Asset Pricing with the Awareness of New Risks

Christian Heyerdahl-Larsen* Philipp Illeditsch† Petra Sinagl‡

June 11, 2026

ABSTRACT

Recessions cause delayed declines in output, followed by recoveries with abnormally high growth. We develop an equilibrium model in which investors become aware of a new aggregate risk but must learn its severity from realized macroeconomic outcomes. The crisis-shape process generates a gradual deterioration and recovery in output, while belief revisions about crisis severity raise risk aversion when outcomes are worse than expected. The model predicts that, after recession onset, risk premia and return volatility do not rise immediately but instead increase gradually before peaking, generating hump-shaped dynamics unlike those in many non-expected-utility models (Ai and Bansal, 2018). When investors learn that the crisis is severe, this generates declining valuation ratios. These dynamics match empirical patterns in output, risk compensation, volatility, and valuations around U.S. recessions.

JEL Classification: G12, E32, E44, G01.

Keywords: hump-shaped risk premia, hump-shaped volatilities, U-shaped output, new priced risk during crises, crisis heterogeneity

*Heyerdahl-Larsen is with the BI Norwegian Business School. *Email*: christian.heyerdahl-larsen@bi.no.

†Illeditsch is with the Mays Business School, Texas A&M University, College Station, TX 77843. *Email*: pilleditsch@mays.tamu.edu

‡Sinagl is with the Tippie College of Business, University of Iowa, Iowa City, IA 52242. *Email*: petra-sinagl@uiowa.edu.

We thank Daniel Andrei, Adlai Fisher, Lorenzo Garlappi, Leyla Jianyu Han (discussant), Burton Hollifield, Chen Li (discussant), Stavros Panageas (discussant), Tim Kroencke (discussant), Julien Pénasse, Dongho Song, Colin Ward (discussant), Jing Xu (discussant), Ram Yamarthy (discussant), Yucheng Yang (discussant), Chao Ying (discussant), Tom Oskar Karl Zeissler (discussant) for helpful comments and suggestions. This paper has benefited from presentations at the University of British Columbia, University of Iowa, Texas A&M University, the 2025 SGF Conference, the 2024 CUHK-RAPS-RCFS Conference, the 2024 Northern Finance Association Conference, the 2024 European Finance Association Conference, the 8th Young Scholars Finance Consortium, the 2024 China International Conference in Finance, the 2024 Midwestern Finance Association Conference, the 2024 Eastern Finance Association Conference, the 6th World Symposium on Investment Research, the 2023 FIRS Conference, and the 30th Finance Forum.

I. Introduction

Crises often begin with the awareness of a new risk whose ultimate consequences are not yet known. In early 2020, news of a novel virus circulated widely before the recession began, but investors, firms, and households did not yet know whether the outbreak would remain contained or develop into a global macroeconomic shock. Similarly, before the Global Financial Crisis, public attention to terms such as *"subprime," "housing crisis,"* and *"mortgage crisis"* rose before the full severity of the downturn became clear. Figure 1 illustrates this pattern using Google search activity around the Global Financial Crisis and the COVID-19 recession. These examples motivate the central idea of the paper: investors may become aware of a new aggregate risk before they fully understand its severity, persistence, and implications for future output.

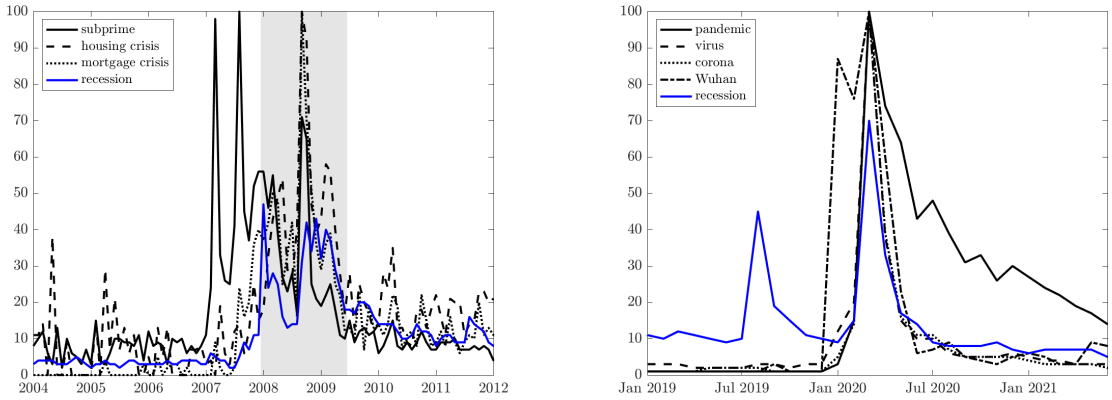

(a) The Great Recession

(b) The COVID-19 Pandemic

Figure 1. Awareness of New Risks. This figure displays Google Trends search interest for crisis terms related to two distinct events: the Global Financial Crisis in Panel (a) and the COVID-19 recession in Panel (b). Data come from Google Trends, which starts in 2004. Numbers represent search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for the term.

Asset prices are forward-looking, so one might expect markets to react immediately when a new recessionary risk becomes salient. Yet the empirical behavior of fundamentals and asset prices around recessions is gradual rather than instantaneous. Output typically does

not collapse on impact. It deteriorates with delay, reaches a trough, and then recovers with temporarily elevated growth. Financial markets display a related dynamic pattern: risk premia and return volatility tend to rise over the course of the downturn, while valuation ratios decline and recover only gradually. These patterns are difficult to reconcile with models in which crises are represented as one-time jumps in fundamentals or in which the arrival of bad news is priced fully and immediately. The central question of this paper is whether the dynamic shape of recessions itself is a priced state variable.

We use the term *awareness* to describe the arrival of a new priced aggregate risk. Awareness means that investors recognize the presence of a new risk factor that can affect future output, but it does not mean that they immediately know the size or persistence of the shock. This distinction is important. Awareness determines when the new risk becomes priced; learning determines how investors revise the perceived severity of that risk as the crisis unfolds. Thus, the paper studies not only the pricing of a new risk, but also the pricing of the evolving crisis shape that investors learn from realized macroeconomic outcomes.

