t$, and so is the funds rate, $i^{\mathrm{obs}}_t = i_t + e^i_t$, when it exceeds 25 basis points. In the 36 quarters in which the funds rate was at or below 25 basis points, the observation contributes only the probability that the model's policy rate lies below that threshold,

$$\log \Phi\left(\frac{0.0025 - i_t}{\sqrt{s_i^2 + h^2_{i,t}/12}}\right),$$

where s_i is the standard deviation of the measurement error and $h^2_{i,t}/12$ the variance of the approximation error in the model's policy rate.

I evaluate the likelihood with a filter that uses the same discrete approximation of the state space as the solution of the model. Between quarters, the distribution of the states (a^r, ξ, g, i) evolves according to the model's transition probabilities, including disasters. In the solution, the disaster intensity of (16) is bounded between $\lambda_{\min} = 0.24\%$ and $\lambda_{\max} = 5\%$ a year, which keeps it within a plausible range in states of very low or very high risk. Conditional on these states, the output discrepancy, the growth of the trend over the quarter and the previous quarter's output level are jointly normal. Their means and variances are updated with each quarter's observations, and their distribution is collapsed to its

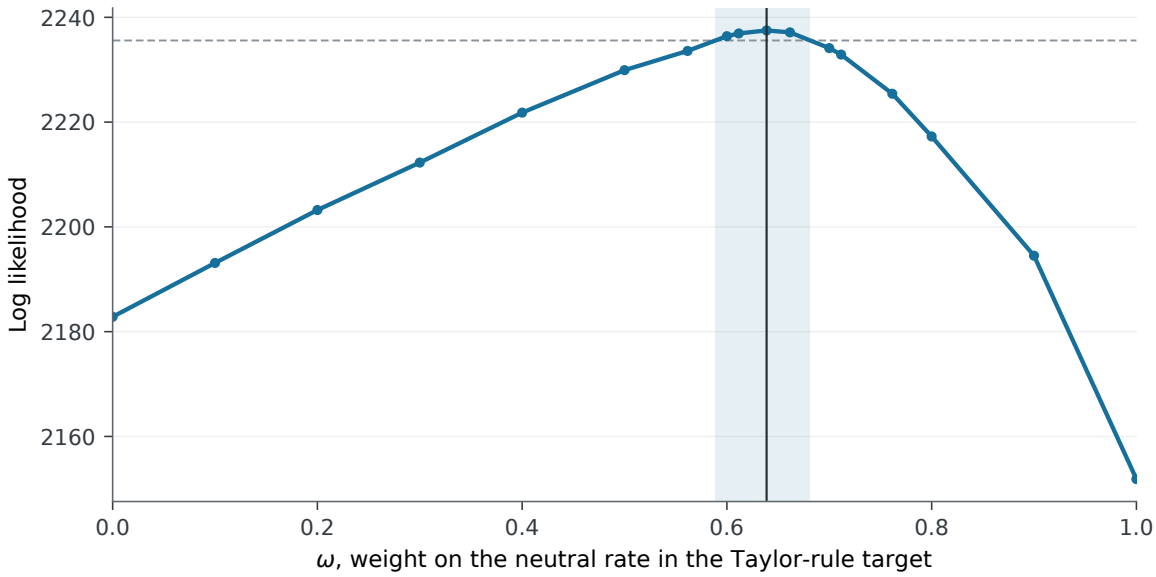


Figure 28: The likelihood and the weight on the neutral rate

Notes: Log likelihood at each value of ω, holding the other parameters at their estimates. The vertical line marks the estimate, the dashed line the maximum less 1.92, and the shaded band the values of ω within it.

first two moments for each value of the states before it is carried forward. The variance of each observation also includes the error from approximating the corresponding model quantity. The growth equation treats the previous quarter's GDP measurement error as a new draw. This approximation simplifies the filter, and it matters little because that error has a standard deviation of only 0.2 percent.

The measurement errors of the CBO gap and the funds rate are bounded below at 10 basis points, the smallest error with which quarterly averages measure end-of-quarter model quantities; both estimates lie at this bound. The noise in the disaster measure is fixed at the approximation error of the disaster-risk state. The innovation of the output discrepancy has an estimated standard deviation of 1.18 percent a quarter, and the measurement error of inflation one of 1.43 percentage points.

Figure 28 shows the log likelihood as a function of ω, holding the other parameters at their estimates. The interval in Table 2 contains the values within 1.92 points of the maximum. Because it holds the other parameters fixed, it is a conditional interval rather than a full profile. The likelihood is nearly flat in the size of the policy shock, whose standard error exceeds its estimate: with shocks of about 5.7 basis points, the data say little about their size.

