

# Uncertainty Shocks are Downside Severity Shocks

Todd B. Walker\*

March 2026

## Abstract

The CBOE Volatility Index (VIX) is a widely used measure of market uncertainty that has become a foundational component of dynamic macroeconomic models. This paper shows that a simpler, more transparent metric—the normalized price of an at-the-money forward put option or Downside Protection Premium (DPP)—tracks the VIX almost perfectly (0.99 monthly correlation) and permits a direct decomposition of uncertainty into the probability of a market shortfall and the conditional severity of that shortfall. Using an entropic tilting procedure to recover the risk-neutral density, I document that nearly all time-series variation in downside risk during crises is driven by changes in the severity of shortfall, and not by changes in the probability of shortfall. Vector autoregressions demonstrate that standard macro “uncertainty shocks” are, in fact, structural tail-severity shocks. This finding suggests that theoretical macro-finance models should shift away from symmetric stochastic volatility shocks and toward time-varying disaster-severity structures.

---

\*Department of Economics, Indiana University, [walkertb@iu.edu](mailto:walkertb@iu.edu)

## 1 INTRODUCTION

The CBOE Volatility Index (VIX/VXO) is the benchmark market-based measure of uncertainty in macroeconomics, and is routinely used to identify uncertainty shocks, discipline models, and summarize financial stress. Yet the VIX is difficult to interpret. It is constructed from a strip of out-of-the-money options and quoted in annualized volatility units; changes in the VIX do not map transparently into economic fundamentals like the cost of insuring against losses, the likelihood of a bad state, or the expected size of that loss. Moreover, the index depends on a wide cross-section of option quotes that can make it sensitive to pricing distortions in relatively illiquid tail strikes, especially during stress episodes. For example, on August 5, 2024, the VIX spiked over 180% to nearly 66 just prior to market open. As Todorov and Vilkov (2024) note, this spike was largely driven by dealer quote-setting behavior and widening bid-ask spreads, highlighting that “the calculation of VIX makes it vulnerable to a widening of bid-ask spread quotes regardless of a fundamental rise in underlying volatility.”

This paper proposes a simple alternative: the price of a 30-day at-the-money-forward S&P 500 put divided by the index level. I call this object the *Downside Protection Premium* (DPP). The DPP has a direct economic interpretation. It is the one-month cost of buying downside insurance on one dollar of equity exposure against finishing below the forward benchmark. It is constructed from the most liquid part of the put surface, and because it is a ratio, it is readily comparable across time and assets.

The paper makes two contributions.

First, I show that, to a macroeconomist, the information content of the DPP is identical to that of the VIX. Once aggregated to monthly frequency, the DPP contains essentially the same information as the VIX for identifying macro uncertainty shocks. Over 1990–2025, the correlation between the monthly 30-day DPP and the VIX is 0.993, with similarly strong comovement at higher frequencies. Bloom-style (2009) uncertainty shocks identified from monthly DPP data are nearly identical to those obtained from the VIX. Likewise, in a standard four-variable uncertainty VAR, replacing the VIX with the DPP leaves the impulse responses of unemployment, inflation, and short-term interest rates virtually unchanged. For the purposes of the macro uncertainty literature, the DPP delivers the uncertainty proxy as the VIX.

The second contribution is interpretive. Because the DPP is the price of a forward-ATM put, it admits a transparent decomposition into two objects: the risk-neutral probability of finishing below the forward benchmark and the expected proportional shortfall conditional on that event. To quantify these components, I use the minimum-relative-entropy approach of Buchen and Kelly (1996) and Stutzer (1996) to recover a daily risk-neutral distribution that is closest to a time-varying nonparametric estimate of the physical return distribution while matching the observed downside-insurance price. The main empirical result is that time variation in the DPP is driven overwhelmingly by *conditional shortfall severity*, not by the probability of a shortfall. The shortfall-probability component is smooth, persistent, and only weakly cyclical; the severity component is crisis-sensitive, strongly countercyclical, and closely aligned with the VIX and other benchmark uncertainty measures.

This distinction matters for the interpretation of macro uncertainty shocks. If DPP-based shocks and VIX-based shocks are observationally equivalent in standard VARs, and if DPP movements are largely severity movements, then the uncertainty shocks identified by the literature are better understood as

repricings of downside-tail severity than as broad increases in symmetric volatility. What macroeconomists often label an “uncertainty shock” is *not* an increase in time-varying volatility (as it is typically modeled), but a substantial repricing of severe downside states.

The results have a structural implication. Much of the DSGE uncertainty literature models uncertainty as a shock to the conditional variance of an otherwise symmetric process. That formulation is tractable, but it is too coarse relative to the option-implied evidence in this paper. The relevant time variation in market-based uncertainty is not a diffuse broadening of the entire distribution. It is an asymmetric movement in downside states, concentrated in the expected severity of losses once those states occur. A more appropriate structural representation therefore separates the probability of downside states from their severity, bringing macro “uncertainty shocks” closer to downside-tail or disaster-severity shocks than to purely symmetric stochastic-volatility shocks.

The message of this paper is straightforward: at monthly frequency, the DPP is a transparent substitute for the VIX, and the shocks it identifies are best interpreted as movements in downside-tail severity, which explains the title of the paper.

**1.1 RELEVANCE TO EXISTING LITERATURE.** The VIX remains the standard market-based proxy for uncertainty in macroeconomics. Bloom (2009) treats implied stock-market volatility as a proxy for uncertainty shocks and demonstrates their importance; Basu and Bundick (2017) identify uncertainty shocks via the VXO in a VAR and then use those empirical responses to discipline a New Keynesian model; Leduc and Liu (2016) use the VIX directly as the uncertainty variable in a monthly VAR and show that uncertainty shocks resemble aggregate demand shocks; Ferrara and Guérin (2018) study the macroeconomic effects of high-frequency uncertainty shocks; Caldara et al. (2016) disentangle financial and uncertainty shocks in an SVAR using option-implied volatility and credit conditions; Bhattarai, Chatterjee, and Park (2020) use identified U.S. VIX fluctuations to quantify international spillovers; Lhuissier and Tripier (2021) identify uncertainty shocks via joint rises in the VIX and credit spreads; Caggiano, Castelnovo, and Pellegrino (2017) and Pellegrino, Castelnovo, and Caggiano (2023) emphasize the state dependence of uncertainty shocks; Alfaro, Bloom, and Lin (2024) quantify a finance uncertainty multiplier; and Diercks, Hsu, and Tamoni (2024) study cascading uncertainty shocks.<sup>1</sup> These papers confirm that the VIX/VXO is an excellent instrument for macroeconomic shocks. Several of these papers even emphasize state dependence, but none separately identify probabilities from severity as is done here.

On the structural modeling side, standard DSGE treatments typically operationalize uncertainty as stochastic volatility in which standard exogenous disturbance remains mean-zero and symmetric, with a conditional variance that evolves over time. Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez, and Uribe (2011) model time-varying volatility in the real interest rate faced by a small open economy; Christiano, Motto, and Rostagno (2014) allow the volatility of cross-sectional idiosyncratic uncertainty to fluctuate over time in a monetary DSGE model with a financial accelerator; Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015) estimate tax and spending processes with time-

---

<sup>1</sup>Related work studies uncertainty shocks more broadly rather than using the VIX/VXO directly as the core observable uncertainty variable; see, for example, Berger, Dew-Becker, and Giglio (2020), Bianchi, Kung, and Tirsikh (2023), and Bok, Mertens, and Williams (2025).

varying volatility and study fiscal-volatility shocks in an otherwise standard New Keynesian model; Born and Pfeifer (2014) analyze policy risk in an estimated New Keynesian model with time-varying monetary and fiscal-policy uncertainty; and Cesa-Bianchi and Fernández-Corugedo (2018) study mean-preserving shocks to the variance of aggregate TFP and to the dispersion of entrepreneurs' idiosyncratic productivity in a financial-accelerator DSGE. Related work places uncertainty shocks in sticky-price New Keynesian environments and studies how their transmission depends on nominal rigidities and model structure (Basu and Bundick, 2017; Oh, 2020). In this literature, the primitive uncertainty shock is therefore a shock to conditional variance or dispersion, consistent with a shock to the VIX.

Results herein suggest that this conditional-variance representation is not able to disentangle changes in downside risk from shortfall severity. Most papers assume a symmetric broadening of the distribution via stochastic volatility. I argue that macro uncertainty, as identified using the VIX and DPP, looks more like the time-varying disaster-risk literature of Gourio (2012) and Wachter (2013) (and subsequent work) with an important refinement. Relative to standard disaster-risk formulations that emphasize time variation in the probability of a disaster, the option-implied evidence here points more strongly toward time variation in the severity of bad states once they occur while probabilities are much smoother.

Like this paper, the finance literature contains several examples arguing that decomposing the VIX into various theoretical components has value. Bekaert et al. (2013) decompose the VIX into a component related to expected stock-market volatility and a component more closely associated with risk aversion, while Bekaert and Hoerova (2014) show that the variance premium predicts stock returns whereas conditional stock-market variance predicts economic activity. Bollerslev et al. (2009) quantify the variance risk premium as the gap between option-implied and realized variance and show that it predicts future returns. Drechsler and Yaron (2011) provide an equilibrium model with a time-varying variance premium, and Adrian and Rosenberg (2008) show that short- and long-run volatility components carry distinct prices of risk. Pflueger et al. (2020) show that market-based measures of perceived risk affect discount rates and future real activity, not merely the expected dispersion of outcomes. Taken together, these papers imply that the VIX is not a pure measure of macroeconomic uncertainty because it conflates physical uncertainty with pricing-kernel variation. My contribution is to push that logic further by showing that, within option-implied downside risk itself, crisis variation is driven primarily by changes in shortfall severity rather than by changes in shortfall probability.<sup>2</sup>

## 2 ATM FORWARD PUTS DIVIDED BY UNDERLYING

Consider an at-the-money (AtM) put option, defined as where the strike  $K$  equals the forward price  $F_{t,t+\tau}$  implied by no-arbitrage ( $K = F_{t,t+\tau} = S_t \exp((r_f - d)\tau)$ ), where  $S_t$  is the price of the underlying asset at time  $t$ ,  $r_f$  is the continuously compounded risk-free rate,  $d$  is the continuous dividend and  $\tau$  is time-to-maturity.<sup>3</sup> Let  $P_t^{\text{ATM}}(\tau)$  be the price of this put option. The forward AtM put divided by underlying is given

<sup>2</sup>For related evidence on asymmetry and tail risk, see Segal et al. (2015) and Bollerslev and Todorov (2011).

<sup>3</sup>Defining at-the-money as setting the strike  $K$  equal to the no-arbitrage forward price  $K = F_{t,t+\tau}$  is standard convention for dividend-paying indices such as the SPX (see e.g. Carr and Wu, 2009).

by

$$\chi_t^\tau \equiv \frac{P_t^{\text{AtM}}(\tau)}{S_t} \quad (1)$$

For much of the analysis, I focus on a duration of 30 calendar days denoted in years as  $\tau = 30/365 \approx 0.08219$  in order to make direct comparison to the VIX which also has a 30-day duration.

Unlike the VIX, which is quoted in annualized volatility units, the AtM-Forward put ratio,  $\chi_t^\tau = P_t^{\text{AtM}}(\tau)/S_t$ , is denominated directly as the cost of downside protection per dollar of spot. It is simply the one-month cost of buying downside protection on one dollar of equity exposure. For a \$100 notional at date  $t$ , the monthly premium is  $100 \cdot \chi_t^{30}$ . It reflects both the probability of downside outcomes and the market price assigned to those outcomes. Moreover, because  $\chi_t^\tau$  is a ratio, it is measured in dimensionless units (“percent of spot”) and is therefore directly comparable across time and assets. For these reasons, I refer to this object as the **Downside Protection Premium (DPP)**.