We develop a tractable equilibrium model in which crises are episodes of a new aggregate risk whose effects unfold gradually. In normal times, output growth is i.i.d. Gaussian. When the economy enters the new-risk regime, output is affected by a crisis-impact process, η_t, that mean-reverts toward a crisis driver, x_t. The crisis driver follows a U-shaped path: it pushes output below its normal path during the downturn and then gradually returns toward normal during the recovery. This specification captures two robust features of recessions: delayed deterioration and subsequent bounce-back growth. It also allows crises to differ in depth and persistence through an episode-specific severity parameter, ε.

The proposed model features learning about this crisis-severity parameter. Investors know when the new-risk regime begins, but they do not initially observe ε. They infer it from the realized path of output. If output deteriorates more than expected, investors revise upward their estimate of crisis severity. These negative belief revisions raise risk aversion through the habit channel and increase discount rates. If the crisis turns out to be mild,

beliefs adjust in the opposite direction and the valuation effects are weaker. The learning mechanism therefore connects the gradual revelation of macroeconomic conditions to the gradual movement in risk premia, volatility, and valuation ratios.

The full-information economy is nested as a special case of the model. When posterior uncertainty about ε is zero at the onset of the crisis, investors know the expected crisis path immediately. This case isolates the direct pricing effect of the new crisis-risk source. It shows that the arrival of the crisis state introduces a new Brownian shock that affects aggregate consumption and becomes priced in equilibrium. As the crisis impact strengthens, both the exposure to this new risk and the market price of the risk rise; as the crisis fades, both decline. This mechanism generates hump-shaped risk premia and return volatility even without learning. Learning adds an additional force: when persistent bad outcomes reveal that the crisis is more severe than initially believed, risk aversion rises further and valuation ratios decline.

We discipline the model using macroeconomic crisis dynamics only. Specifically, we estimate the parameters governing the crisis driver and the crisis-impact process by simulated method of moments. The targeted moments summarize the local-projection response of the output gap to NBER recession onset, the timing of the output trough and recovery, the cross-crisis dispersion in these durations, the frequency of crisis periods, and the increase in output-growth volatility during crises. We do not target equity risk premia, stock-market volatility, or valuation ratios. Asset-pricing moments are therefore out-of-sample implications of the estimated macro crisis process.

The estimated model matches the gradual decline and recovery of output around recessions and generates asset-pricing dynamics that resemble the data. Risk premia are initially negative after recession onset, reflecting the fact that prices do not immediately incorporate the full severity of the downturn. As the recession unfolds and realized output growth disappoints relative to investors' beliefs, risk aversion rises and expected excess returns increase. This produces a delayed hump-shaped pattern: risk premia first fall, then rise as

macroeconomic conditions deteriorate, peak around the period of greatest stress, and decline during the recovery. Return volatility follows a similarly hump-shaped pattern, increasing gradually as the new crisis risk becomes more important and then reverting as the crisis dissipates. This timing differs from many non-expected-utility models in which risk compensation rises sharply upon the arrival of bad news or uncertainty (Ai and Bansal, 2018). In the learning economy, severe crises also generate declining valuation ratios. When investors learn that the crisis is more persistent than initially expected, realized growth arrives as worse-than-expected news, risk aversion increases, and the price-dividend ratio falls.

We document these patterns empirically using postwar U.S. data. Around NBER business-cycle peaks, output displays a delayed decline followed by recovery. Expected excess returns exhibit an S-shaped pattern: they are initially low or negative around the onset of the recession, rise gradually as the downturn unfolds and macroeconomic conditions deteriorate, and then decline during the recovery. Return volatility follows a related delayed increase and subsequent reversal, while valuation ratios fall and recover only gradually. Local projections confirm that the output gap, expected returns, and volatility all respond dynamically to recession onset rather than through immediate one-time adjustments. Survey of Professional Forecasters data provide direct support for the learning channel: forecast errors are negative around recession peaks and predict subsequent forecast revisions, indicating that expectations about real activity adjust gradually as new information arrives.

The paper contributes to the asset-pricing literature by emphasizing that markets price the shape of crises. Existing models often focus on the arrival of disasters, time-varying disaster probabilities, long-run risks, ambiguity, or habit formation. Our mechanism is different. A crisis is not modeled as an instantaneous collapse in fundamentals. Instead, it is a new-risk episode whose severity is learned over time from the realized path of output. This allows the model to jointly account for delayed output deterioration, hump-shaped risk premia, hump-shaped volatility, and recessionary valuation declines. The framework also clarifies why awareness of a new risk is not sufficient to generate recession-like asset-price

dynamics: large valuation effects arise when subsequent realizations reveal that the new-risk episode is severe and persistent.

This paper is organized as follows. Section II discusses the related literature. Section III documents the empirical behavior of output, survey expectations, risk premia, return volatility, and valuation ratios around recessions. Section IV presents the model. Section V describes the data. Section VI estimates the crisis-dynamics parameters and studies the model's asset-pricing implications. Section VII concludes.

II. Review of literature

Empirical facts. Recessions lead to gradual and prolonged declines in consumption (Barro and Ursúa, 2008). This trend also applies to asset prices. As noted by Muir (2017), asset prices during crises often follow a U-shaped trajectory, with cumulative returns declining by approximately 40%. However, about half of this loss is typically recovered within a few years. An analysis of 42 recessions in 14 countries since 1951 shows that both prices and dividends generally begin to decline at the start of a recession and remain significantly low even a dozen quarters after the recession ends (Kroencke, 2022).