D.2 Model and data moments

Table 8 compares the estimated model with 21 moments of the data over 1994Q4–2025Q2. Of the 21 moments, 12 lie within two bootstrap standard deviations of the data. Among the asset-price moments, only the volatility of the safe return lies outside this range: the model's safe return is much smoother than in the data. The mean safe return is higher in the model, 0.28 against 0.09 percent a quarter, but within the sampling uncertainty of the data. The table also reports the mean dividend yield, which is not among the 21 moments. It is much higher in the model, about 6.2 percent a year against 1.8 percent in the data, so the model values equity low relative to its dividends. The main shortfall concerns the dynamics of the policy rate. The estimated rule adjusts quickly, so the model's funds rate is less persistent than in the data, 0.91 against 0.98, and its changes are less autocorrelated, 0.39 against 0.69. The model's funds rate also moves less overall and more closely with inflation. Inflation is less volatile in the model than in the data, with a standard deviation of 0.27 against 1.46 percentage points, and the measurement error in inflation absorbs most of the variation in measured inflation.

D.3 The onset of the financial crisis

Section 5.4 fits a single disaster-risk shock to the onset of the financial crisis. The shock shifts log disaster intensity by δ from a date τ in 2007Q3 onwards and then decays with a two-year half-life. Starting from the 2007Q2 state with the funds rate at 5.25%, I choose $\delta \in [0, 0.80]$ and $\tau \in [0, 0.5]$ years to minimise the sum of squared differences between the six observed and model changes, in detrended private output, $\log(C + I)$, and the federal funds rate from 2007Q2 to each of the next three quarters, each scaled by the standard deviation of the observed changes in that variable. The model changes are quarterly averages of the expected responses under the estimated rule, with every parameter at its estimate. The counterfactual applies the fitted shock from the same state under the same rule with $\omega = 0$.

D.4 Mean-preserving news and a realised disaster

Figure 29 repeats Figure 10 for the mean-preserving disaster-risk shock, which raises productivity growth between disasters to keep expected growth unchanged. The neutral rate falls by 96 basis points on impact, and the responses have the same shape as under the shock with a fixed drift. Under the standard Taylor rule, the output gap falls by 0.39 log percentage points on impact, against 0.20 under the estimated rule.

Table 8: Model and data moments

Moment	Data	Model	Distance
Equity return, SD	8.75	8.80	0.06
Safe return, mean	0.09	0.28	1.12
Equity excess return, mean	1.81	1.95	0.20
Safe return, SD	0.67	0.19	−4.43
Safe return, autocorrelation	0.89	0.89	0.11
GDP growth, SD	1.17	1.06	−0.27
GDP cycle, SD	1.31	1.56	0.81
Inflation, mean	2.17	2.05	−0.43
Inflation, SD	1.46	0.27	−3.22
Inflation, autocorrelation	0.69	0.92	1.60
Disaster measure, SD	5.81	5.09	−0.72
Disaster measure, autocorrelation	0.73	0.82	2.06
Funds rate, mean	2.55	3.16	0.92
Funds rate, SD	2.26	0.98	−5.91
Funds rate, autocorrelation	0.98	0.91	−8.03
Funds rate, fourth autocorrelation	0.79	0.52	−2.77
Funds-rate change, autocorrelation	0.69	0.39	−4.42
Growth and funds rate, correlation	0.11	0.04	−0.66
Funds rate and inflation, correlation	−0.00	0.87	4.20
Rate change and lagged growth, correlation	0.14	0.01	−1.14
Rate change and lagged inflation, correlation	0.33	−0.09	−2.67
Dividend yield, mean log	−4.01	−2.78	26.25

Notes: Data cover 1994Q4–2025Q2. Means and standard deviations are in percent: returns and GDP growth are quarterly, inflation and the funds rate annual; the GDP cycle is HP-filtered log GDP. The distance is the difference between model and data divided by the bootstrap standard deviation of the data moment. Model moments are computed as in Table 3. The dividend yield, the log of trailing twelve-month real dividends divided by the current real ex-dividend price, is not among the 21 moments; its data cover 1995Q3–2025Q2, and its distance uses a moving-block bootstrap of the data mean.

Figure 30 shows the responses to a realised disaster from the steady state, which lowers productivity by 25% and raises the disaster intensity from 3.0% to 4.55% a year. The expected recovery of productivity offsets the higher intensity, so the neutral rate changes little on impact, and it is 43 basis points higher after two years. A rule that does not track it therefore leaves policy too loose. Under the standard Taylor rule, inflation rises by 88 basis points on impact and is still 87 basis points above target after five years, against 39 and 33 basis points under the estimated rule. The output gap barely moves under any rule, because the disaster lowers productivity, and with it the natural level of output, as well as actual output.