**2.1 A DECOMPOSITION** The ratio (1) has a useful decomposition that will be exploited throughout the paper. The price of a European put option is given by the discounted expected payoff under the risk-neutral measure  $\mathbb{Q}$ ,  $P_t = e^{-r_f \tau} \mathbb{E}_t^{\mathbb{Q}}[\max(K - S_{t+\tau}, 0)]$  where  $S_{t+\tau}$  is the asset price at expiration, and  $\mathbb{E}_t^{\mathbb{Q}}$  denotes the conditional expectation under  $\mathbb{Q}$  given information at time  $t$ . Since the strike of the AtM-forward put is defined by  $K = F_{t,t+\tau} = S_t e^{(r_f - d)\tau}$ , the price over spot is given by

$$\begin{aligned} \chi_t^\tau &\equiv \frac{P_t^{\text{AtM}}(\tau)}{S_t} = e^{-r_f \tau} \mathbb{E}_t^{\mathbb{Q}} \left[ \max \left( \frac{F_{t,t+\tau}}{S_t} - \frac{S_{t+\tau}}{S_t}, 0 \right) \right] = e^{-r_f \tau} \mathbb{E}_t^{\mathbb{Q}} \left[ \max \left( e^{(r_f - d)\tau} - \frac{S_{t+\tau}}{S_t}, 0 \right) \right] \\ &= e^{-r_f \tau} e^{(r_f - d)\tau} \mathbb{E}_t^{\mathbb{Q}} \left[ \max \left( 1 - \frac{S_{t+\tau}}{S_t e^{(r_f - d)\tau}}, 0 \right) \right] = e^{-d\tau} \mathbb{E}_t^{\mathbb{Q}} [(1 - \tilde{R}_t^\tau)^+] \end{aligned}$$

where  $\tilde{R}_t^\tau \equiv S_{t+\tau}/F_{t,t+\tau}$  is the forward-normalized gross return. Thus,  $\chi_t^\tau$  is the discounted risk-neutral expected proportional shortfall of the asset relative to its forward value. Note that the relevant discount factor is  $e^{-d\tau}$ , not  $e^{-r_f \tau}$ , because dividing by  $S_t$  converts the strike into the carry-adjusted forward benchmark. Expanding  $(1 - \tilde{R}_t^\tau)^+ = (1 - \tilde{R}_t^\tau) \mathbf{1}\{\tilde{R}_t^\tau < 1\}$  and assuming  $\tilde{R}_t^\tau$  admits a risk-neutral density yields the following decomposition

$$\chi_t^\tau = e^{-d\tau} \left( \underbrace{\mathbb{Q}_t(\tilde{R}_t^\tau < 1)}_{\text{Prob(Shortfall)}} \underbrace{\mathbb{E}_t^{\mathbb{Q}}[1 - \tilde{R}_t^\tau \mid \tilde{R}_t^\tau < 1]}_{\text{Amt(Shortfall)}} \right) \quad (2)$$

The first term is the risk-neutral probability that the asset finishes below its forward benchmark at horizon  $\tau$  (Prob(Shortfall)) and the second term is the expected proportional size of that shortfall conditional on a shortfall occurring (Amt(Shortfall)).

Thus, an increase in the Downside Protection Premium,  $\chi_t^\tau$ , can occur because market prices imply a higher likelihood of finishing below the forward benchmark, because the expected shortfall in those downside states becomes more severe, or because both components rise simultaneously. It captures not only whether downside states are more likely under  $\mathbb{Q}$ , but also how costly those downside states are expected to be when they occur. Technically, under the risk-neutral distribution, it is the first lower

partial moment around the forward benchmark. Relative to far out-of-the-money put prices or higher-order lower partial moments, it is less sensitive to extreme disaster states and more sensitive to the overall expected downside loss around the forward benchmark.

**2.2 CONSTRUCTING A SYNTHETIC ATM-FORWARD PUT PRICE** To construct the DPP,  $\chi_t^\tau$ , I use option price data on the S&P 500 Index (SPX) sourced from DeltaNeutral (<https://deltaneutral.com/>) that span January 1990 to December 2025. This data consists of end-of-day quotes on all traded strikes and maturities. VIX data are from the Chicago Board Options Exchange (CBOE, <https://www.cboe.com>). The DPP,  $\chi_t^\tau$  for  $\tau = 30/365, 60/365, 90/365$ , and the VIX are derived from the same underlying asset (SPX), allowing for a direct and theoretically consistent comparison. SPX options are European-style, which eliminates the early-exercise premium associated with American-style options and allows for an easier connection to Black-Scholes theory discussed below.

Due to the discrete nature of strikes and time-to-maturity, a precise at-the-money forward put price ( $P_t^{\text{ATM}}$ ) for a target expiration of  $\tau$  must be approximated. I implement a simple two-step interpolation procedure for each quote date. First, options are filtered to require a minimum daily volume of one trade and a minimum of two days to expiry to reduce short-dated noise. A bid-ask spread filter ( $\leq 50\%$ ) is applied to remove extreme quotes and mitigate potential microstructure noise or stale quotes. These filters yield high coverage (8,950 valid days across maturities, or 98% of trading days with coverage increasing with time). Puts are then divided into short-term ( $T_{\text{short}} \leq \tau$ ) and long-term ( $T_{\text{long}} > \tau$ ) subsets based on their time-to-expiration  $T$  (in days). I select the short-term expiration date closest to  $\tau$  (where  $T_{\text{short}} \leq \tau$ ) and the long-term expiration date closest to  $\tau$  (where  $T_{\text{long}} > \tau$ ). The search is constrained to no more than two times the target maturity (60, 120, and 180 days, respectively).

Within each selected expiration date (short and long), I identify two put options that bracket the AtM-forward point for that maturity. Specifically, for each expiration  $T$ , I compute the no-arbitrage forward price  $F_{t,t+T}$  from near-AtM put-call parity using both put and call quotes at the same strike and a daily risk-free rate. This construction absorbs index dividends and ensures that the AtM condition is defined relative to the forward rather than the spot. I then define the forward-minus-strike distance  $\text{FMS}(T) \equiv F_{t,t+T} - K$ . I select one strike with the smallest positive value  $\text{FMS}(T) > 0$  and one strike with the largest negative value  $\text{FMS}(T) < 0$ .

Let  $P_{\text{short}}^+$  and  $P_{\text{short}}^-$  denote the prices of these short-term puts with forward-minus-strike values  $U_{\text{short}}^+ > 0$  and  $U_{\text{short}}^- < 0$ , respectively, where  $U_{\text{short}}^\pm \equiv \text{FMS}(T_{\text{short}})$ . Similarly, let  $P_{\text{long}}^+$  and  $P_{\text{long}}^-$  be the prices for the long-term puts with  $U_{\text{long}}^\pm \equiv \text{FMS}(T_{\text{long}})$ . An estimated AtM-forward price for each expiration date is computed via linear interpolation based on the distance to the forward:  $P_{\text{ATM, short}} = w_{\text{short}}^+ P_{\text{short}}^+ + w_{\text{short}}^- P_{\text{short}}^-$ , where

$$w_{\text{short}}^+ = \frac{|U_{\text{short}}^-|}{|U_{\text{short}}^-| + U_{\text{short}}^+}, \quad w_{\text{short}}^- = 1 - w_{\text{short}}^+$$

and analogously for the long-term expiration.

Using these estimated AtM-forward put prices, the second step forms an AtM put price for target

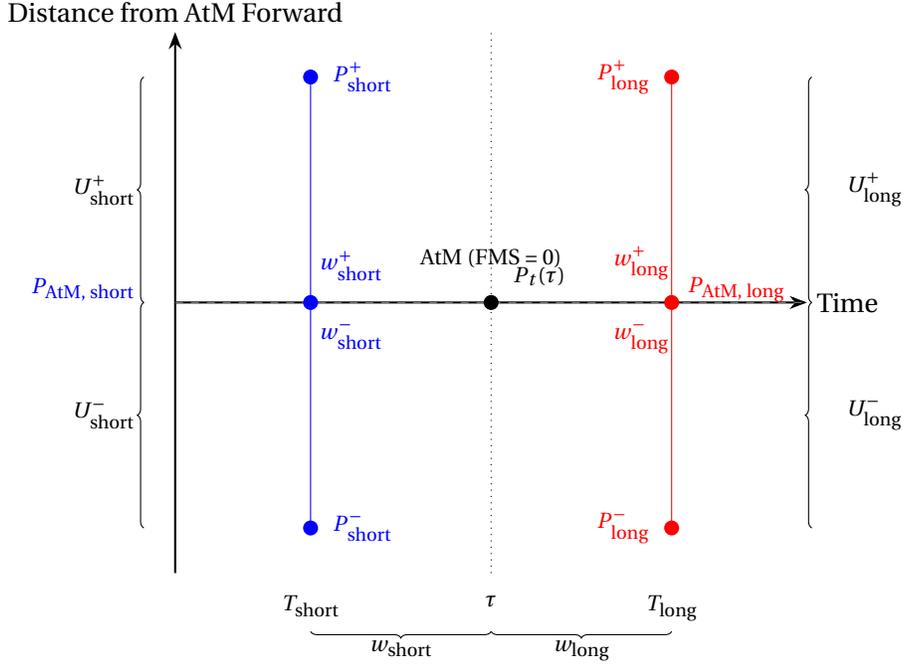


Figure 1: Two-step interpolation procedure for estimating the AtM-forward put price at target expiration  $\tau$ . The first step interpolates between strikes using FMS-based weights ( $w_{\text{short}}^+$ ,  $w_{\text{short}}^-$ ,  $w_{\text{long}}^+$ ,  $w_{\text{long}}^-$ ) to estimate AtM-forward prices for short ( $T_{\text{short}}$ ) and long ( $T_{\text{long}}$ ) expirations. The second step interpolates in time using weights ( $w_{\text{short}}$ ,  $w_{\text{long}}$ ) to estimate the AtM-forward price at  $\tau$ .

maturity  $\tau$ , denoted  $P_t(\tau)$ , via linear interpolation *in time* between  $P_{\text{AtM, short}}$  and  $P_{\text{AtM, long}}$ :  $P_t(\tau) = w_{\text{long}}P_{\text{AtM, long}} + w_{\text{short}}P_{\text{AtM, short}}$ , where

$$w_{\text{long}} = \frac{\tau - T_{\text{short}}}{T_{\text{long}} - T_{\text{short}}}, \quad w_{\text{short}} = 1 - w_{\text{long}}$$

The constructed price  $P_t(\tau)$  is a synthetic AtM-forward put with exactly  $\tau$  days to expiration on date  $t$ . It is important to note that the construction is intentionally straightforward, relying on basic linear interpolation from just four near-AtM put prices and involves no complex volatility surface smoothing. Figure 1 provides a schematic depicting the construction.