The gradual, U-shaped response of output and other macroeconomic variables to recessions—rather than an immediate reaction—has been widely documented. Basu, Candian, Chahrour, and Valchev (2021) identify shocks to equity risk premia and show that output responds to these shocks in a similarly U-shaped pattern. Christiano, Eichenbaum, and Evans (2005) find that consumption and investment react in a hump-shaped manner following an expansionary monetary policy shock. In a Bayesian DSGE framework, Smets and Wouters (2007) generate hump-shaped responses of aggregate demand, driven by variable capital utilization and fixed production costs. Beaudry and Portier (2006) show that both consumption and stock prices respond with a pronounced hump shape to long-run TFP shocks. Finally, Brunnermeier, Palia, Sastry, and Sims (2021) estimate a structural VAR linking financial

variables to economic activity and report clear hump- or dip-shaped patterns.

Our model incorporates the U-shaped pattern of output and predicts a gradual price decline, followed by a recovery period. In our model, the arrival of crises is exogenous. Supporting this, Jordà, Schularick, and Taylor (2011) find it plausible that crises emerge unpredictably. They also note considerable variations in crises regarding their effects on output and consumption, as well as their duration, features that our model effectively mirrors.

Recent empirical work shows that financial markets appear to react sluggishly to the onset of recessions. Gómez-Cram (2022) documents that stock returns are predictably negative for several months following the start of a recession and become unusually high only afterward. Using a state-space model that identifies recession turning points in real time from macroeconomic data, the author shows that returns exhibit strong momentum during recessions but mild reversals during expansions. This pattern implies that investors tend to adjust their expectations about the business cycle only gradually. Supporting this interpretation, Gómez-Cram (2022) analyzes analyst earnings forecasts and finds that forecasts are excessively optimistic when expected returns are low at the onset of recessions and are subsequently revised downward sharply. These results suggest that stock prices incorporate recession-related information with a delay, generating the gradual price decline observed in the early stages of downturns.

Theoretical Perspectives on Recessions and Risk Premia. Kroencke (2022) shows that innovations in expected returns are highly volatile during recessions and illustrates that these facts are difficult to explain within standard asset pricing theories. Simulating "recessions" using frameworks like the Bansal and Yaron (2004) long-run risk model, the Campbell and Cochrane (1999) habit model, and the Wachter (2013) model of rare disasters, he observes that none of these models adequately capture the observed variances in stock prices or price changes.[1]

Risk premia are substantially higher in recessions than in expansions (Moench and Stein,

[1]For a detailed review of the implications of disaster risk for asset pricing, see Tsai and Wachter (2015).

2025; Muir, 2017; Lustig and Verdelhan, 2012). Muir (2017) adds that risk premia spike dramatically in financial crises, defined specifically as a banking panic or banking crises, but rise only modestly in recessions or wars. Muir (2017) argues that standard consumption-based asset pricing models do not reconcile these facts because the overall drop in consumption and the increase in consumption volatility is fairly similar across financial crises and recessions and is largest during wars.[2]

Our model generates and matches the heightened volatility observed during recessions. Specifically, we use the observed variance ratio between crisis and non-crisis periods as a key moment to estimate the model parameters driving each crisis in our sample. The model implies not only significant increases in risk premia but also elevated return volatility, aligning closely with empirical evidence.

Nakamura, Steinsson, Barro, and Ursúa (2013) estimate an empirical model of consumption disasters, which generates an equity premium from disaster risk that is substantially smaller than in disaster models. They conclude that an unrealistically large value of the inter-temporal elasticity of substitution is necessary to explain stock-market crashes at the onset of disasters. Gourio (2012) introduces time-varying disaster risk into a standard real business cycle model, which is also capable of generating a U-shaped reaction of macroeconomic variables and asset prices. His approach, however, relies on leverage to generate volatility of cash flows and returns, and it does not address the volatility of the unlevered return on capital.

Ghaderi, Kilic, and Seo (2022)'s model of slowly unfolding disasters explains the gradual response of asset prices to economic shocks through information processing. In their framework, agents learn about the time-varying intensity of consumption jumps, which increases during disasters, and the gradual recognition of a sustained transition to a recessionary state generates persistent asset-price effects. Our model also features learning, but the object of

[2]More recent literature offers clues on the potential mechanism driving the higher expected returns observed during recessions. Ai and Bhandari (2021) show that when idiosyncratic risk to human capital is not fully insurable, the anticipation of lack of risk sharing in the future can raise workers' current marginal utilities during recessions.

learning and the economic mechanism are different. Learning occurs after the economy enters a new-risk regime whose real effects unfold gradually through the crisis-impact process. Investors learn not merely whether a bad state has arrived or whether disaster intensity has increased, but how severe and persistent the realized crisis path will be. Thus, the gradual response of asset prices in our model is tied directly to the observed output dynamics of crises. Persistent downturn paths generate worse-than-expected macroeconomic outcomes, raise risk aversion, and produce valuation declines. This distinguishes our mechanism from standard learning frameworks in which the main force is the monotone resolution of uncertainty about a fixed latent state.

The post-crisis period is associated with abnormally high economic growth, also known as the bounce-back effect in level (Nakamura, Steinsson, Barro, and Ursúa, 2013; Kim, Morley, and Piger, 2005). Classical asset-pricing models, including the Ghaderi, Kilic, and Seo (2022)'s slowly unfolding disasters model, do not generate this recovery period. Beeler and Campbell (2012) show the long-run risk model produces persistence but not mean reversion in the level of consumption. Hasler and Marfè (2016) highlight the importance of recoveries that follow disaster events in explaining the observed shape of the term structures of equity return.

Our model contributes to the existing body of literature studying the relationship between crises and their ensuing impacts on asset prices and economic activity. We show that during economic recessions, the connection between asset prices and fundamentals becomes significantly stronger. Using a novel general equilibrium model, we explain why asset prices do not respond immediately to the introduction of new risks, even in the face of an anticipated economic slowdown. Our model predictions can quantitatively match the asset price reactions, which manifest in a hump-shaped pattern of expected returns and return volatility. The paper sets itself apart from other existing theories by demonstrating that the model can mirror actual recession dynamics, both in terms of changes in observed levels of output and risk premia.

III. Empirical facts

In this section, we document empirical stylized facts about the behavior of output and asset prices around crises. We first examine the dynamics of real activity and then turn to the evolution of asset prices and risk premia around the onset of recessions.