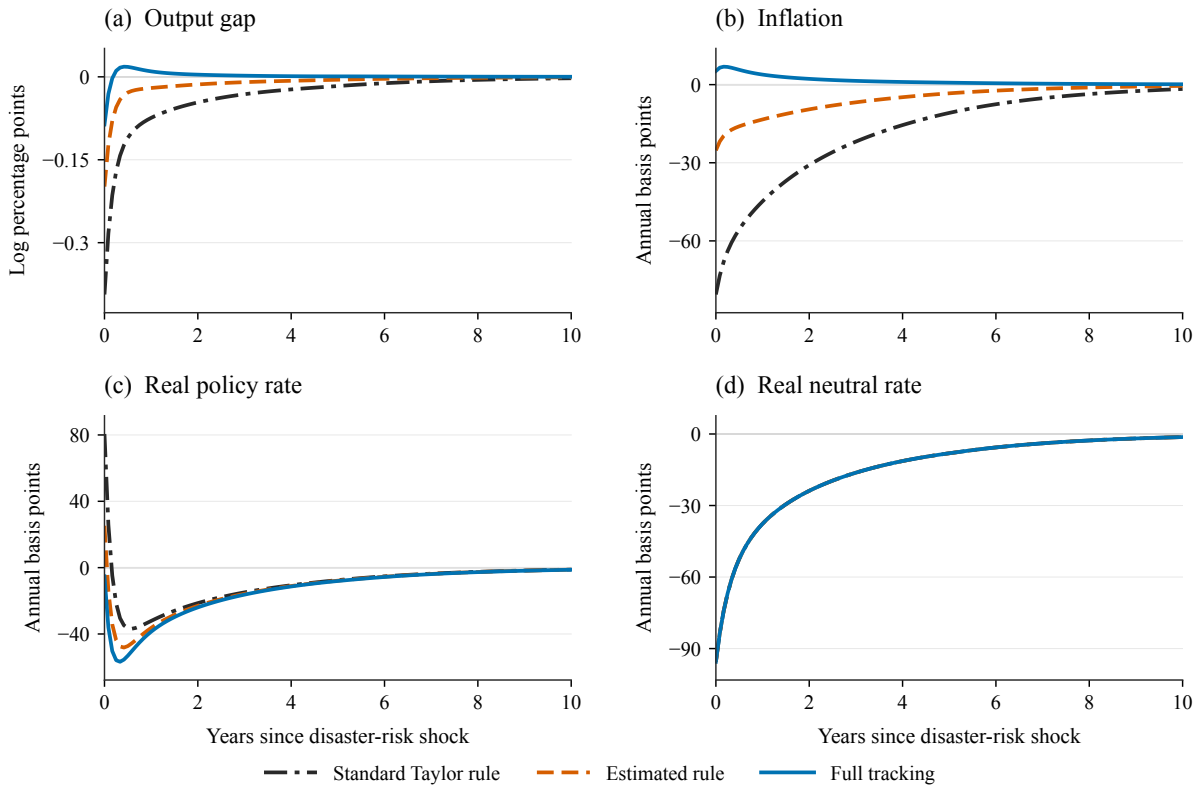


Figure 29: Monetary responses to a mean-preserving disaster-risk shock

Notes: As in Figure 10, except that productivity growth between disasters rises to keep expected growth unchanged.

D.5 Comparisons with other measures

I first explain why the two risk-adjusted neutral rates differ in the financial crisis. I then compare them with other short-run and longer-horizon estimates of the US neutral rate, and with measures of monetary policy. The comparison runs from 1995Q1 to 2025Q2.

The model and the semi-structural estimate. The two risk-adjusted estimates differ most when the option-based probability is far above its mean (Figure 14). In the model, the disaster intensity cannot exceed $\lambda_{\max} = 5\%$ a year. With productivity and trend growth at the steady state, the intensity reaches this level once the option-based probability exceeds 11.0%. As the measure rises from its mean to that level, the model's neutral rate falls from 1.18% to −0.13%, or by 27.6 basis points per percentage point on average, the loading reported in Section 5.6. Beyond it, the model's neutral rate barely responds to the measure, while the semi-structural loading applies at any level. The option-based probability reached 36.6% in 2008Q4, which is one reason the semi-structural estimate fell