**2.3 INCREASING VOLUME OVER TIME** Figure 2 plots the 21-day average put and call volumes for SPX over time, which illustrates a substantial increase in trading activity, particularly since the early 2000s. Put volume (top, red) has surpassed call volume (bottom, blue), with notable spikes during periods of market uncertainty, such as the 2008 financial crisis and the 2020 pandemic. Moreover, the average put volume across relative moneyness is heavily concentrated to near-the-money strikes. The peak of this distribution occurs in the 0% to 2.5% out-of-the-money (OTM) range, which accounts for 23.52% of the total average volume. Cumulatively, over 60% of all put trading activity is situated within the 0% to 10% OtM range, while volume for puts that are slightly in-the-money (-2.5% to 0%) comprises 11.22%. This concentration underscores the central role of near-ATM puts as the primary vehicles for downside pro-

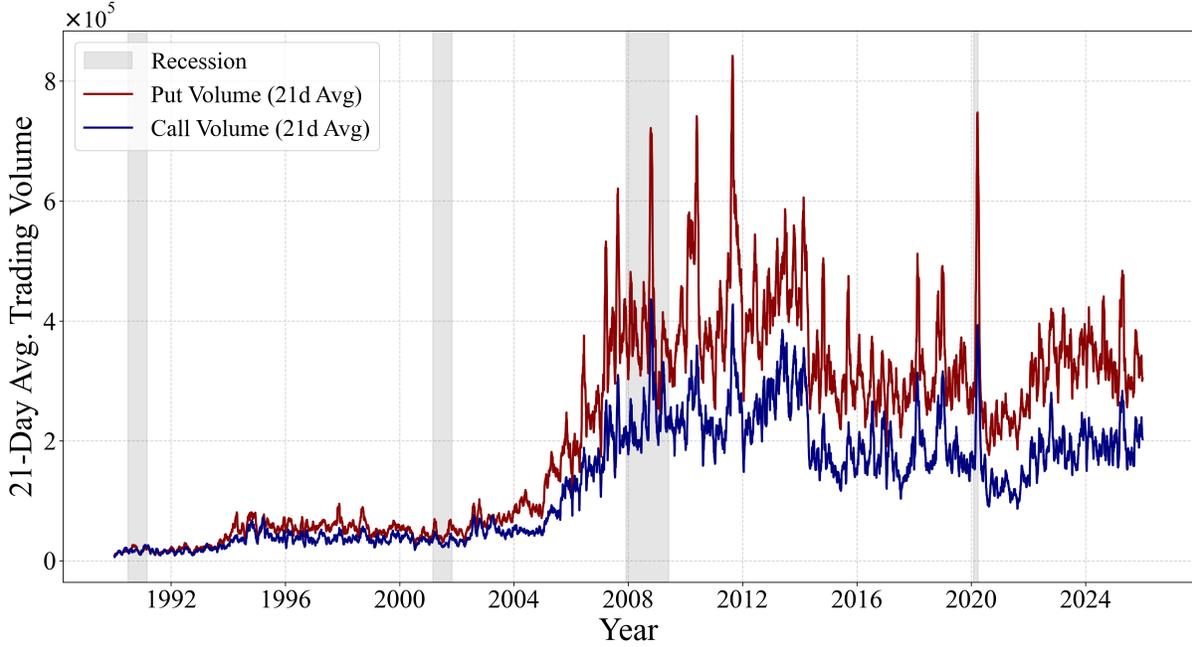


Figure 2: Option Volume: 21-day average put and call volumes over time.

tection, reflecting both maximum sensitivity to implied volatility (Vega) and high liquidity. Conversely, the sharp decline in volume for deeper in-the-money and far out-of-the-money strikes—dropping to just 0.25% for the -20% ItM bin—suggests a significant drop-off in liquidity. Such thin liquidity at distant strikes can exacerbate price shocks and contribute to extreme market phenomena, such as the VIX spike observed on August 5, 2024 (Todorov and Vilkov, 2024).

**2.4 SUMMARY STATISTICS** Table 1 reports summary statistics for the monthly averaged DPP  $\chi_t^\tau$  over the 1990–2025 sample at 30-, 60-, and 90-day horizons. The average DPP scales roughly with the square root of time, consistent with a Black-Scholes approximation discussed in the next section. The mean  $\chi_t^\tau$  rises from 1.85% at 30 days to 2.75% at 60 days and 3.44% at 90 days, yielding ratios  $\chi^{60}/\chi^{30} \approx 1.48$  and  $\chi^{90}/\chi^{30} \approx 1.85$  that are close to the theoretical benchmarks  $\sqrt{2} \approx 1.41$  and  $\sqrt{3} \approx 1.73$ . All three series display substantial right-skewness, driven by rare crisis spikes, and the skewness is more pronounced at short horizons. This suggests that short-dated insurance is particularly exposed to sudden crash risk, whereas longer-dated insurance reflects a smoother assessment of medium-horizon downside conditions. Finally, the series are highly persistent, with AR(1) coefficients rising from 0.85 to 0.90 as maturity increases, downside-insurance premia contain an important slow-moving component.

### 3 CONNECTION TO THE VIX & UNCERTAINTY SHOCKS

This paper has two goals: first, to demonstrate that the Downside Protection Premium (DPP),  $\chi_t^\tau$  of equation (1), has the same information content as the VIX when used to identify macroeconomic shocks at

Table 1: Summary Statistics for the Monthly Downside Protection Premium ( $\chi_t^\tau$ )

Maturity ( $\tau$ )	Mean	Std. Dev.	Min	Median	Max	Skewness	AR(1)
30 Days	0.0185	0.0076	0.0079	0.0166	0.0621	2.0128	0.8531
60 Days	0.0275	0.0103	0.0131	0.0247	0.0843	1.7873	0.8817
90 Days	0.0344	0.0120	0.0175	0.0310	0.0988	1.6870	0.8973

*Note:* The sample spans January 1990 to December 2025. All statistics are calculated from the monthly averaged daily constant-maturity series. The AR(1) refers to the first-order autocorrelation coefficient of the monthly series.

monthly frequency; and second, to provide a more nuanced interpretation of these shocks. This section establishes the first point.

The DPP is tightly linked to short-horizon implied volatility, which helps explain its close empirical connection to the VIX. This connection is especially transparent using the Black–Scholes calculus. Recall the price of a European put option on an index with dividend yield  $d$

$$P_{BS} = Ke^{-r_f\tau}\Phi(-d_2) - S_t e^{-d\tau}\Phi(-d_1)$$

$$d_1 = \frac{\ln(S_t/K) + (r_f - d + \sigma^2/2)\tau}{\sigma\sqrt{\tau}} \quad d_2 = d_1 - \sigma\sqrt{\tau}$$

An AtM option defined relative to the forward,  $K = F_{t,t+\tau} = S_t e^{(r_f-d)\tau}$ , implies  $Ke^{-r_f\tau} = S_t e^{-d\tau}$  and  $d_1 = (\sigma/2)\sqrt{\tau}$ ,  $d_2 = -(\sigma/2)\sqrt{\tau}$ . Dividing by  $S_t$  and using the symmetry  $\Phi(-x) = 1 - \Phi(x)$  gives

$$\chi_t^\tau = e^{-d\tau} [\Phi(d_1) - \Phi(-d_1)] = e^{-d\tau} [2\Phi(d_1) - 1]$$

For short maturities,  $x = d_1 = \sigma\sqrt{\tau}/2$  is small, so the linear approximation<sup>4</sup>  $\Phi(x) \approx 0.5 + x/\sqrt{2\pi}$  implies  $2\Phi(d_1) - 1 \approx \sigma\sqrt{\tau}/\sqrt{2\pi}$ , and hence

$$\chi_t^\tau \approx e^{-d\tau} \cdot \frac{\sigma\sqrt{\tau}}{\sqrt{2\pi}} \quad \sigma_t \approx \chi_t^\tau \cdot \frac{\sqrt{2\pi}}{e^{-d\tau}\sqrt{\tau}} \quad \text{VIX}_t = 100 \cdot \sigma_t \approx 876 \cdot \chi_t^{30} \quad (3)$$

For short horizons ( $\tau \approx 30/365$ ) and typical dividend yields ( $d \approx 1.5\% - 2.5\%$  annualized),  $e^{-d\tau} \approx 0.998 - 0.999$ , so the scaling factor,  $100\sqrt{2\pi}/(e^{-d\tau}\sqrt{\tau})$ , is approximately 876. For example, if  $\chi_t^{30} = 0.02$  (2% of spot), the approximation implies  $\sigma \approx 0.02 \times 8.76 \approx 17.5\%$  annualized, or  $\text{VIX} \approx 17.5$ —close to observed values in low-volatility regimes. This provides a simple parametric mapping between the DPP and the VIX using the Black-Scholes calculus.<sup>5</sup>

The mapping in (3) should be interpreted as a Black–Scholes approximation. The DPP is tied to the

<sup>4</sup>For a forward AtM option,  $d_1 = -d_2$ . The Taylor expansion of the standard normal CDF around  $x = 0$  is

$$\Phi(x) = \frac{1}{2} + \frac{x}{\sqrt{2\pi}} - \frac{x^3}{6\sqrt{2\pi}} + \mathcal{O}(x^5).$$

At  $\tau \approx 0.0822$  and  $\sigma = 0.40$ , the higher-order terms are negligible for the normalized price.

<sup>5</sup>The relationship  $\chi_t^\tau \approx 0.4\sigma\sqrt{\tau}$  also mirrors the well-known market rule of thumb that an AtM straddle is worth approximately  $0.8\sigma\sqrt{\tau}S_t$ . Since  $P \approx C$  at the forward,  $\chi_t^\tau$  is essentially half the value of a normalized forward straddle.

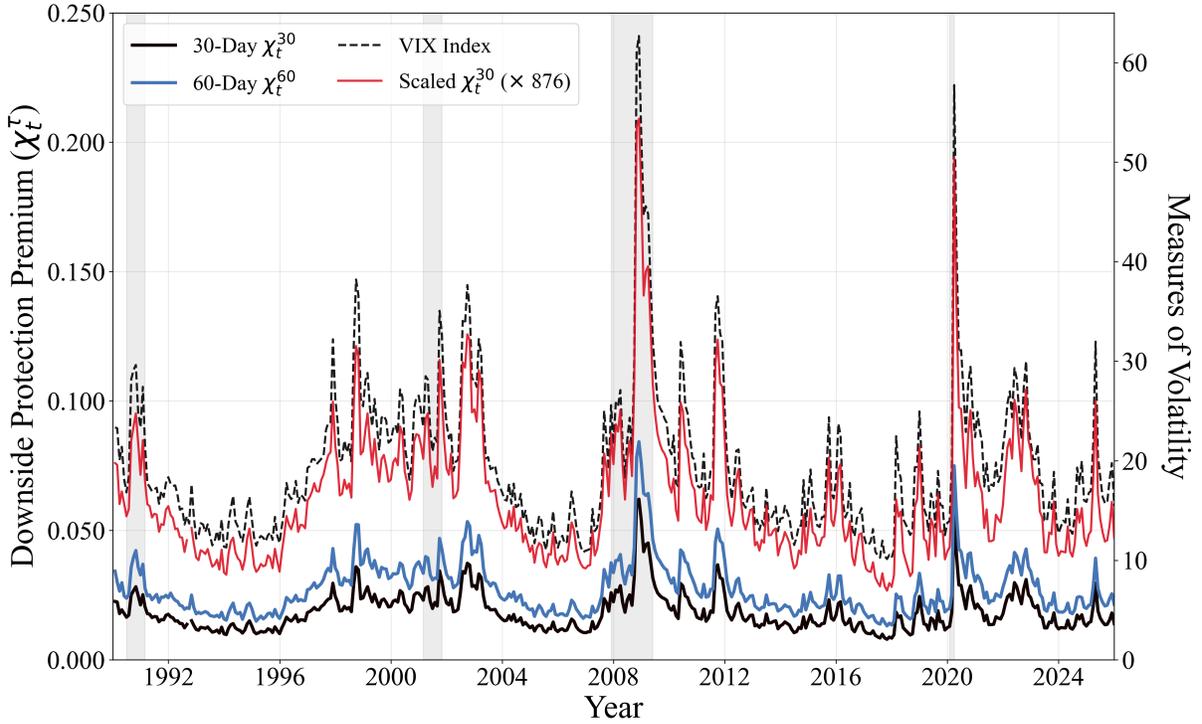


Figure 3: Monthly aggregated  $P_t(\tau)/S_t$  for SPX at 30, 60, and 90-day target expirations (1990–2025). NBER recession periods are shaded in grey.

price of a single forward-ATM put and therefore primarily reflects local implied volatility at the forward strike, whereas the VIX is constructed from a strip of out-of-the-money options and recovers a model-free measure of risk-neutral variance. As a result, any wedge between the rescaled DPP and the VIX reflects departures from the lognormal benchmark assumed in Black–Scholes, including skewness, tail thickness, and jump risk. This wedge is closely related to the cross-strike shape of the implied volatility surface. If the smile were flat, ATM implied volatility would summarize the strip well and the mapping in (3) would be nearly exact. In equity index data, the smile is downward sloping in the put wing, so low-strike downside insurance is priced at higher implied volatilities than the forward-ATM contract. This raises the strip-based variance measure relative to the ATM-based DPP benchmark and implies that the rescaled DPP will generally lie below the VIX in levels. However, if the *shape of the smile is relatively stable over time*, then this wedge will be approximately constant, so the DPP will track the time-series movements in the VIX closely even if it is consistently lower in level.<sup>6</sup>

Figure 3 plots the DPP,  $\chi_t^\tau = P_t^{\text{ATM}}(\tau)/S_t$ , for 30 and 60-day target expirations from 1990.1 through 2025.12, with NBER recession periods shaded in grey (left y-axis). End-of-day prices were utilized to construct daily values for  $\chi_t^\tau$  as discussed in the previous section, which were subsequently aggregated

<sup>6</sup>There are many ways to demonstrate this point. I could plot the volatility smile over time and confirm a consistent downward sloping put wing; plot the time-varying deviation in the VIX and the approximation (3); show correlation in the monthly aggregated time series between the VIX and DPP. I chose the last option as it is most relevant for this paper, but also computed the first two. Results, available upon request, were consistent in that the DPP tracks the VIX nearly identically over time.

to monthly values via arithmetic means. The right  $y$ -axis plots the monthly aggregated scaled volatility estimate of (3), given by  $876 \cdot \chi_t^{30}$  along with the VIX.