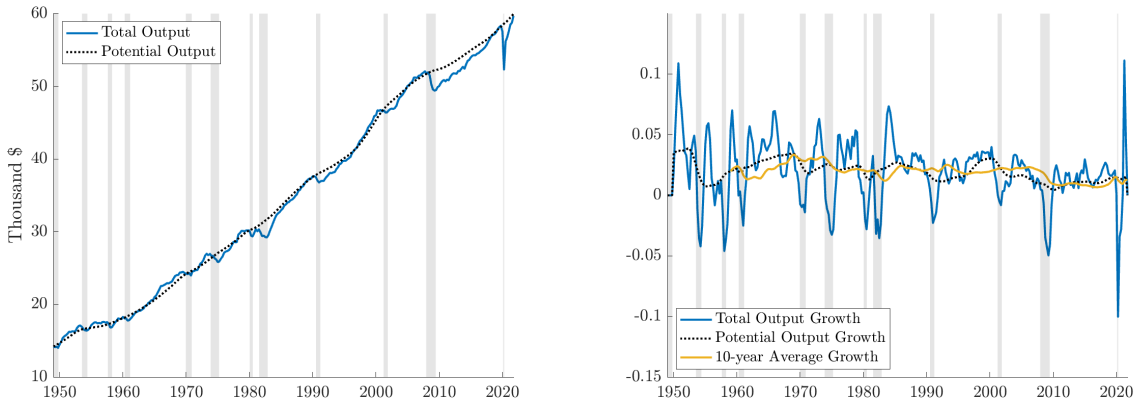

(a) Output Level (1949-2021)

(b) Output Growth (1949-2021)

Figure 2. Real GDP per capita (1949-2021) Panel (a) displays the evolution of the quarterly real and potential output in the United States (per capita) between 1949 to 2021. In panel (b), the figure shows the year-on-year change in the quarterly GDP and potential output (the blue and dashed black line, respectively) as well as the 10-year average total GDP growth rate (yellow line). Shaded areas represent NBER recessions.

A. The U-shaped output during crises

Figure 2 shows the evolution of US production over the past 70 years, highlighting that recessions have consistently created significant economic disruptions. These periods generate large output gaps due to substantial drops in total output, while there are no similar positive shocks. This asymmetric feature, first documented by Neftci (1984a), is a reason why the National Bureau of Economic Research (NBER) focuses on identifying and studying recessions. While this asymmetry is well-documented in the data, our model in Section IV explicitly incorporates it as a novel feature.

We use NBER business-cycle peaks as the observable counterpart of the subset of model

new-risk episodes that are ultimately classified as recessions. The NBER peak is not necessarily the first moment at which investors become aware of a risk. Rather, it marks the start of recessionary dynamics among episodes that deteriorate sufficiently and persist long enough to be classified as recessions. In practice, the NBER peak is likely to occur somewhat later than the true transition into the new-risk state in the model, since investors may become aware of a new aggregate risk before it is fully reflected in realized macroeconomic activity. Nevertheless, the NBER peak provides the cleanest and most consistently measured empirical proxy for the onset of recessionary dynamics. This timing is useful for our mechanism because the model predicts that output does not collapse on impact. Instead, the peak is followed by delayed deterioration, a trough, and subsequent recovery.

B. Asset prices around crises

How do financial markets behave around the onset of recessions? Figure 3 examines the dynamics of equity returns, return volatility, and stock market valuations around NBER business-cycle peaks. The figure aggregates postwar recession episodes and plots the average behavior of these variables from 12 months before to 24 months after the peak that marks the start of a recession.

Subsequent realized excess returns display a distinctive non-monotonic pattern. Excess returns decline in the months leading up to the business-cycle peak and reach their lowest levels around the onset of the recession. Only several months later do realized returns begin to rise, remaining elevated during the early phase of the downturn before gradually reverting toward their unconditional levels during the recovery. This produces a pronounced dip followed by a delayed hump-shaped increase in expected returns.

Return volatility increases initially at the onset of recessions, continues to rise and remains elevated during the early phase of the downturn before gradually declining as macroeconomic conditions stabilize. Stock market valuations, measured by the price–dividend ratio, start falling around the onset of recessions and typically reach their lowest levels roughly one year