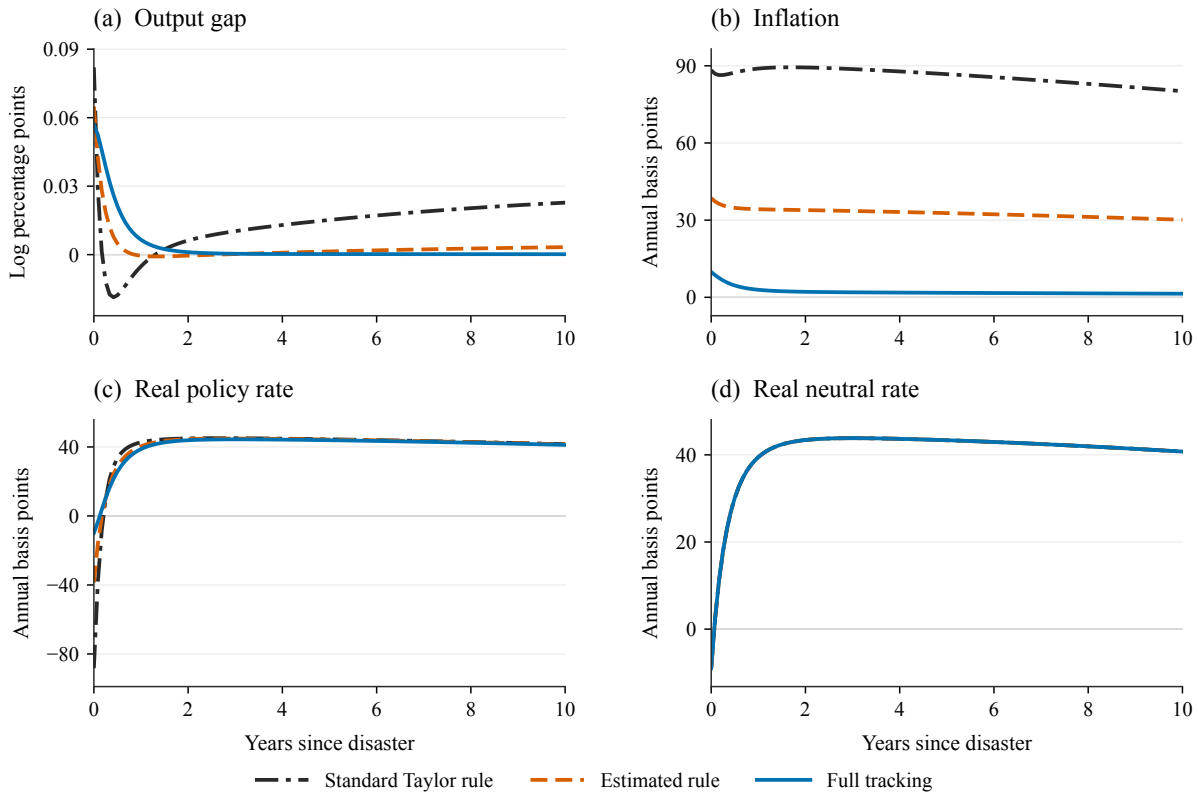


Figure 30: Monetary responses to a realised disaster

Notes: A disaster at date zero lowers productivity by 25% from the steady state. Responses average over subsequent shocks and are measured against each rule's own path without the disaster. Units are as in Figure 10. Full tracking sets $\omega = 0.99$.

much further than the model's.

Short-run neutral rates. Figure 31 compares my estimates with the short-run neutral rates of Cúrdia et al. (2015) and Del Negro et al. (2017b), and with the Laubach–Williams estimate. All four short-run rates turn negative in the financial crisis, while the Laubach–Williams estimate stays positive. In 2008Q4, the semi-structural rate is −5.76% and the model's rate −1.25%, against −2.49% for Cúrdia et al. and −2.59% for Del Negro et al. The Laubach–Williams estimate is still 0.73%. My estimates therefore agree with these structural benchmarks that the short-run neutral rate fell below zero in the crisis, but not on how far it fell.

Differences are larger in the pandemic. In 2020Q2, the two risk-adjusted rates are −1.45% and −1.18%, whereas the updated Del Negro et al. series falls to −30.80%, far outside the range of the other estimates. The Cúrdia et al. series ends in 2016 and does not

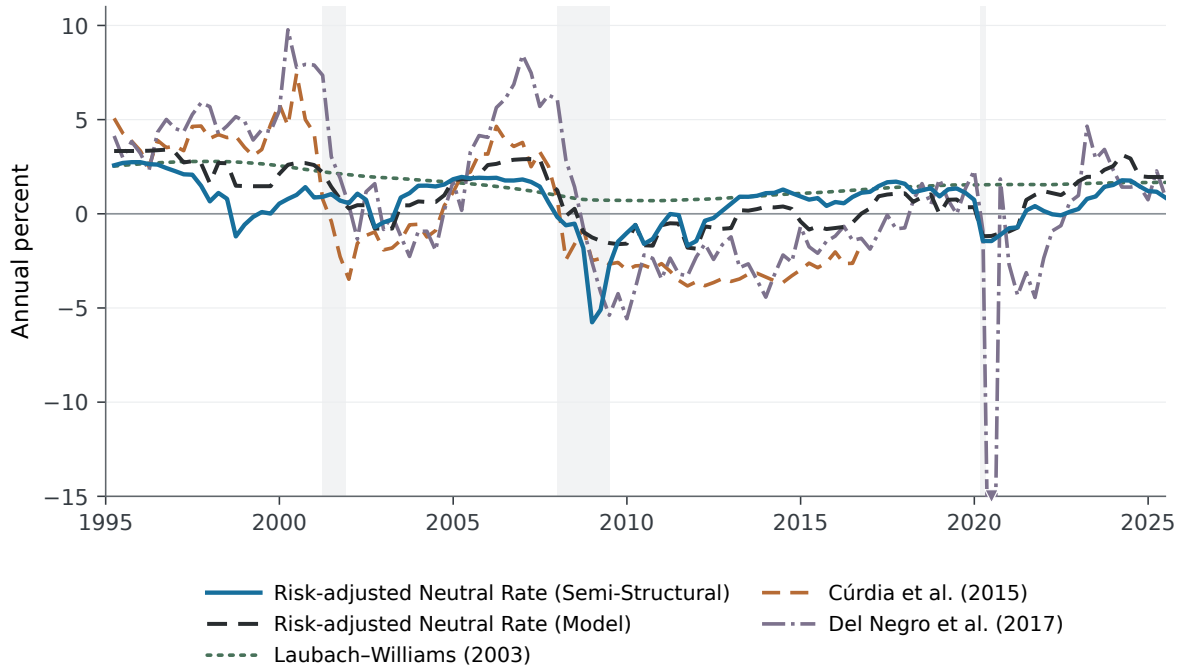


Figure 31: Risk-adjusted and published short-run neutral rates

Notes: Annual percent, 1995Q1–2025Q2. Recessions are shaded. My estimates are those of Figure 14. Cúrdia et al. is the authors' 2016 update, which ends in 2016Q3. Del Negro et al. is the May 2026 update; its 2020Q2 value of −30.80% lies below the plotted range. Laubach–Williams is the current two-sided estimate.