Figure 3 confirms that the VIX exceeds the approximation  $876 \cdot \chi_t^{30}$  consistent with the discussion above. What is relevant for this paper is the correlation between the two time series. The 30-day DPP ( $\chi_t^{30}$ ) shows very strong correlation with the VIX across several frequencies. Over the 1990–2025 sample, the correlation is 0.989 at the daily level and increases with temporal aggregation (arithmetic average), reaching 0.991 (weekly), 0.992 (biweekly), and 0.993 when aggregated monthly, the most relevant frequency for macroeconomic analysis. Lead-lag patterns at monthly frequency confirm this pattern. The correlation peaks at lag 0 (0.9931), with nearly symmetric decay on both sides: 0.8461 at lag  $-1$  month and 0.8469 at lag  $+1$  month. This symmetry and rapid drop-off indicate that the two series move essentially in lockstep. Subsample analysis at monthly frequency further confirms the link. The correlation is 0.9932 before 2008 and 0.9935 afterward; it reaches 0.9951 during NBER recession months and remains high at 0.9899 during expansions. Finally, the correlation in monthly changes is also strong: the correlation between  $\Delta \log(\chi_t^{30})$  and  $\Delta \log(\text{VIX}_t)$  is 0.9823, and the correlation between  $\Delta \chi_t^{30}$  and  $\Delta \text{VIX}_t$  is 0.9877. The particularly strong comovement in recessions—when uncertainty shocks are expected to have the largest real effects—suggests that  $\chi_t^{30}$  reliably captures the same uncertainty signal as the VIX across economic regimes.

**3.1 CONSTRUCTING AN UNCERTAINTY SHOCK** In the macro uncertainty literature, the VIX has been used in several closely related ways: as the recursively ordered uncertainty variable in VARs à la Basu and Bundick (2017); as a direct observable proxy for uncertainty in monthly VAR evidence (Leduc and Liu, 2016); at higher frequency to construct uncertainty innovations (Ferrara and Guérin, 2018); and, in some cases, jointly with financial variables such as credit spreads to identify uncertainty shocks under additional restrictions (Lhuissier and Tripier, 2021). These approaches all treat implied stock-market volatility as a benchmark empirical measure of uncertainty. To connect the DPP to the statistical identification used in the uncertainty-shock literature, we follow Bloom (2009) precisely and define uncertainty shocks as large positive deviations of a volatility proxy from its smooth trend. While this section focuses on Bloom’s (2009) original construction, given the correlation established in the previous section, the DPP  $\chi_t^{30}$  could replace the VIX in the studies above with no change in inference.

Specifically, we aggregate daily observations of the VIX and the normalized at-the-money insurance premium  $\chi_t^{30}$  to monthly means, apply the Hodrick–Prescott filter with smoothing parameter  $\lambda = 129,600$ , and treat the resulting cyclical component as the shock series. We then implement a Bloom-style event identification rule by selecting the peak month within each contiguous episode where the cyclical component exceeds a 1.65 standard deviation threshold.

Figure 4 compares the resulting statistical shocks for the VIX (using  $\log(\text{VIX})$ ) and for  $\chi_t^{30}$ . The largest peak months coincide across the two measures, including March 2020 (COVID-19 onset) and November 2008 (Global Financial Crisis peak), with broad alignment in the next tier of episodes (e.g., September 2011, September 2002, September 1998). Differences are within 2% of the arbitrary 1.65 standard deviation threshold; that is, if we extended the threshold to  $1.65 \pm 0.033$ , the identified shocks would line up exactly. Thus, the same reduced-form uncertainty shock episodes that drive real outcomes in the

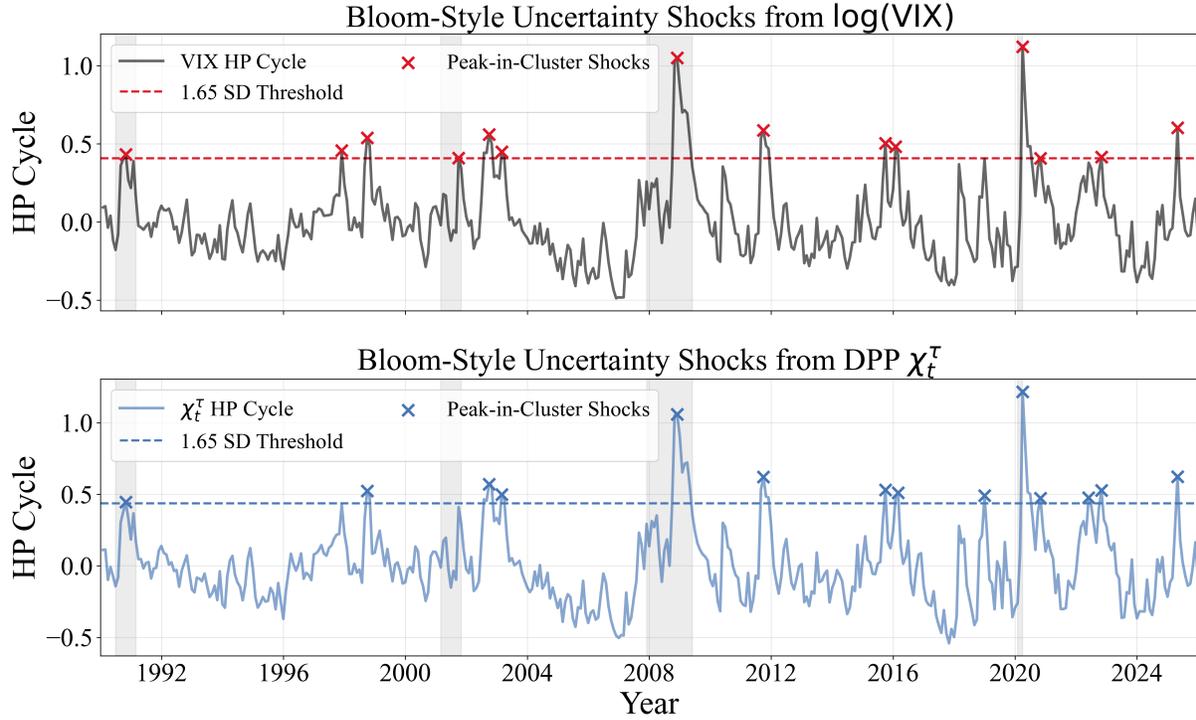


Figure 4: Bloom-style statistical uncertainty shocks constructed from monthly data (1990–2025). Panel A applies the Hodrick–Prescott filter ( $\lambda = 129,600$ ) to monthly  $\log(\text{VIX})$  and plots the cyclical component; Panel B applies the same filter to the monthly DPP  $\chi_t^{\tau}$  aggregated to monthly levels. In each panel, shocks are identified as peak months within contiguous episodes where the HP cycle exceeds a 1.65 standard deviation threshold (dashed line). Differences are within 2% of the arbitrary 1.65 standard deviation threshold. NBER recession periods are shaded in grey.

VIX-based literature also appear as sharp spikes in the Downside Protection Premium.

Taken together, these results establish that the monthly aggregated Downside Insurance Premium is not merely correlated with the VIX, but would produce a nearly identical uncertainty shock series under the same statistical identification procedure. The close correspondence implies that empirical findings obtained using VIX-based shocks in the existing literature would be robust to using the DPP instead.

#### 4 EMPIRICAL DECOMPOSITION

As derived in Section 2,  $\chi_t^{\tau}$  separates into (i) the *risk-neutral probability* of finishing below the forward and (ii) the *expected shortfall severity* conditional on that event,  $\mathbb{Q}(\tilde{R}_t^{\tau} < 1) \cdot \left(1 - \mathbb{E}_t^{\mathbb{Q}}[\tilde{R}_t^{\tau} \mid \tilde{R}_t^{\tau} < 1]\right)$ . Consequently, a rise in  $\chi_t^{\tau}$  can reflect a higher probability of shortfall under  $\mathbb{Q}$ , a deeper left tail (larger conditional shortfalls), or both.

This decomposition makes the DPP more informative than the VIX for macroeconomic analysis because it separates two conceptually distinct sources of downside concern that are conflated when studying the VIX alone. The VIX is a reduced-form object; it summarizes the overall level of risk-neutral vari-

ance, but does not reveal whether a rise reflects a greater likelihood of adverse outcomes, a worsening of their expected severity. By contrast, the DPP admits a transparent economic breakdown into *probability of shortfall* and *amount of shortfall*. For macroeconomists, this distinction is useful because these two margins can map into different underlying mechanisms and different transmission channels as discussed in detail below. In this sense, the DPP retains the empirical success of the VIX in identifying uncertainty shocks, while also allowing one to ask more nuanced questions about transmission of the shock.

In order to operationalize these concepts, we need an estimate of the risk-neutral density,  $\boldsymbol{\pi}^{\mathbb{Q}}$ . Here we follow the entropic (i.e., minimum relative entropy) approach of Buchen and Kelly (1996) and Stutzer (1996) and obtain  $\boldsymbol{\pi}^{\mathbb{Q}}$  as the probability vector closest to  $\boldsymbol{\pi}^{\mathbb{P}}$  in Kullback–Leibler (KLIC) divergence,

$$I(\boldsymbol{\pi}^{\mathbb{Q}}, \boldsymbol{\pi}^{\mathbb{P}}) = \sum_{i=1}^N \pi_i^{\mathbb{Q}} \ln \left( \frac{\pi_i^{\mathbb{Q}}}{\pi_i^{\mathbb{P}}} \right) \quad (4)$$

The guiding economic principle is parsimony—among all state-price distortions that satisfy observed pricing restrictions, we select the one that departs as little as possible from baseline beliefs in an information-theoretic sense. This criterion delivers a unique mapping from beliefs to prices and therefore a unique risk-neutral density. The optimization minimizes (4) subject to the following constraints.

**1. Normalization.**

$$\sum_{i=1}^N \pi_i^{\mathbb{Q}} = 1, \quad \pi_i^{\mathbb{Q}} \geq 0 \quad (5)$$

which imposes standard restrictions on discrete probability spaces.

**2. Martingale Restriction.** A no-arbitrage condition on the returns

$$\sum_{i=1}^N \pi_i^{\mathbb{Q}} R_i = e^{(r_f - d)\tau} \quad (6)$$

The expected risk-neutral return of the asset must equal the discounted return.

**3. Downside Protection Premium ( $\chi_t^\tau$ ).**

$$\sum_{i=1}^N \pi_i^{\mathbb{Q}} \max(e^{(r_f - d)\tau} - R_i, 0) = e^{-r_f \tau} \left( \frac{P_t^{\text{AtM}}(\tau)}{S_t} \right) \equiv e^{-r_f \tau} \chi_t^\tau \quad (7)$$

where  $P_t^{\text{AtM}}(\tau)$  is the time- $t$  price of an AtM-forward put with strike  $K = F_{t,t+\tau} = S_t e^{(r_f - d)\tau}$  and maturity  $\tau$ , and  $\chi_t^\tau$  is this price divided by underlying. This restriction is a moment condition that utilizes the market price of an AtM-forward put to uniquely identify the risk-neutral probability of downside states.