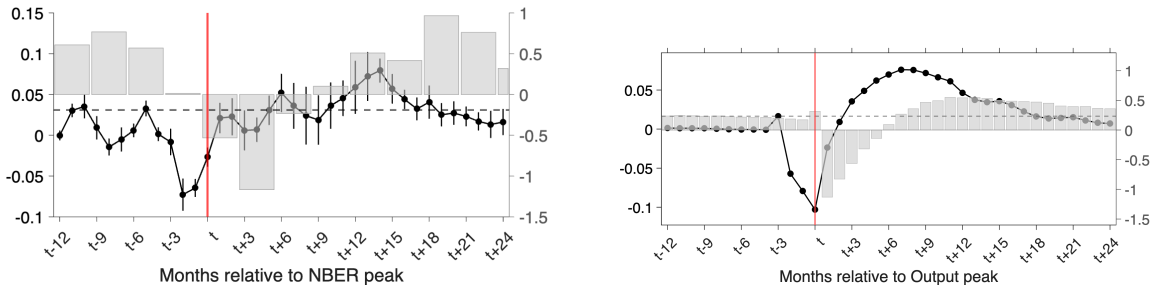

(a) Data: future 3-month excess returns

(b) Model: future 3-month excess returns

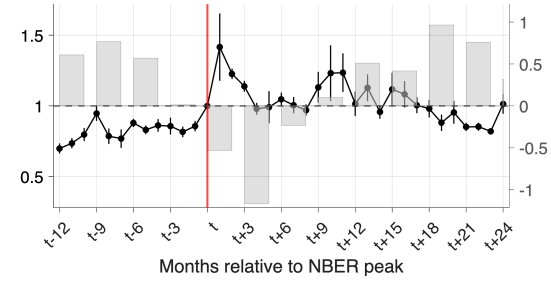

(c) Data: return volatility

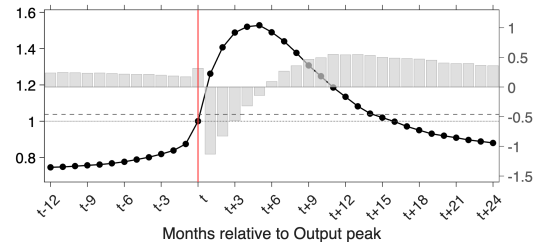

(d) Model: return volatility

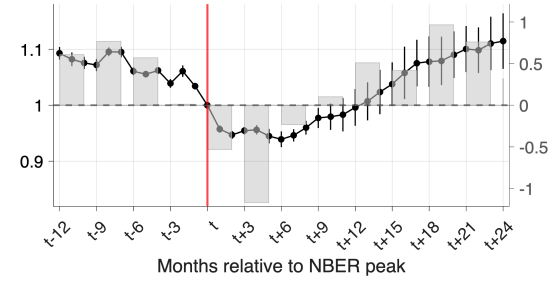

(e) Data: price-dividend ratio

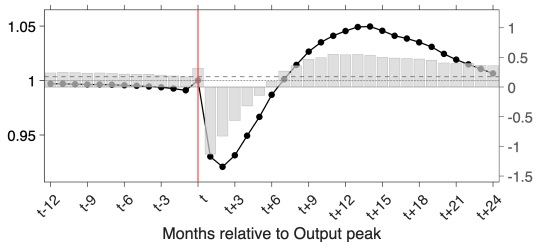

(f) Model: price-dividend ratio

Figure 3. Data vs. Model: Asset prices around crises. The figure compares asset-price dynamics around empirical NBER business-cycle peaks with the corresponding dynamics around simulated model crisis peaks. Panels (a), (c), and (e) use the data and report future realized 3-month excess returns, return volatility (measured using the square root value of the 'svar' variable from Goyal, Welch, and Zafirov (2024)), and the price–dividend ratio, respectively. Panels (b), (d), and (f) show the same objects in the model. In the simulations, crises arrive through Poisson jumps, and each simulated episode is aligned around its peak output date, the model counterpart of the NBER peak. Model parameters are described in the calibration section. Black dots report average values across episodes at each event month, and vertical lines show HAC standard errors, using 3 lags for the return panel and 12 lags for volatility and the price–dividend ratio. Return volatility and the price–dividend ratio are normalized to one at the peak month. The dashed horizontal line is the unconditional sample mean of the corresponding left-axis variable. Gray bars show output growth on the secondary right axis: quarterly real GDP growth in the data and monthly simulated output growth in the model. The horizontal axis measures months relative to the output level peak date. The prior mean is centered at the true crisis severity, $\hat{\varepsilon}_0 = \bar{\varepsilon} = 0.0166$, with prior variance $\nu_0 = \sigma_\varepsilon^2 = 0.3886^2$. The remaining model parameters are described in the calibration section.

after the peak, similar to the behavior of realized output. Valuations then recover gradually as economic conditions improve.

These patterns highlight three key empirical regularities: recessions are associated with sharp contractions in real activity, combined with increases in return volatility, and non-monotonic dynamics in realized returns.

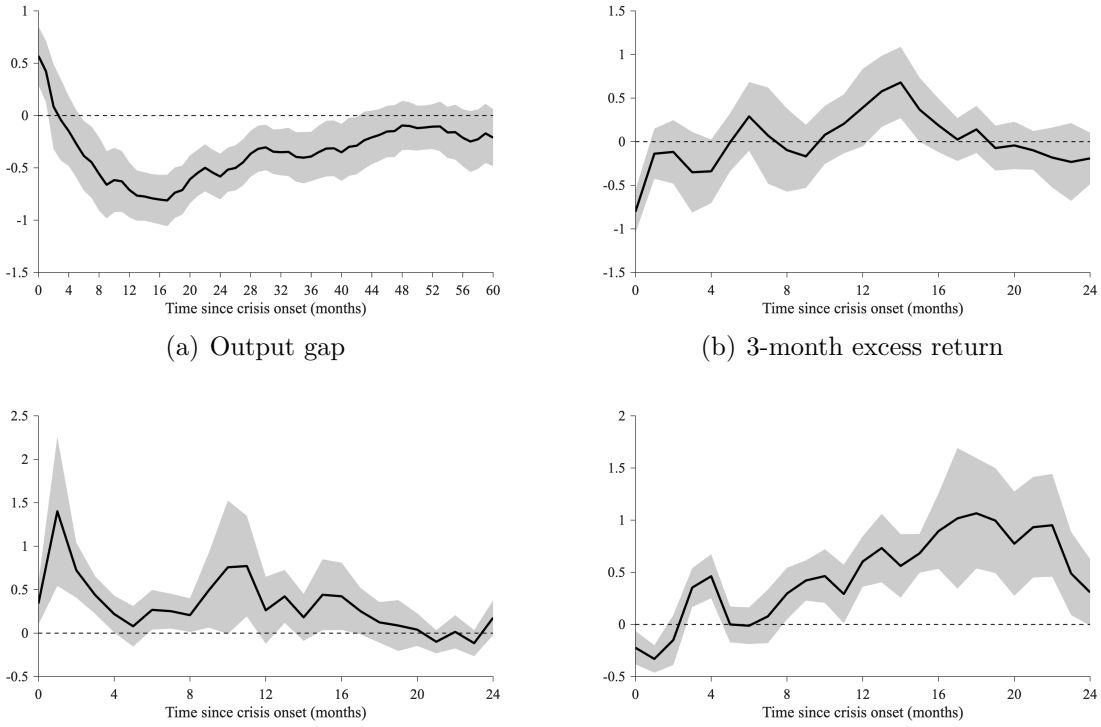

(a) Output gap

(b) 3-month excess return

(c) Return volatility

(d) Risk premium (Marfè and Pénasse, 2024)

Figure 4. Local projections. This figure reports impulse responses to the onset of recessions estimated using local projections (Jordà and Taylor, 2025). The shock is a dummy for NBER business-cycle peaks. Each panel plots the response of a different variable: the output gap (panel a), the 3-month excess equity return (panel b), return volatility (panel c), and the expected excess return measure of Marfè and Pénasse (2024) (panel d). The responses are estimated from regressions of the form $y_{t+h} = \alpha_h + \beta_h \text{Crisis}_t + \varepsilon_{t+h}$, where h denotes the horizon in months. The coefficient β_h measures the dynamic response of the variable y following the onset of a recession. The dots show the estimated impulse responses and the shaded areas denote 68% confidence bands based on heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. All variables are standardized prior to estimation so that magnitudes are comparable across panels.