cover the pandemic. The estimates also differ in what they measure. Del Negro et al. focus on the convenience yield that investors attach to the safety and liquidity of Treasuries, and Cúrdia et al. estimate the efficient rate of interest, the real rate in a counterfactual economy free of nominal rigidities and other distortions. My estimates instead show how risk moves the neutral rate.

Long-run rates. Figure 32 compares my estimates with longer-horizon benchmarks: the US estimate of Holston et al. (2017), the trend rate of Del Negro et al. (2017b), and the estimates of Lubik and Matthes (2015) and Christensen and Rudebusch (2019). These benchmarks describe where the real rate is expected to settle once temporary shocks have passed. The Lubik–Matthes estimate, for example, is a five-year-ahead real rate, and the Christensen–Rudebusch estimate is the expected average short real rate five to ten years ahead, inferred from inflation-indexed debt. My estimates instead measure a short-run neutral rate, which moves with current risk, so they diverge from these benchmarks whenever risk is high.

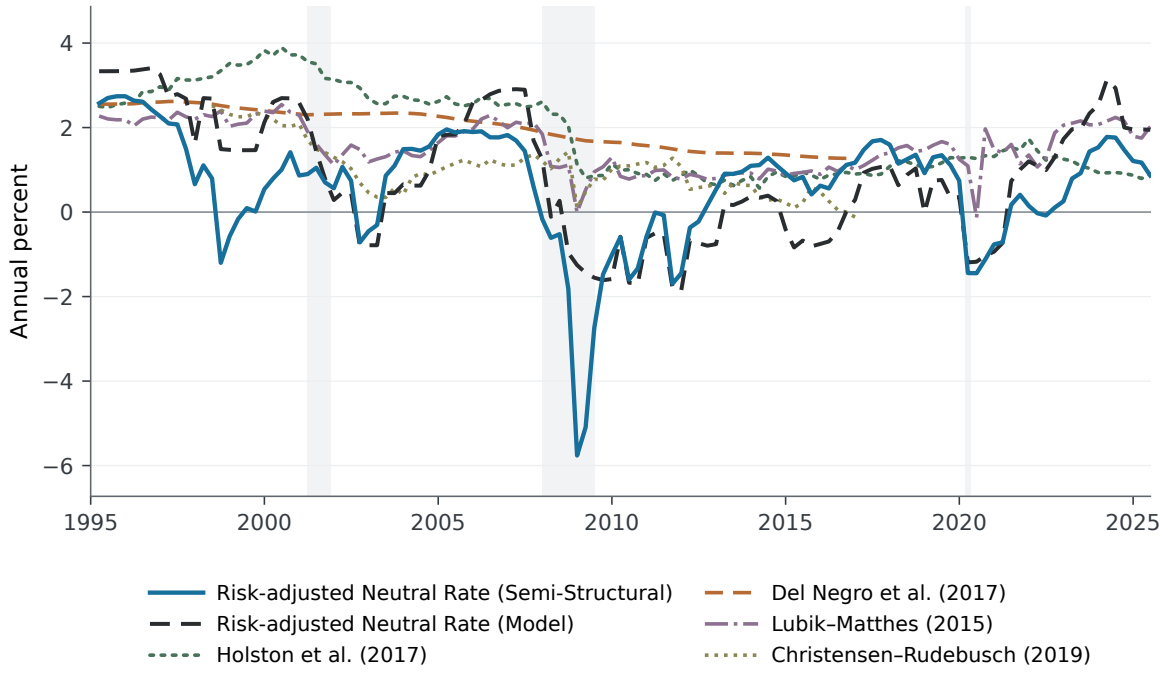


Figure 32: Risk-adjusted neutral rates and longer-horizon benchmarks

Notes: Annual percent, 1995Q1–2025Q2. Recessions are shaded. My estimates are as in Figure 31. Holston et al. is the current US estimate, with the COVID adjustment of Holston et al. (2023); Del Negro et al. is the 2017 trend estimate, which ends in 2016Q4; Lubik–Matthes uses the revised prior of Lubik and Matthes (2023); and Christensen–Rudebusch is available over 1998Q2–2016Q4.

The gap persists after the financial crisis. Over 2010–2019, the semi-structural rate averages 0.58% and the model's rate −0.14%, against 0.92% for Holston et al. and 1.06% for Lubik–Matthes. The model's rate stays near or below zero through much of the recovery, even though its crisis trough is shallower than that of the semi-structural estimate.