From (4)–(7), form the Lagrangian

$$\mathcal{L} = \sum_{i=1}^N \pi_i^{\mathbb{Q}} \ln \left( \frac{\pi_i^{\mathbb{Q}}}{\pi_i^{\mathbb{P}}} \right) + \lambda_0 \left( 1 - \sum_{i=1}^N \pi_i^{\mathbb{Q}} \right) + \lambda_1 \left( e^{(r_f-d)\tau} - \sum_{i=1}^N \pi_i^{\mathbb{Q}} R_i \right) + \lambda_2 \left( e^{-r_f\tau} \chi_t^\tau - \sum_{i=1}^N \pi_i^{\mathbb{Q}} g(R_i) \right) \quad (8)$$

where  $g(R_i) \equiv \max(e^{(r_f-d)\tau} - R_i, 0)$ . Taking the first-order condition with respect to  $\pi_i^{\mathbb{Q}}$  and solving yields

$$\pi_i^{\mathbb{Q}} = \pi_i^{\mathbb{P}} \exp \left( \lambda_0 - 1 - \lambda_1 R_i - \lambda_2 g(R_i) \right) \quad (9)$$

Imposing normalization determines the constant,<sup>7</sup> so the solution can be written compactly as

$$\pi_i^{\mathbb{Q}} = \frac{\pi_i^{\mathbb{P}} \exp \left( -\lambda_1 R_i - \lambda_2 g(R_i) \right)}{\sum_{j=1}^N \pi_j^{\mathbb{P}} \exp \left( -\lambda_1 R_j - \lambda_2 g(R_j) \right)} \quad (10)$$

The entropic tilting approach has attractive numerical properties. Because the solution takes an exponential form and the objective is strictly convex, the optimization is globally concave and admits a unique solution. In practice, the problem reduces to solving a small system of nonlinear equations for the two multipliers  $\lambda_1$  and  $\lambda_2$  (typically via Newton–Raphson or a quasi-Newton method), which is fast and stable even for moderately large supports ( $N \approx 200$ –500 states). The exponential reweighting naturally avoids negative probabilities and preserves the support of the baseline distribution. These features make daily estimation over long samples computationally efficient and robust to tail behavior in the physical measure.

The Lagrangian described above characterizes the entropic distortion for a static distribution at a generic date  $t$ . The main empirical contribution of the paper, however, comes from the time-series evolution of these structural objects. To recover the history, I solve the entropic program sequentially for each trading day in the sample. At each date  $t$ , the physical distribution  $\boldsymbol{\pi}^{\mathbb{P}}$  over  $\tau$ -horizon outcomes is estimated nonparametrically (discussed below). Given this baseline  $\boldsymbol{\pi}^{\mathbb{P}}$ , the daily constraints—pinned to the observed  $\chi_t^\tau$  and the risk-free rate—identify the closest risk-neutral measure  $\boldsymbol{\pi}^{\mathbb{Q}}$  in KLIC distance. Repeating this procedure daily across the 1990–2025 sample yields a time series that is then aggregated to monthly frequency via arithmetic means.

<sup>7</sup>From (9),

$$\sum_{i=1}^N \pi_i^{\mathbb{Q}} = \exp(\lambda_0 - 1) \sum_{i=1}^N \pi_i^{\mathbb{P}} \exp(-\lambda_1 R_i - \lambda_2 g(R_i))$$

The normalization constraint  $\sum_i \pi_i^{\mathbb{Q}} = 1$  therefore implies

$$\exp(\lambda_0 - 1) = \left[ \sum_{i=1}^N \pi_i^{\mathbb{P}} \exp(-\lambda_1 R_i - \lambda_2 g(R_i)) \right]^{-1}$$

Substituting this expression back into (9) yields

$$\pi_i^{\mathbb{Q}} = \frac{\pi_i^{\mathbb{P}} \exp(-\lambda_1 R_i - \lambda_2 g(R_i))}{\sum_{j=1}^N \pi_j^{\mathbb{P}} \exp(-\lambda_1 R_j - \lambda_2 g(R_j))}$$

**Estimating the Physical Distribution  $\pi^{\mathbb{P}}$**  At each trading date  $t$ , the physical distribution of  $\tau$ -horizon gross returns,  $R_{t,t+\tau} \equiv S_{t+\tau}/S_t$ ,  $x_{t,\tau} \equiv \log R_{t,t+\tau}$ , is estimated nonparametrically using a rolling window of the previous 10-years of historical data. The objective is to construct a conditional baseline distribution of returns that reflects the information set available to traders at date  $t$  and that is defined on a fixed discrete support.

For a target maturity  $\tau$  (e.g.,  $\tau = 30/365$ ), the forward price  $S_{t+\tau}$  is defined at the exact calendar date  $t + \tau$ . Because  $t + \tau$  may not coincide with a trading day, the log price ( $\log S_{t+\tau}$ ) is obtained by linear interpolation between the nearest observed trading dates bracketing  $t + \tau$ . Specifically, if  $t_0 \leq t + \tau \leq t_1$  denote the closest trading dates before and after  $t + \tau$ , then

$$\log S_{t+\tau} = (1 - \omega) \log S_{t_0} + \omega \log S_{t_1}, \quad \omega = \frac{t + \tau - t_0}{t_1 - t_0}$$

This construction ensures that each trading date  $t$  generates a  $\tau$ -calendar-day return, thereby aligning the physical return horizon exactly with the option-based pricing restriction imposed above.

Let  $W$  denote the window length (baseline  $W = 2520$  trading days). At date  $t$ , the sample,  $\{x_{t-j,\tau}\}_{j=1}^W$  is used to estimate the conditional physical distribution. The rolling window allows the distribution to evolve over time and ensures that  $\pi_t^{\mathbb{P}}$  reflects the prevailing volatility and tail environment rather than a long-run unconditional average. This empirical distribution of  $\{x_{t-j,\tau}\}_{j=1}^W$  is partitioned into  $K$  quantile bins (baseline  $K = 300$ ) with bin probabilities proportional to the number of observations falling into each bin. To avoid zero-probability states and ensure numerical stability of the Kullback–Leibler minimization, Dirichlet pseudo-counts are added to each bin. Let  $n_k$  denote the raw bin count. The smoothed probabilities are

$$\pi_{k,t}^{\mathbb{P}} = \frac{n_k + \alpha}{\sum_{j=1}^K (n_j + \alpha)}$$

where  $\alpha$  is a positive constant (baseline  $\alpha = 1$ ). This step guarantees strictly positive support and prevents numerical instability arising from empty bins.

The resulting probability vector  $\pi_t^{\mathbb{P}} = (\pi_{1,t}^{\mathbb{P}}, \dots, \pi_{K,t}^{\mathbb{P}})$  defines the discrete physical measure on the corresponding support  $\{R_{k,t}\}_{k=1}^K$ . The object  $\pi_t^{\mathbb{P}}$  represents a time-varying, nonparametric estimate of the conditional physical return distribution over horizon  $\tau$ .

Results reported below are robust to alternative parameterizations, including shorter and longer rolling windows ( $W \in \{3000, 4000\}$ ), different bin counts ( $K \in \{200, 300, 400\}$ ), and milder or stronger Dirichlet smoothing ( $\alpha \in \{0.5, 1.0\}$ ). However, substantially shorter windows (e.g.,  $W = 252$ ) produce more volatile estimates of the downside-insurance shadow price  $\lambda_{2,t}$ . This increased volatility arises because very short windows allow the left tail of the physical distribution to fluctuate excessively from day to day, requiring large compensating adjustments in the entropic tilt to match the observed DPP. If the physical distribution is overly sensitive to recent outliers, the optimizer must choose extreme values of  $\lambda$  to satisfy the moment condition imposed by the market price of the AtM forward put, leading to economically implausible shifts in risk preferences. Assuming at least five years of knowledge when computing the physical distribution is important for the results of the paper. This assumption is justified because by

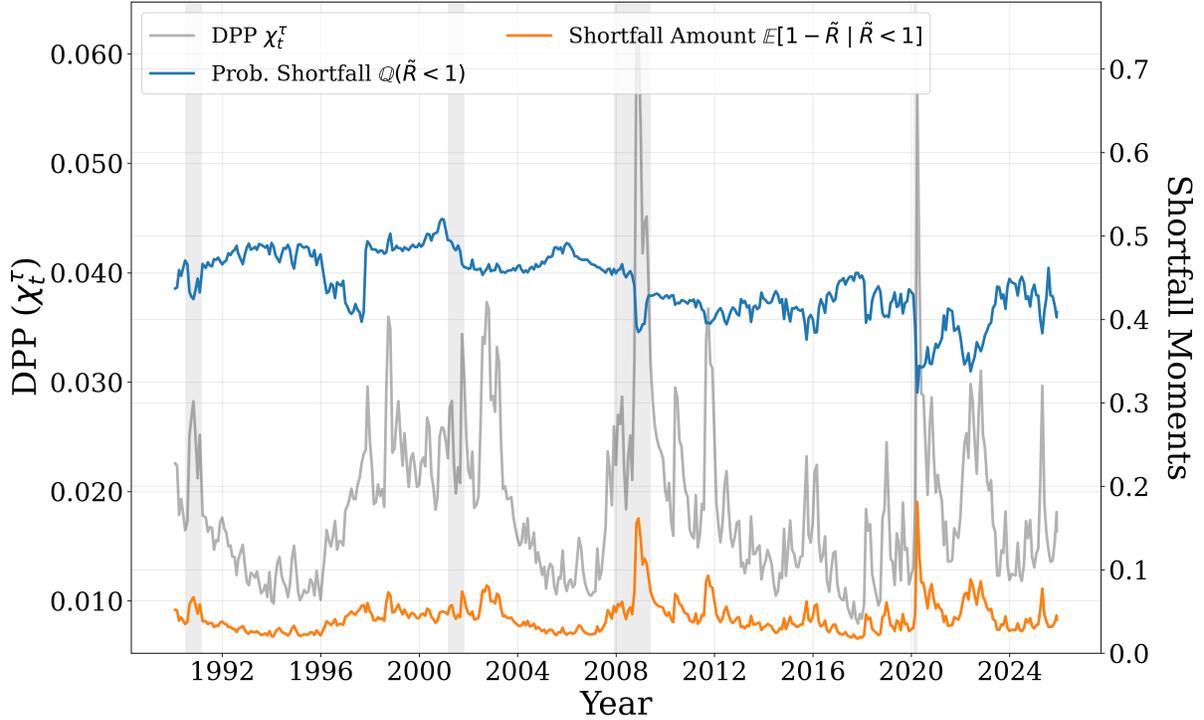


Figure 5: Decomposition of downside insurance:  $\chi_t^r$  (gray, left axis), risk-neutral shortfall probability  $Q_t(\tilde{R}_t^r < 1)$  (blue, right axis), and conditional shortfall severity  $E_t^Q[1 - \tilde{R}_t^r | \tilde{R}_t^r < 1]$  (orange, right axis). NBER recessions are shaded in gray. Monthly averages from daily entropic estimates, 1990–2025.

1990 traders were well aware of the possibility of fat left tails thanks to the 1987 market crash. As shown by Kozlowski et al. (2020), traders’ estimates of tail risk are highly persistent. Market participants do not purge events like the 2008 financial crisis and the 1987 crash from their baseline expectations after a few years.

**4.1 DECOMPOSITION INTO SHORTFALL PROBABILITY AND SEVERITY** As noted above, the DPP  $\chi_t^r$  can be decomposed into two economically distinct components: the probability of finishing below the forward benchmark  $\text{Prob}(\text{Shortfall})$ ,  $Q_t(\tilde{R}_t^r < 1)$ , and the expected proportional shortfall conditional on that event  $\text{Amt}(\text{Shortfall})$ ,  $E_t^Q[1 - \tilde{R}_t^r | \tilde{R}_t^r < 1]$ .

Table 2: Monthly Decomposition of Downside Protection Premium

Variable	Mean	Std. Dev.	Median	Skewness	AR(1)
DPP, $\chi_t^r$	0.0185	0.0076	0.0166	1.9980	0.8553
Shortfall probability, $Q_t(\tilde{R}_t^r < 1)$	0.4412	0.0377	0.4443	-0.4872	0.9527
Conditional severity, $E_t^Q[1 - \tilde{R}_t^r   \tilde{R}_t^r < 1]$	0.0429	0.0201	0.0385	2.5521	0.8284

Note: Monthly averages of daily entropic estimates, 1990–2025.