C. *Local Projections*

To formally characterize these dynamics, we estimate local projections following Jordà and Taylor (2025). For each horizon h, we estimate

$$y_{t+h} = \alpha_h + \beta_h \text{Peak}_t + \varepsilon_{t+h},$$

where Peak_t is an indicator equal to one at the NBER peak date. The coefficient β_h traces the impulse response of variable y following the beginning of a crisis. We compute heteroskedasticity- and autocorrelation-consistent standard errors using Bartlett weights and standardize all variables prior to estimation to facilitate comparison across panels.

Figure 4 reports the resulting impulse responses. Real activity follows the familiar U-shaped pattern: the output gap declines sharply after the crisis begins, reaches its trough roughly one year later, and then gradually recovers. Expected returns exhibit delayed dynamics. Future realized excess returns decline initially but rise several months later, peaking when the output gap is most depressed before reverting toward normal levels. The expected excess return measure of Marfè and Pénasse (2024) displays a similar pattern. In contrast, return volatility increases at crisis onset and remains elevated for some time before gradually declining. These dynamics reinforce the stylized facts documented in the event-study evidence: crises generate a hump-shaped rise in return volatility and risk premia.

D. *Survey expectations about output growth*

Figure 5 examines the behavior of Survey of Professional Forecasters (SPF) expectations around NBER recession peaks. The figure plots conditional averages of forecast errors and forecast revisions relative to the recession peak quarter. Forecast errors are defined as the difference between realized growth and the SPF forecast for the same horizon.

The panels show a clear and systematic pattern around recessions. Forecast errors are strongly negative around the recession peak, indicating that forecasters initially underes-