Monetary-policy measures. Figure 33 compares my neutral rates with two measures of the policy setting, both converted into real rates with the same inflation adjustment as in Section 5.2. The federal funds rate is the conventional proxy, but it stops at zero. The shadow rate of Wu and Xia (2016) also picks up asset purchases and forward guidance at the lower bound. Both measures describe the setting of policy, whereas my estimates describe the benchmark against which that setting is judged. The gap between the two gives the stance. In the financial crisis, both measures lie above the semi-structural estimate, so policy was contractionary relative to it on either measure.

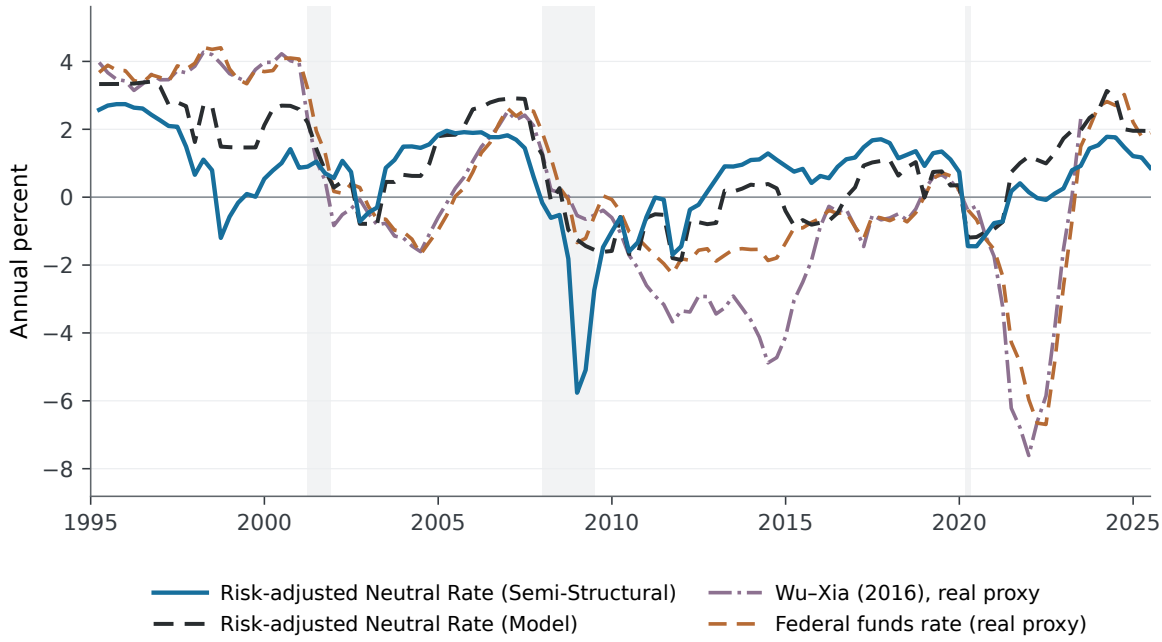


Figure 33: Risk-adjusted neutral rates and real-policy proxies

Notes: Annual percent, 1995Q1–2025Q2. Recessions are shaded. My estimates are as in Figure 31. Each real policy rate is the quarterly mean nominal rate minus four-quarter trailing GDP-deflator inflation. The Wu–Xia shadow rate ends in 2023Q2.

D.6 The semi-structural model

The state-space model of Section 5.6 builds on Laubach and Williams (2003). Its state vector is $s_t = (y_t^\star, g_t^\star, z_t, \widetilde{\sigma}_t^2, \widetilde{\lambda}_t, x_t)'$: log potential output, annual trend growth, a persistent residual component of the neutral rate, the option-implied variance and disaster indices, and the output gap in percentage points. The neutral rate is $r_t^\star = \alpha_L + g_t^\star + z_t - \theta_\lambda \widetilde{\lambda}_t$, as in (24); the variance index is filtered but does not enter the rate. The tildes distinguish the option-based indices from the physical risk of the structural model. In this linear Gaussian system, they are unbounded indices rather than probabilities. The states follow

$$\begin{aligned} y_t^\star &= y_{t-1}^\star + g_{t-1}^\star/4 + \eta_t^y, & g_t^\star &= g_{t-1}^\star + \eta_t^g, \\ z_t &= \rho_z z_{t-1} + \eta_t^z, & \widetilde{\lambda}_t &= 0.002 + \rho_\lambda \widetilde{\lambda}_{t-1} + \eta_t^\lambda, \end{aligned} \tag{38}$$

with the variance index following the same first-order process around a long-run mean of 0.01. Aggregate demand follows the IS equation

$$x_t = a_y x_{t-1} + 100 a_r (r^{\text{real}}_{t-1} - r^{\star}_{t-1}) + \eta^x_t, \qquad a_y = 0.8, \quad a_r = -0.4, \tag{39}$$

in which a real rate above the neutral rate lowers the gap in the next quarter. The real-rate input is the quarterly mean funds rate minus the four-quarter trailing mean of GDP-deflator inflation, the proxy of Section 5.2. I fix $\rho_z = 0.995$ and $\rho_\lambda = \rho_\sigma = 0.9$, so the long-run mean of the disaster index is 0.02, and the intercept is $\alpha_L = 0.0168 + 0.02\theta_\lambda$, which holds the neutral rate at the indices' long-run means fixed when the loading changes.