Table 2 reports summary statistics for the monthly decomposition over 1990–2025. The central result is that *most time-series variation in the DPP reflects movements in conditional shortfall severity rather than in shortfall probability*. The risk-neutral probability of shortfall is high on average,  $\mathbb{Q}_t(\tilde{R}_t^T < 1) = 0.441$ , and persistent, with a monthly AR(1) coefficient of 0.953. This object varies within a relatively narrow range, from 0.312 to 0.520, and exhibits only mild negative skewness. This suggests that option prices embed a slowly moving baseline assessment of downside exposure that evolves only gradually over time. In contrast, conditional shortfall severity is far more volatile and crisis-sensitive. With a mean of 4.29%, its standard deviation is nearly half as large as its level (2.01 percentage points), its skewness is 2.55, with a maximum of 18.13%. Thus, unlike the shortfall probability, conditional severity displays pronounced right-tail behavior and large episodic spikes. Finally, DPP  $\chi_t^T$  comoves much more closely with conditional severity than with shortfall probability. The monthly correlation between  $\chi_t^T$  and conditional severity is 0.968, whereas the correlation between  $\chi_t^T$  and shortfall probability is weak and slightly negative (-0.264). Given the multiplicative decomposition of the DPP, these moments imply that crisis spikes in downside insurance prices are driven primarily by changes in the expected magnitude of downside outcomes rather than by changes in the probability of merely finishing below the forward benchmark.

Figure 5 confirms these findings. The shortfall probability remains persistently elevated throughout the sample and moves relatively smoothly across business-cycle phases. By contrast, conditional severity exhibits sharp spikes concentrated in periods of acute financial stress, especially during the global financial crisis and the COVID-19 episode.

The mean shortfall probability is less sensitive to business-cycle states, averaging 44.15% in expansions and 43.91% in recessions. By contrast, conditional severity rises from 4.00% in expansions to 7.15% in recessions, an increase of roughly 79%. The insurance premium itself rises from 1.74% to 3.02%, or about 74%. Thus, nearly the entire cyclical increase in downside insurance prices reflects a recession-induced increase in expected shortfall severity rather than an increase in the probability of shortfall.

Taken together, these patterns reveal a sharp separation in risk-neutral downside beliefs. The market’s assessment of the likelihood of finishing below the forward benchmark is smooth and highly persistent, whereas its assessment of the magnitude of losses in those downside states is volatile and crisis-dominant. Time variation relevant for macro-finance applications comes primarily from changes in expected downside severity. More broadly, these results suggest that what macroeconomists often label an “uncertainty shock” may, in practice, be less about a higher probability of mild bad outcomes and more about a repricing of the expected severity of tail events.

## 5 UNCERTAINTY MEASURES IN VARs AND STRUCTURAL MODELING

I now compare the DPP, probability and severity of shortfall to five benchmark measures that have been widely used directly as uncertainty proxies in empirical macro-finance: macro uncertainty, real uncertainty, financial uncertainty, economic policy uncertainty, and the VIX. The first three measures are the updated *macro*, *real*, and *financial* uncertainty indexes. These series implement the methodology of Jurado et al. (2015), who define uncertainty as the conditional volatility of the unforecastable component of a large panel of variables, and the later update framework in Ludvigson et al. (2021). The repos-

itory <https://www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes> provides monthly indexes at forecast horizons  $h = 1, 3, 12$ . The macro index is constructed from a broad macroeconomic panel, the real index isolates uncertainty in real activity measures, and the financial index summarizes uncertainty in a large panel of financial variables including returns, yields, spreads, valuation ratios, and other asset-market indicators. In contrast to VIX-style option-implied measures, these indexes are intended to capture latent forecast-error uncertainty. The fourth benchmark is the U.S. *Economic Policy Uncertainty* index, downloaded from FRED under the series code USEPUINXD and aggregated to monthly values. This measure is based on the newspaper-count methodology of Baker et al. (2016). It is designed to proxy uncertainty related specifically to economic policy rather than general macroeconomic volatility. In the Baker–Bloom–Davis construction, the index rises when newspapers more frequently discuss the economy, policy, and uncertainty jointly. The fifth benchmark is the *VIX*, following Bloom (2009) and the subsequent literature, as described above.

Table 3: Correlations of Shortfall Components with Benchmark Uncertainty Measures

Benchmark measure	$\chi_t^r$	$\mathbb{Q}_t(\tilde{R}_t^r < 1)$	$\mathbb{E}_t^{\mathbb{Q}}[1 - \tilde{R}_t^r   \tilde{R}_t^r < 1]$
<i>Full sample</i>			
VIX	0.993	-0.281	0.966
Macro uncertainty ( $h = 1$ )	0.581	-0.481	0.666
Real uncertainty ( $h = 1$ )	0.374	-0.594	0.508
Financial uncertainty ( $h = 1$ )	0.780	-0.300	0.762
EPU	0.473	-0.474	0.565
<i>Recessions</i>			
VIX	0.995	-0.736	0.980
Macro uncertainty ( $h = 1$ )	0.775	-0.818	0.818
Real uncertainty ( $h = 1$ )	0.605	-0.873	0.714
Financial uncertainty ( $h = 1$ )	0.805	-0.692	0.802
EPU	0.639	-0.735	0.706
<i>Expansions</i>			
VIX	0.989	-0.226	0.949
Macro uncertainty ( $h = 1$ )	0.341	-0.489	0.478
Real uncertainty ( $h = 1$ )	0.213	-0.569	0.382
Financial uncertainty ( $h = 1$ )	0.718	-0.279	0.712
EPU	0.333	-0.446	0.451

Table 3 reports correlations. Three implications stand out. First, the DPP is strongly positively correlated with all benchmark measures, with the tightest relationship unsurprisingly appearing for the VIX. Other than the VIX, the DPP is most strongly associated with financial uncertainty (0.780), followed by macro uncertainty (0.581), EPU (0.473), and real uncertainty (0.374). This ranking is intuitive as the DPP is extracted from equity option prices, it should be most closely connected to financial-market based uncertainty measures, somewhat less tightly connected to broad macro uncertainty, and more weakly connected to uncertainty measures tied specifically to real activity or policy news. Second, across all five

benchmarks, conditional shortfall severity is positively correlated with the uncertainty proxies, while shortfall probability is negatively correlated. The pattern holds for macro uncertainty (0.666 versus  $-0.481$ ), real uncertainty (0.508 versus  $-0.594$ ), financial uncertainty (0.762 versus  $-0.300$ ), and EPU (0.565 versus  $-0.474$ ). Third, these relationships become even sharper during recessions. In recession months, the correlation between severity and the benchmark measures rises substantially to 0.980 for the VIX, 0.818 for macro uncertainty, 0.714 for real uncertainty, 0.802 for financial uncertainty, and 0.706 for EPU. At the same time, the shortfall-probability component becomes even more negatively correlated with the benchmark measures, with values ranging from  $-0.692$  to  $-0.873$ . This recession-expansion contrast further confirms that when uncertainty matters most for macroeconomic outcomes, the relevant information is not a large change in the likelihood of crossing the forward benchmark, but rather a sharp repricing of how costly downside states are expected to be. In expansions, the same qualitative ranking remains, but the correlations are uniformly weaker, especially for macro uncertainty, real uncertainty, and EPU. Financial uncertainty remains the closest non-VIX benchmark in both states, consistent with the market-based nature of the DPP.

Table 3 shows that the DPP is not only closely related to standard benchmark uncertainty measures, but also provides a more refined interpretation of what these measures are capturing. In particular, the common variation between widely used uncertainty proxies and option-implied downside insurance appears to reflect fluctuations in expected downside *severity*, not in the baseline probability of a shortfall. This distinction is invisible in reduced-form aggregate measures such as the VIX, EPU, or forecast-error uncertainty indexes, but it is central for understanding the nature of the shocks that these variables proxy for in macro-finance applications.

**5.1 INTERPRETING UNCERTAINTY SHOCKS IN VARs** A central use of uncertainty measures in macroeconomics is as an observed proxy in reduced-form VARs and local projections. In that literature, a variable such as the VIX,  $\log$  VIX, or a smoothed cyclical component of implied volatility is inserted directly into the system, and an innovation to that series is interpreted as an “uncertainty shock.” This approach has proved empirically useful, but the underlying economic content of the identified shock is often left ambiguous. In particular, a one-dimensional implied-volatility proxy does not reveal whether market participants are revising the likelihood of adverse states or the severity of losses conditional on those states. The results above clearly indicate that the macro shock identified by the VIX in VARs is *not a diffuse rise in near-threshold uncertainty but a repricing of downside tail severity*. How does this materialize in a VAR?

I estimate a monthly four-variable vector autoregression (VAR) following the identification strategy of Leduc and Liu (2016). The model is designed to study the dynamic effects of uncertainty shocks on macroeconomic variables. The endogenous variables are ordered as follows: natural logarithm of the monthly average of the CBOE Volatility Index (VIX); Unemployment rate (UNRATE); Year-over-year CPI inflation:  $100 \times [\log(\text{CPI}_t) - \log(\text{CPI}_{t-12})]$ ; and 3-month Treasury bill rate (TB3MS). The recursive (Cholesky) identification assumes that uncertainty ( $\log$  VIX) is contemporaneously exogenous to the macroeconomic variables, while the macro variables may respond within the month to uncertainty innovations. This ordering is standard in the uncertainty literature. I estimate the VAR using frequentist

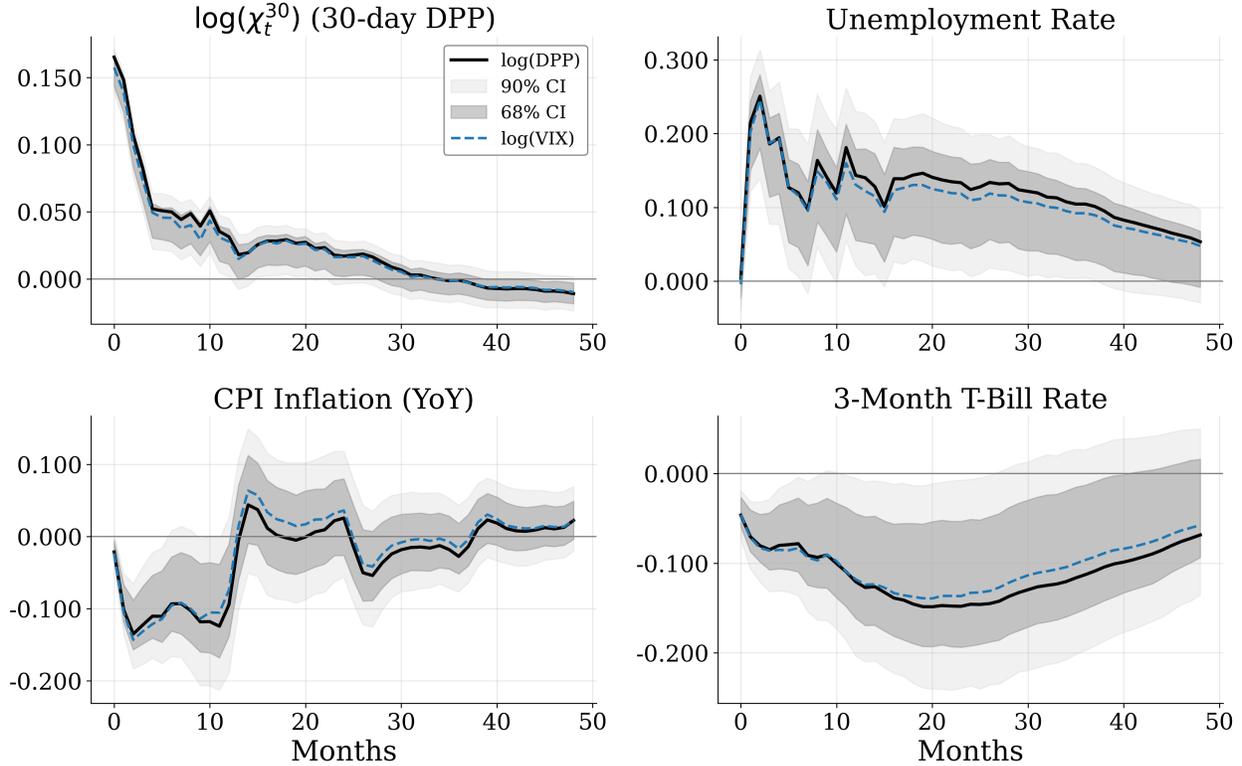


Figure 6: Impulse responses to a one-standard-deviation shock in  $\log(\text{DPP})$  (solid black) versus  $\log(\text{VIX})$  (dashed blue). Dark gray shading: 68% bootstrap confidence band; light gray shading: 90% bootstrap confidence band. Monthly four-variable VAR, 1990–2025, recursive identification.

methods. The VAR is estimated with a constant term and lag length is selected automatically using the Akaike Information Criterion (AIC) with a maximum of 24 lags. Point estimates of impulse response functions (IRFs) are obtained from the orthogonalized residuals under the recursive identification scheme. Confidence intervals are constructed via residual-based bootstrap with 1,000 replications.