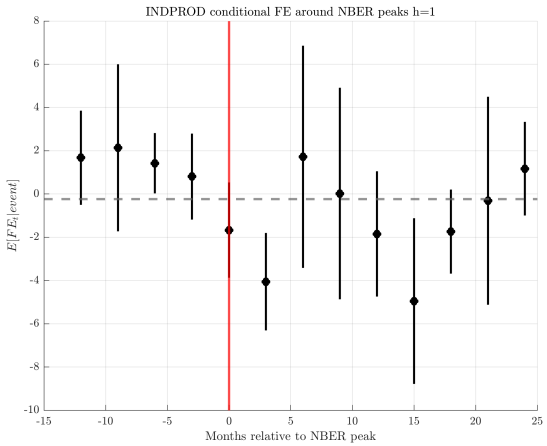

(a) Industrial production, horizon 1

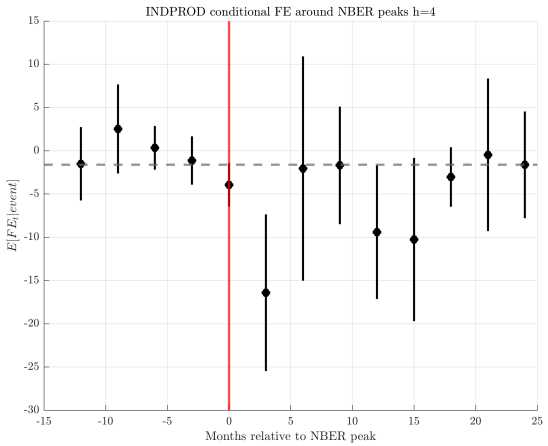

(b) Industrial production, horizon 4

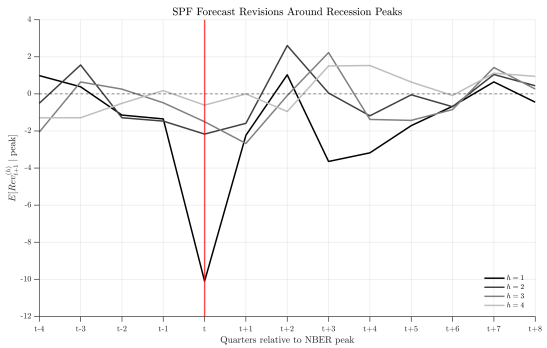

(c) Industrial production forecast revisions

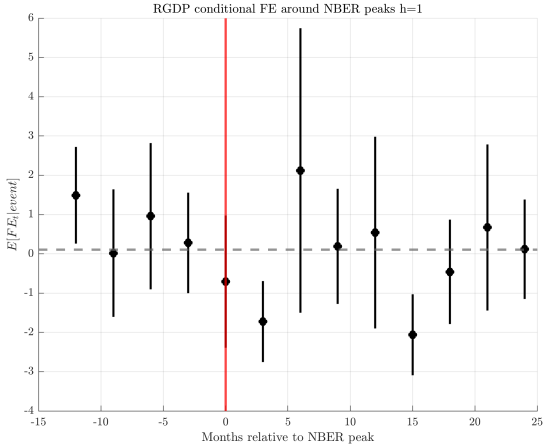

(d) Real GDP, horizon 1

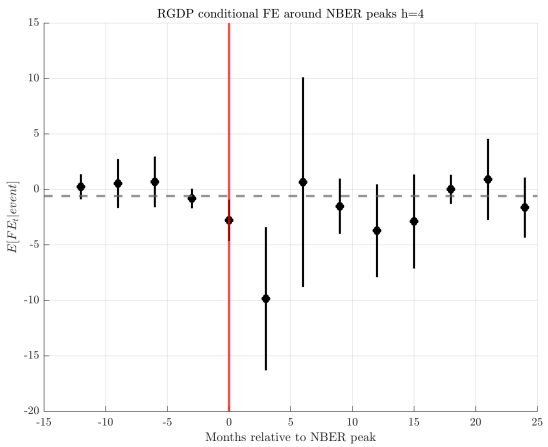

(e) Real GDP, horizon 4

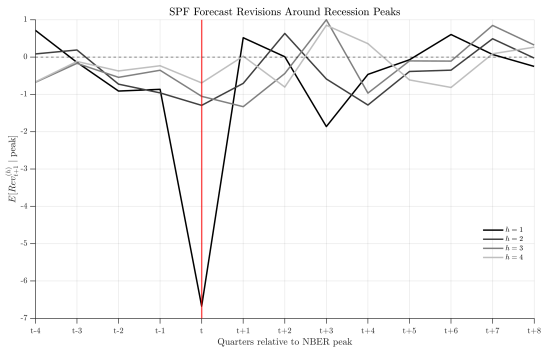

(f) Real GDP forecast revisions

Figure 5. Forecast errors and forecast revisions around NBER recession peaks.
The figure plots conditional averages of Survey of Professional Forecasters (SPF) forecast errors and forecast revisions around NBER recession peaks. Event time is measured in quarters relative to the peak ($t = 0$). Forecast errors are defined as $FE_t^{(h)} = \text{Realized}_t^{(h)} - \text{Forecast}_t^{(h)}$. Panels (a), (b), (d), and (e) show forecast errors for horizons $h = 1$ and $h = 4$; panels (c) and (f) show subsequent forecast revisions. The top row reports industrial production and the bottom row real GDP.

Table I: Forecast Errors and Future SPF Forecast Revisions: Data and Model

	Dependent variable: $Rev_{t+1}^{(h)}$			
	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Panel A. Industrial Production				
FE_t	0.390***	0.156*	0.154**	-0.043
	(3.24)	(1.82)	(2.11)	(-1.22)
Constant	-0.754***	-0.176	-0.219	0.014
	(-3.05)	(-0.89)	(-1.41)	(0.10)
N	226	225	224	218
R^2	0.090	0.042	0.069	0.008
Panel B. Real GDP				
FE_t	0.347**	0.102	0.119**	0.003
	(2.00)	(1.52)	(2.36)	(0.10)
Constant	-0.429**	-0.145	-0.238***	-0.103*
	(-2.52)	(-1.34)	(-2.63)	(-1.70)
N	225	224	223	217
R^2	0.054	0.022	0.053	0.000
Panel C. Model Simulations: Path-Level Time-Series Regressions				
FE_t	0.217***	0.165***	0.095***	0.049***
	(7.68)	(8.03)	(8.02)	(8.00)
Avg. R^2	0.203	0.179	0.168	0.164

This table reports predictive regressions of future fixed-target SPF forecast revisions on the current one-step-ahead forecast error:

$$Rev_{t+1}^{(h)} = \alpha_h + \beta_h FE_t + u_{t+1}.$$

For each horizon h, the dependent variable is the next-survey revision in the forecast for a fixed target quarter:

$$Rev_{t+1}^{(h)} = \text{Forecast}_{t+1}^{(h)} - \text{Forecast}_{t}^{(h+1)}.$$

The forecast error is the realized current-quarter growth rate minus the previous forecast for that same growth rate:

$$FE_t = \text{Realized}_t - \text{Forecast}_{t-1,t}.$$

Panels A and B use SPF forecasts and realized outcomes for Industrial Production and Real GDP. Panel C reports the corresponding model-implied evidence from simulated fixed-target output-growth forecasts. In Panel C, the regression is estimated separately within each simulated path. In Panels A and B, Newey–West t-statistics are reported in parentheses. In Panel C, the reported t-statistics test whether the mean path-level coefficient is different from zero, using Newey–West standard errors computed from the sequence of estimated path-level coefficients. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

A positive coefficient on FE_t implies that positive growth surprises today predict upward revisions in future forecasts, consistent with a belief-correction mechanism.

timate the severity of the downturn. In the quarters following the peak, forecast errors become positive as the economy recovers more quickly than expected. This pattern is visible for both industrial production and real GDP and is particularly pronounced at longer horizons. The panels that plot forecast revisions show a similar dynamic: forecasts are revised downward around the recession peak and subsequently revised upward as the recovery unfolds. Taken together, the figures indicate that survey expectations adjust gradually as new macroeconomic information arrives.

Table I provides complementary regression evidence on the relationship between forecast errors and subsequent forecast revisions. Specifically, we estimate predictive regressions of future SPF forecast revisions on the current one-step-ahead forecast error. Across industrial production and real GDP, positive growth surprises today tend to predict upward revisions in future forecasts. The coefficient on the forecast error is positive and statistically significant for several horizons, particularly for industrial production and real GDP at short and intermediate horizons.

The model reproduces this same belief-correction pattern. In Panel C, we construct fixed-target output-growth forecasts from the simulated learning economy and estimate the same regression path by path. The mean path-level coefficient on FE_t is positive at all horizons and statistically significant using Newey–West standard errors computed from the sequence of simulated path-level estimates. The coefficients decline with the forecast horizon, consistent with the idea that current growth surprises are most informative about near-term revisions but still contain information about future output dynamics. Thus, the model does not merely generate gradual forecast errors around recessions; it also matches the regression-based evidence that unexpected output realizations lead agents to revise their beliefs about the future path of the economy.

This evidence supports the learning mechanism in the model. When realized growth is stronger than expected, agents infer that the crisis is less severe or less persistent than previously believed and revise future growth forecasts upward. Conversely, negative growth sur-

prises lead to downward revisions. The positive forecast-error coefficient therefore provides direct empirical and simulated evidence that expectations adjust gradually as information about crisis severity is revealed over time.

Overall, the evidence suggests that professional forecasters update expectations gradually in response to realized macroeconomic outcomes. Forecast errors systematically predict future revisions, consistent with a belief-correction mechanism in which agents learn about the severity and persistence of macroeconomic downturns over time. This evidence motivates the information structure in our proposed model: investors observe that a new-risk episode has begun, but they infer its severity only gradually from realized output dynamics.