The observations are log GDP, the change in GDP-deflator inflation, the squared VIX, the option-based disaster probability and log CBO potential output, each with its own measurement error, and the filter uses only the entries observed at each date. Log GDP equals potential plus the gap, and the two indices are observed with noise. The change in inflation equals κx_t with $\kappa = 0.0015$, so constant inflation is consistent with a zero gap. I estimate θ_λ by Gaussian maximum likelihood over 1986Q1–2026Q2, together with seven standard deviations: those of the innovations to potential output, the gap and the two indices, and of the measurement errors in inflation and the two indices. The innovation standard deviations of trend growth and the residual component, and the measurement errors of GDP and CBO potential, are fixed. I report smoothed estimates. The 90 percent interval for θ_λ in Section 5.6 is a model-based bootstrap. The bands in Figure 14 are conditional on the estimated parameters and exclude parameter and specification uncertainty. The likelihood does not use the published Laubach–Williams estimate.

E Main Derivations

This appendix gives the main steps behind the results of Sections 2 and 4. The full arguments, with the regularity conditions and the general-EIS case, are in the Online Appendix.

E.1 The two-period model

The equity premium and the neutral rate. With the preferences in (3), the household's stochastic discount factor is

$$\mathcal{M} = \frac{e^{-\rho} c_0 c_1^{-\gamma}}{\mathbb{E}[c_1^{1-\gamma}]},$$

and both assets satisfy $\mathbb{E}[\mathcal{M} e^{r^m}] = \mathbb{E}[\mathcal{M}] e^{r^f} = 1$. In equilibrium $c_1 = z_1 = \mathcal{Q} e^{r^m}$. Dividing one pricing condition by the other removes the scale of $\mathcal{M}$ and leaves a relation between exponential moments of the equity return,

$$\mathbb{E}[e^{(1-\gamma) r^m}] = e^{r^f}\, \mathbb{E}[e^{-\gamma r^m}].$$

Taking logarithms and using $c_{r^m}(\theta) = c_G(\theta) - \theta \log \mathcal{Q}$ gives the safe rate, $r^f = c_G(1-\gamma) - c_G(-\gamma) - \log \mathcal{Q}$. Subtracting it from the log expected equity return, $\log \mathbb{E}[e^{r^m}] = c_G(1) - \log \mathcal{Q}$, gives the premium $\mathcal{W} = c_G(1) + c_G(-\gamma) - c_G(1-\gamma)$ and the price $\log \mathcal{Q} = g - r^f - \mathcal{W}$ of Theorem 1. Output then follows from (5): $\log y_0 = \rho + \log \mathcal{Q} = (g + \rho - \mathcal{W}) - r^f$, so output is at potential when $r^f = r^\star = g + \rho - \mathcal{W}$, as Proposition 2.1 states.

Two properties of the premium follow directly. First, writing each CGF as its cumulant series and adding the three evaluations gives $\mathcal{W} = \sum_{n \geq 2} \Upsilon_n \kappa_n$ with $\Upsilon_n = [1 + (-\gamma)^n - (1-\gamma)^n]/n!$, the expansion in (7). Because the first-order term cancels, a shift in mean growth does not change the premium. Second, a CGF is convex, so its derivative is increasing and

$$\mathcal{W} = \int_0^1 \left[c_G'(s) - c_G'(s-\gamma) \right] ds \geq 0,$$

with strict inequality whenever output growth is risky. Online Appendix A.4 gives the moment conditions and the equilibrium selection behind the pricing equalities.

E.2 The quantitative model

The neutral rate (Proposition 4.1). Let $V_t = \exp(a_t^p + v_t)$ denote certainty-equivalent consumption, so that continuation utility is $U_t = V_t^{1-\gamma}/[(1-\gamma)\rho]$. With a unit elasticity of

intertemporal substitution, the aggregator of Section 4.1 implies the stochastic discount factor

$$M_t = \exp\Big(\int_0^t f_U\, du\Big)\frac{C_0}{C_t}\Big(\frac{V_t}{V_0}\Big)^{1-\gamma},$$

where f_U is the derivative of the aggregator with respect to continuation utility. Marginal utility therefore depends on continuation value as well as on consumption. When $\gamma > 1$, a state in which V_t is low is valued more highly than its consumption alone implies. Because households can trade the safe bond, the drift of M_t must equal $-r_t M_t$. Applying Itô's formula with jumps to M_t, and using the equation for v_t to eliminate f_U, gives the safe rate at any allocation,

$$r_t = \rho + \frac{\mathbb{E}_t[dc_t]}{dt} - \tfrac{1}{2}\sigma_{c,t}^2 - (\gamma-1)\sigma_{c,t}\sigma_{v,t} - \lambda_t\Big[J_{c,t} + e^{(1-\gamma)J_{v,t}}(e^{-J_{c,t}} - 1)\Big], \tag{40}$$