Figure 6 replaces the  $\log(\text{VIX})$  metric with the natural log of the Downside Protection Premium,  $\log(\chi_t^{30})$  and plots the impulse response with 68% / 90% confidence intervals to a one standard deviation shock in  $\log(\text{DPP})$ . The blue dashed line plots the impulse response estimated with  $\log(\text{VIX})$  used as the shock. The responses line up nearly exactly across the two specifications. A shock that raises  $\log(\text{DPP})$  produces the same qualitative and quantitative pattern as a shock that raises  $\log(\text{VIX})$  and is very similar to Leduc and Liu (2016). A persistent increase in unemployment, a hump-shaped rise then fall in inflation, and an initial decline followed by a gradual rise in short-term rates.

The close alignment is not surprising given the high correlation between  $\log(\text{DPP})$  and  $\log(\text{VIX})$  over the sample, but it carries an important interpretive implication. Because  $\chi_t^{30}$  isolates the *price of insurance against downside moves below the forward level* (i.e., the risk-neutral expected loss conditional on falling below the forward), a shock to  $\log(\chi_t^{30})$  has a cleaner economic interpretation (in my opinion) as it primarily reflects a change in the market's assessment of *tail severity* rather than a broad rescaling of near-the-money volatility. A positive innovation to  $\log(\text{DPP})$  is an increase in the proportional cost of

insuring against downside equity risk. Because the VAR is estimated in logs, an impulse of roughly 0.16 plotted in Figure 6 corresponds to an exact proportional increase of  $e^{0.16} - 1 \approx 17.35\%$ . For the DPP, this means that if the baseline monthly insurance cost is 1.85% of spot, the shock raises it to about 2.17% of spot, or from \$1.85 to \$2.17 per \$100 of insured equity exposure. The fact that the macro dynamics are virtually indistinguishable suggests that the VIX shock commonly identified in the literature is, to a first approximation, *also* a tail-severity shock — even though the VIX itself is a symmetric volatility measure.

Figure 7 reports the responses obtained when the uncertainty variable is replaced by the natural logarithm of the risk-neutral shortfall probability.<sup>8</sup> In response to an increase in the probability of shortfall, the unemployment rate falls, inflation increases and interest rates mildly rise. Why? The shortfall-probability component is a threshold object that is continuously updated. It measures the risk-neutral probability of finishing below the forward benchmark at date  $t$ , not the depth of losses conditional on entering the downside tail. Episodes of elevated downside insurance are driven primarily by increases in conditional shortfall severity. As the left tail deepens, probability mass shifts farther into the tail. For this reason, innovations to shortfall probability do not recover the benchmark contractionary uncertainty shock. Rather, they capture local redistributions of risk around the forward threshold, which can be associated with stronger inflation and labor-market outcomes even as downside-tail severity moves in the opposite direction.<sup>9</sup>

Figure 8 makes the decomposition particularly transparent. The vertical line at  $\tilde{R}_t^T = 1$  marks the forward benchmark, so the height of each CDF at that threshold is exactly the risk-neutral shortfall probability,  $Q_t(\tilde{R}_t^T < 1)$ . Moving from a typical date, 2006-06-30, to a crisis date, 2008-10-31, the value of the CDF at the threshold *falls*, from roughly 0.481 to 0.379. Thus, the crisis episode is not characterized by a higher probability of merely finishing below the forward benchmark. At the same time, the crisis distribution exhibits a significantly heavier downside tail below the threshold. This is reflected in the large increase in conditional shortfall severity, from  $s = 0.027$  on 2006-06-30 to  $s = 0.161$  on 2008-10-31. Downside states become much worse conditional on occurring, even though they are not more likely at the forward threshold itself. The figure therefore illustrates the central empirical message of the decomposition: periods that look like “uncertainty shocks” in aggregate measures such as the VIX are better understood as episodes in which the left tail deepens sharply, rather than episodes in which the probability of a modest shortfall simply rises.

**5.2 INTERPRETING UNCERTAINTY SHOCKS IN STRUCTURAL MODELS** Standard DSGE treatments of uncertainty typically model “uncertainty shocks” as innovations to the conditional variance of an otherwise

<sup>8</sup>Because the DPP is multiplicative in probability and severity, shocks to these components are not orthogonal structural primitives of the benchmark model. Instead, they isolate individual margins of downside risk and therefore serve as a diagnostic comparison rather than a formal structural identification. Even with that caveat, the figure is revealing as the response to a shortfall-probability shock differs sharply from the benchmark log(VIX) response.

<sup>9</sup>It is interesting to note that the response observed in Figure 7 is sensitive but robust to excluding recession dates from the sample. Responses are qualitatively the same but no longer statistically significant.

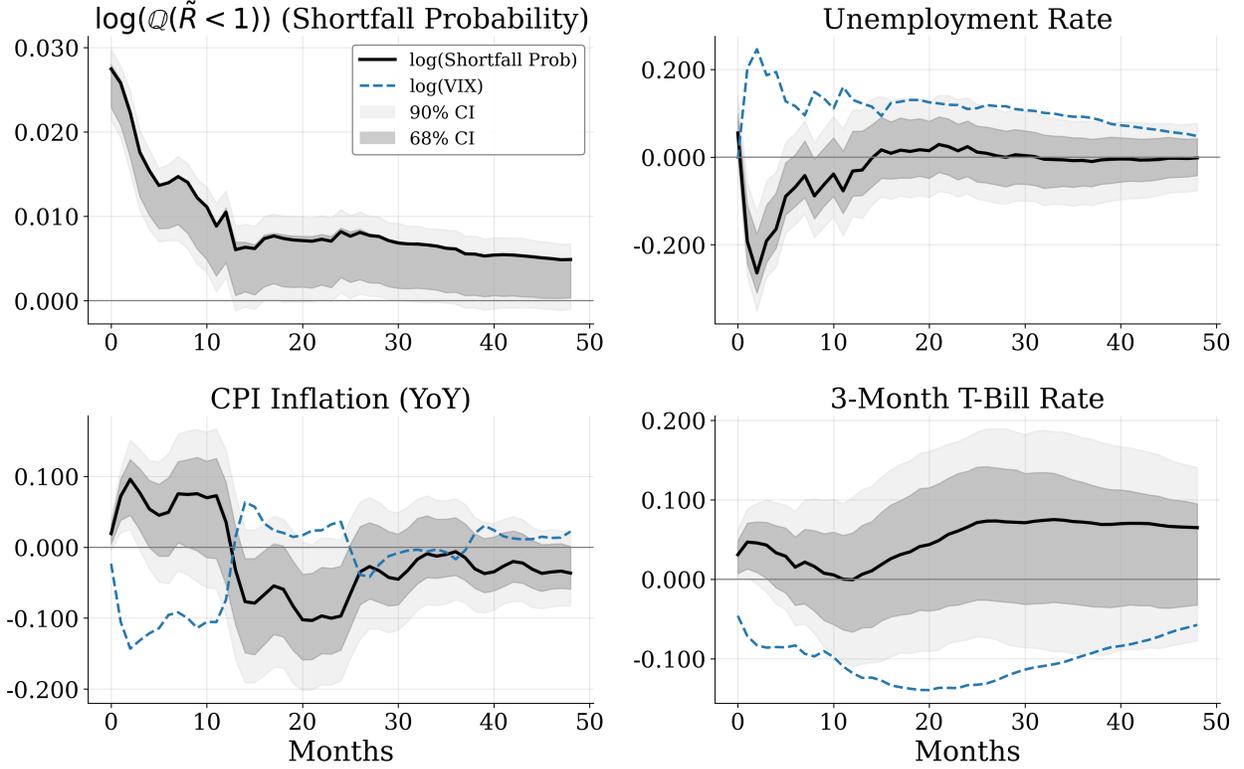


Figure 7: Impulse responses to a one-standard-deviation shock in the log risk-neutral shortfall probability,  $\log(Q_t(\tilde{R}_t^t < 1))$  (solid black), versus the benchmark  $\log(\text{VIX})$  shock (dashed blue). Dark gray shading: 68% bootstrap confidence band; light gray shading: 90% bootstrap confidence band. Monthly four-variable VAR, 1990–2025, recursive identification.

symmetric process

$$\ln Z_t = (1 - \rho_z) \ln \bar{Z} + \rho_z \ln Z_{t-1} + \sigma_{z,t} \varepsilon_{z,t}, \quad (11)$$

$$\ln \sigma_{z,t} = (1 - \rho_{\sigma_z}) \ln \bar{\sigma}_z + \rho_{\sigma_z} \ln \sigma_{z,t-1} + \sigma_{\sigma_z} \varepsilon_{\sigma_z,t}, \quad (12)$$

where  $\varepsilon_{z,t}$  and  $\varepsilon_{\sigma_z,t}$  are i.i.d. innovations and uncertainty is identified exclusively with the second-moment parameter  $\sigma_{z,t}$ . In this framework, a rise in uncertainty means that large positive and large negative innovations become symmetrically more likely. This representation has proved tractable and influential, and it underlies much of the modern DSGE literature on uncertainty shocks.

The evidence in this paper suggests that this representation is too coarse relative to the option-implied object that actually tracks the VIX and reproduces standard uncertainty-shock episodes in macro VARs. The empirical variation in the Downside Protection Premium is not driven primarily by a diffuse broadening of the risk-neutral distribution, nor by a materially higher probability of a modest shortfall below the forward benchmark. Instead, it is driven overwhelmingly by changes in the conditional severity of downside outcomes. In the data, what rises in periods identified as “uncertainty shocks” is not the probability of a mildly bad state, but the expected depth of losses once a downside state occurs. This