where $\sigma_{c,t}$ and $\sigma_{v,t}$ are the exposures of log consumption and of v_t to ordinary productivity shocks, and $J_{c,t}$ and $J_{v,t}$ are their changes in a disaster. The expected change in consumption includes the expected loss from disasters, $\lambda_t J_{c,t}$, which the last bracket removes. At the natural allocation, $C_t = A_t$, so $\sigma_c = \sigma_A$, $J_c = J_A$, $\sigma_v = \sigma_A \partial_{a^r} v_t^n$ and $J_v = \Delta_J v_t^n$, and expected consumption growth equals expected productivity growth. Substituting these into (40) gives (17) with the risk adjustment (18). The same identity also yields the IS relation (21), as the difference between (40) at the actual allocation, where $c_t = \log A_t + x_t - \log \Psi_t$, and at the natural allocation. Online Appendix D.5 gives the equation for v^n and the proof that it is linear in trend growth.

The Phillips curve. Firm j chooses its own inflation rate $\pi_t(j)$, which moves its relative log price $q_t = \log(P_t(j)/P_t)$ at rate $\pi_t(j) - \pi_t$. Its value $\mathcal{J}$ solves a Bellman equation whose flow payoff is real profits net of adjustment costs, $Y_t[e^{(1-\varepsilon)q} - mc_t e^{-\varepsilon q} - (\pi(j) - \pi^\star)^2/(2\eta)]$, where mc_t is real marginal cost net of the subsidy, and in which future payoffs are valued with M_t. The first-order condition for $\pi(j)$ equates the marginal adjustment cost to the marginal value of a higher relative price, $\mathcal{J}_q = Y_t(\pi_t(j) - \pi^\star)/\eta$. Differentiating the Bellman equation with respect to q and imposing symmetry, $q = 0$ and $\pi(j) = \pi$, gives

$$r_t Y_t \widehat{\pi}_t = \frac{\mathbb{E}_t^Q[d(Y_t\widehat{\pi}_t)]}{dt} + \eta Y_t\Big[\varepsilon\, mc_t - (\varepsilon - 1)\Big],$$

where $\mathbb{E}_t^Q$ is the expectation under the pricing measure defined by M_t. Under this measure, $Y_t\widehat{\pi}_t$, which is proportional to the marginal value of a higher relative price, is priced like an asset that earns the safe rate and pays a flow proportional to the gap between marginal

cost and its flexible-price level.

Two substitutions turn this equation into (22). The first changes the measure. Writing $Y_t = \Psi_t C_t$ and dividing by C_t converts discounting with M_t into the continuation-utility-weighted expectation $\mathbb{E}_t^v$ of the main text: $\Xi_t = e^{\rho t} M_t C_t / C_0$ is a martingale whose increments load only on the shocks to v_t, so it is the change of measure that weights states by continuation utility. The second replaces marginal cost. Labour supply (14), the subsidy and $Y_t/A_t = L_t = e^{x_t}$ give $\frac{\varepsilon}{\varepsilon-1} mc_t = e^{(1+\varphi)x_t}/\Psi_t$, so the flow becomes $\eta(\varepsilon-1)[e^{(1+\varphi)x_t} - \Psi_t]$. Around the natural allocation, $\Psi_t \simeq 1$ and $e^{(1+\varphi)x_t} - 1 \simeq (1+\varphi)x_t$, which gives the slope $\kappa = \eta(\varepsilon-1)(1+\varphi)$ in (23). Online Appendix D.6 gives the Bellman equation and the change of measure in full.

The welfare measure. Households' flow of utility from consumption and hours is $s_t = \log C_t - L_t^{1+\varphi}/(1+\varphi)$. In the natural allocation, $L_t = 1$ and $C_t = A_t$, so $s_t^n = \log A_t - 1/(1+\varphi)$; under a rule, $C_t = A_t e^{x_t}/\Psi_t$ and $L_t = e^{x_t}$. The flow loss relative to the natural allocation is therefore

$$d_t = s_t^n - s_t = \frac{e^{(1+\varphi)x_t} - 1}{1+\varphi} - x_t - \log\Big(1 - \frac{\widehat{\pi}_t^2}{2\eta}\Big) \geq 0, \tag{41}$$

which is zero only when both gaps are zero. The first two terms measure the labour-market distortion, and the last term measures the resources used in price adjustment. With a unit elasticity of intertemporal substitution, scaling consumption by a constant b at every date and state adds $\log b$ to s_t and to v_t, so certainty-equivalent consumption is homogeneous of degree one in consumption. The share of natural consumption that households would give up to be indifferent between the natural allocation and the rule is therefore

$$\text{CE} = 1 - \exp\{-(v^n - v)\},$$

the measure reported in Table 5, with v^n and v evaluated at the same state. Expanding (41) to second order gives $d_t \simeq \frac{1}{2}(1+\varphi)x_t^2 + \widehat{\pi}_t^2/(2\eta)$, the flow loss displayed in Section 5.5. To this order, the lifetime loss is the discounted sum of these squared gaps, with states weighted by continuation utility in the natural allocation. Risk aversion therefore changes the weight on bad states, but not the relative weight on inflation and output gaps. Online Appendix D.13 gives the exact loss equation and the approximation error.