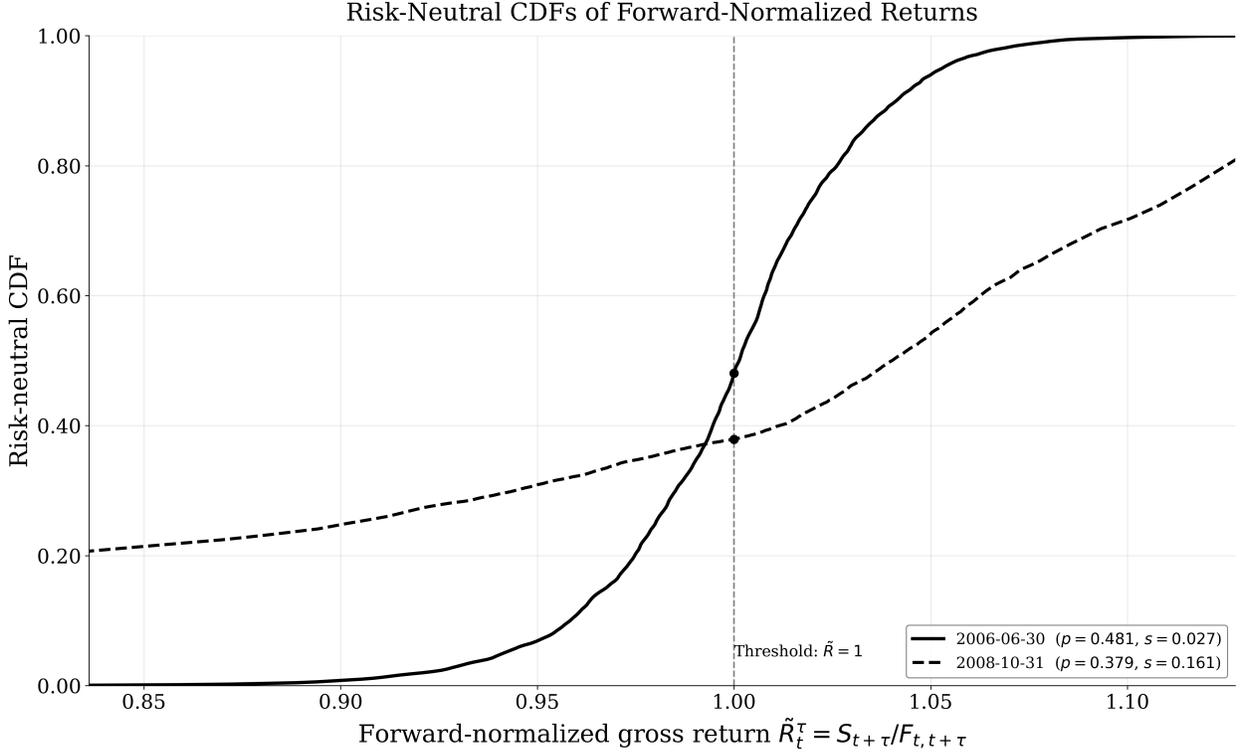


Figure 8: Risk-neutral cumulative distribution functions of forward-normalized returns,  $\tilde{R}_t^\tau = S_{t+\tau}/F_{t,t+\tau}$ , for a calm date (solid black) and a crisis date (dashed black). The vertical dashed line marks the shortfall threshold,  $\tilde{R}_t^\tau = 1$ . The point where each CDF intersects the threshold equals the risk-neutral shortfall probability,  $Q_t(\tilde{R}_t^\tau < 1)$ .

implies that the structural object underlying VIX-identified uncertainty shocks is better understood as a downside-tail shock than as a symmetric stochastic-volatility shock.

This finding places the macro uncertainty literature closer to the time-varying disaster-risk literature of Gourio (2012) and Wachter (2013) than is usually recognized. In Gourio (2012), macroeconomic dynamics are shaped by the possibility of a rare but severe contractionary event; in Wachter (2013), time-varying disaster risk is central for asset pricing. The decomposition in this paper points in the same direction, but with an important refinement. Relative to standard disaster-risk formulations that emphasize time variation in the *probability* of a disaster, the option-implied evidence here points more strongly toward time variation in the *severity* of bad states once they occur.

To preserve tractability while incorporating this evidence, consider a discrete Markov downside-risk structure. Let  $M_t \in \{1, \dots, N\}$  denote a finite-state aggregate risk regime, and let  $s_t \in \{0, 1\}$  indicate the arrival of a downside event. Define the shock process as

$$\ln Z_t = (1 - \rho_z) \ln \bar{Z} + \rho_z \ln Z_{t-1} + \sigma_z \varepsilon_{z,t} - \mathbb{1}(s_t = 1) \kappa(M_t), \quad (13)$$

where  $\varepsilon_{z,t} \sim N(0, 1)$ ,  $\sigma_z > 0$  is a baseline Gaussian volatility, and  $\kappa(M_t) > 0$  is the regime-dependent severity of the downside event. This nests the basic logic of Gourio (2012) as the economy occasionally

faces a discrete contractionary event. The key difference is that the regime  $M_t$  governs the severity of that event, thereby matching the option-implied evidence presented here.

The probability side can be modeled separately. Let

$$p_t \equiv \Pr(s_t = 1 \mid \mathcal{F}_{t-1}) \quad (14)$$

denote the conditional probability of entering the downside state. In the simplest version,  $p_t = \bar{p}$  is fixed and only  $\kappa(M_t)$  varies across regimes. This parsimonious specification is consistent with our empirical results with a shortfall-probability margin that is smooth and highly persistent, and a shortfall-severity margin that is volatile and crisis-dominant. A more flexible specification allows both margins to vary, for example through a Markov transition matrix for  $M_t$  together with state-dependent disaster severity. In either case, the essential point is that the model separates the *arrival* of downside states from the *severity* of downside states.

One final note, this distinction has immediate implications for solution methods. If uncertainty is modeled as a symmetric  $\sigma_{z,t}$  shock, a second-order perturbation is often sufficient, since the relevant object is the conditional variance of an otherwise symmetric innovation. That is why second-order methods have become standard in the DSGE uncertainty literature. The present evidence, however, points to a different structural object: the empirically relevant shock operates through downside asymmetry. In a perturbation framework, first-order approximations impose certainty equivalence and therefore eliminate uncertainty effects altogether. Second-order approximations recover variance effects and precautionary terms, but they remain fundamentally organized around symmetric second moments. Genuine downside-tail asymmetry—including skewness, lower-partial-moment risk, and the differential macroeconomic consequences of bad versus good states—enters naturally only at *third order*. For that reason, if one wishes to preserve a perturbation-based solution method while modeling uncertainty in a way consistent with the option-implied evidence, a third-order approximation is the natural minimum.

The empirical results in this paper therefore suggest a reinterpretation of structural uncertainty shocks. Rather than modeling uncertainty as a shock to the conditional variance of a symmetric innovation, DSGE models should allow for separate probability and severity margins in the downside tail. A Markov disaster-severity structure provides one tractable way to do so. Relative to the standard stochastic-volatility approach, it offers a closer structural counterpart to the option-implied evidence: what macro VARs identify as an uncertainty shock is better interpreted as a repricing of downside-tail severity, in the spirit of Gourio (2012), than as a broad increase in symmetric volatility.

## REFERENCES

- Adrian, Tobias and Joshua Rosenberg**, “Stock returns and volatility: Pricing the short-run and long-run components of market risk,” *Journal of Finance*, 2008, 63 (6), 2997–3030.
- Alfaro, Iván, Nicholas Bloom, and Xiaoji Lin**, “The finance uncertainty multiplier,” *Journal of Political Economy*, 2024, 132 (2), 577–615.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis**, “Measuring economic policy uncertainty,” *Quarterly Journal of Economics*, 2016, 131 (4), 1593–1636.
- Basu, Susanto and Brent Bundick**, “Uncertainty shocks in a model of effective demand,” *Econometrica*, 2017, 85 (3), 937–958.
- Bekaert, Geert and Marie Hoerova**, “The VIX, the variance premium and stock market volatility,” *Journal of Econometrics*, 2014, 183 (2), 181–192.
- , —, and **Marco Lo Duca**, “Risk, uncertainty and monetary policy,” *Journal of Monetary Economics*, 2013, 60 (7), 771–788.
- Berger, David, Ian Dew-Becker, and Stefano Giglio**, “Uncertainty shocks as second-moment news shocks,” *Review of Economic Studies*, 2020, 87 (1), 40–76.
- Bhattarai, Saroj, Arpita Chatterjee, and Woong Yong Park**, “Global spillover effects of US uncertainty,” *Journal of Monetary Economics*, 2020, 114, 71–89.
- Bianchi, Francesco, Howard Kung, and Mikhail Tirsikh**, “The origins and effects of macroeconomic uncertainty,” *Quantitative Economics*, 2023, 14 (3), 855–896.
- Bloom, Nicholas**, “The impact of uncertainty shocks,” *Econometrica*, 2009, 77 (3), 623–685.
- Bok, Brandyn, Thomas M. Mertens, and John C. Williams**, “Macroeconomic drivers and the pricing of uncertainty, inflation, and bonds,” *Journal of Financial Economics*, 2025, 172, 104130.
- Bollerslev, Tim and Viktor Todorov**, “Tails, fears, and risk premia,” *Journal of Finance*, 2011, 66 (6), 2165–2211.
- , **George Tauchen, and Hao Zhou**, “Expected stock returns and variance risk premia,” *Review of Financial Studies*, 2009, 22 (11), 4463–4492.
- Born, Benjamin and Johannes Pfeifer**, “Policy risk and the business cycle,” *Journal of Monetary Economics*, 2014, 68, 68–85.
- Buchen, Peter W. and Michael Kelly**, “The maximum entropy distribution of an asset inferred from option prices,” *Journal of Financial and Quantitative Analysis*, 1996, 31 (1), 143–159.

- Caggiano, Giovanni, Efrem Castelnuovo, and Giovanni Pellegrino**, “Estimating the real effects of uncertainty shocks at the Zero Lower Bound,” *European Economic Review*, 2017, 100, 257–272.
- Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek**, “The macroeconomic impact of financial and uncertainty shocks,” *European Economic Review*, 2016, 88, 185–207.
- Carr, Peter and Liuren Wu**, “Variance risk premiums,” *Review of Financial Studies*, 2009, 22 (3), 1311–1341.
- Cesa-Bianchi, Ambrogio and Emilio Fernández-Corugedo**, “Uncertainty, financial frictions, and nominal rigidities: A quantitative investigation,” *Journal of Money, Credit and Banking*, 2018, 50 (4), 603–636.
- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno**, “Risk shocks,” *American Economic Review*, 2014, 104 (1), 27–65.
- Diercks, Anthony M., Alex Hsu, and Andrea Tamoni**, “When it rains it pours: Cascading uncertainty shocks,” *Journal of Political Economy*, 2024, 132 (2), 694–720.
- Drechsler, Itamar and Amir Yaron**, “What’s vol got to do with it?,” *Review of Financial Studies*, 2011, 24 (1), 1–45.
- Fernández-Villaverde, Jesús, Pablo A. Guerrón-Quintana, Keith Kuester, and Juan F. Rubio-Ramírez**, “Fiscal volatility shocks and economic activity,” *American Economic Review*, 2015, 105 (11), 3352–3384.
- , **Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe**, “Risk matters: The real effects of volatility shocks,” *American Economic Review*, 2011, 101 (6), 2530–2561.
- Ferrara, Laurent and Pierre Guérin**, “What are the macroeconomic effects of high-frequency uncertainty shocks?,” *Journal of Applied Econometrics*, 2018, 33 (5), 662–679.
- Gourio, François**, “Disaster risk and business cycles,” *American Economic Review*, 2012, 102 (6), 2734–2766.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng**, “Measuring uncertainty,” *American Economic Review*, 2015, 105 (3), 1177–1216.
- Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran**, “The tail that wags the economy: Beliefs and persistent stagnation,” *Journal of Political Economy*, 2020, 128 (8), 2839–2879.
- Leduc, Sylvain and Zheng Liu**, “Uncertainty shocks are aggregate demand shocks,” *Journal of Monetary Economics*, 2016, 82, 20–35.
- Lhuissier, Stéphane and Fabien Tripier**, “Regime-dependent effects of uncertainty shocks: A structural interpretation,” *Quantitative Economics*, 2021, 12 (4), 1139–1170.

**Ludvigson, Sydney C., Sai Ma, and Serena Ng**, “Uncertainty and business cycles: Exogenous impulse or endogenous response?,” *American Economic Journal: Macroeconomics*, 2021, 13 (4), 369–410.

**Oh, Joonseok**, “The propagation of uncertainty shocks: Rotemberg versus Calvo,” *International Economic Review*, 2020, 61 (3), 1097–1113.

**Pellegrino, Giovanni, Efram Castelnuovo, and Giovanni Caggiano**, “Uncertainty and monetary policy during the Great Recession,” *International Economic Review*, 2023, 64 (2), 577–606.

**Pflueger, Carolin, Emil Siriwardane, and Adi Sunderam**, “Financial market risk perceptions and the macroeconomy,” *Quarterly Journal of Economics*, 2020, 135 (3), 1443–1491.

**Segal, Gill, Ivan Shaliastovich, and Amir Yaron**, “Good and bad uncertainty: Macroeconomic and financial market implications,” *Journal of Financial Economics*, 2015, 117 (2), 369–397.

**Stutzer, Michael**, “A simple nonparametric approach to derivative security valuation,” *Journal of Finance*, 1996, 51 (5), 1633–1652.

**Todorov, Karamfil and Grigory Vilkov**, “Anatomy of the VIX spike in August 2024,” BIS Bulletin 95, Bank for International Settlements October 2024.

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