IV. The Model

In this section, we develop an exchange economy in which asset prices respond to the dynamic shape of crises. The model is designed to capture three features of downturns. First, output does not fall instantaneously at the beginning of a crisis. Instead, the downturn unfolds gradually, reaches a trough, and is followed by a recovery. Second, crises differ in severity and persistence. Some episodes are shallow and short-lived, while others produce large and long-lasting output losses. Third, investors do not know the ultimate severity of a crisis when it begins. They learn about it from the realized path of output.

The central state variable is a crisis-shape process that governs the timing and severity of downturns and recoveries. A crisis begins when the economy enters a new-risk regime. At that time, a crisis-severity parameter, denoted by ε, determines how strongly the crisis driver pushes output below its normal path. The representative agent does not observe ε directly. Instead, she observes the realized evolution of output and updates her beliefs about the severity and persistence of the crisis.

This learning channel is the main information structure in the model. The version without learning is nested as a special case in which uncertainty about ε is zero at the beginning of

the crisis. In that case, the representative agent knows the crisis path immediately, and the model collapses to the full-information benchmark.

The mechanism is straightforward. During crises, output is affected by a new source of risk. This new risk changes both expected output growth and output volatility. Through the stochastic discount factor, asset prices load on this state. As the downturn deepens, compensation for crisis risk rises and return volatility increases. As the crisis dissipates, both decline. This generates hump-shaped risk premia and return volatility.

Learning adds a second force. When investors do not know the severity of the crisis, realized output contains information about the future path of the downturn. If the crisis turns out to be worse than initially expected, realized growth is persistently disappointing. These negative belief revisions raise risk aversion, amplify discount-rate movements, and generate declines in valuation ratios. Hence, the model links the dynamic shape of recessions to the dynamic shape of risk premia, volatility, and valuations.

A. Output

Our model's main innovation is incorporating output patterns during crises into the output process via two interconnected processes: η_t, defined below in Equation (3) and referred to as the crisis impact, and x_t, defined below in Equation (4) and referred to as the crisis driver. The crisis driver (x_t) captures the expected trajectory of the economic disruption due to the new source of risk,[3] while the crisis impact, η_t, captures the realized propagation of the disruption into output. Specifically, we decompose the output as follows.

$$Y_t = \hat{Y}_t \eta_t, \tag{1}$$

where $\hat{Y}_t$ represents output under normal economic conditions (e.g., $\hat{Y}_t$ can be viewed as potential output when firms operate at full capacity). The process $\hat{Y}_t$ follows a geometric

[3]We use the term 'new risks' to refer to crisis-specific shocks that alter the path of output and risk premia. These are modeled via crisis-specific calibrations to capture observed heterogeneity across recessions.

Brownian motion

$$d\hat{Y}_t = \mu_{\hat{Y}} \hat{Y}_t dt + \sigma_{\hat{Y}} \hat{Y}_t dZ_{\hat{Y},t}, \tag{2}$$

where $\mu_{\hat{Y}}$ and $\sigma_{\hat{Y}}$ are constants and $Z_{\hat{Y},t}$ is a Brownian shock.

The crisis impact, η_t, is a strictly positive stochastic process that depends on a two-state continuous-time Markov process $\omega_t \in \{H, L\}$. In normal economic times, $\omega_t = H$, $\eta_t = 1$, and thus output growth is i.i.d. normally distributed. When a crisis occurs at random time s, $\omega_t = L$, η_t becomes stochastic and typically falls below one, reflecting the negative impact on output Y_t.

Transitions from the normal state H to the crisis state L occur at a random arrival time s that is exponentially distributed with intensity parameter ν. Once a crisis begins at time s, the crisis impact, η_t, evolves for all $t \geq s$ according to:

$$d\eta_t = \kappa_\eta \left(x_t - \eta_t\right) dt + \sigma_\eta \eta_t \left(\lambda - \eta_t\right) dZ_{\eta,t}, \tag{3}$$

where $\kappa_\eta > 0$, $\sigma_\eta > 0$, $\lambda > 1$, and $Z_{\eta,t}$ is a Brownian shock. The instantaneous volatility term, $\sigma_\eta \eta_t (\lambda - \eta_t)$, ensures that η_t remains within the interval $(0, \lambda)$, given that $\lambda > 1$. The initial and terminal conditions for the crisis-impact process are $\eta_s = \eta_\tau = 1$, where τ denotes the end of the crisis defined in Equation (5).

The crisis impact, η_t, is reverting towards the crisis driver, x_t, where x_t is given by

$$x_t = 1 - \left(e^{-\kappa_1 (t-s)} - e^{-\kappa_2 (t-s)}\right) \varepsilon, \quad \forall s \leq t \leq \tau \tag{4}$$

where $0 < \kappa_1 < \kappa_2$. The crisis-severity variable ε is drawn at the beginning of each crisis episode from a normal distribution with mean $\bar{\varepsilon}$ and standard deviation σ_ε, and remains fixed throughout that episode. If the realized draw of ε is negative, the crisis driver lies above one rather than below one, pushing η_t back toward its normal value. Such episodes therefore exit the crisis state quickly once the stopping condition in equation (5) is met, and

can be interpreted as new-risk episodes that do not develop into realized recessions.

The crisis driver x_t follows a U-shaped pattern: it starts at one, decreases to a minimum, and then reverts back towards one, as shown in Figure 6. This reflects temporary output destruction during crises, followed by above-average growth during recovery.[4]

The crisis ends at the stopping time τ, defined as:

$$\tau = \inf\{\, t > s + \Delta, \quad \text{s.t.} \quad \eta_t \geq 1\}, \tag{5}$$

where the strictly positive parameter Δ ensures that η_t does not immediately return to 1 at time s, preventing the emergence of a crisis. Intuitively, Δ captures the time it takes for the crisis to fully impact the economy. The timeline below describes the evolution of output during a crisis starting at time s.

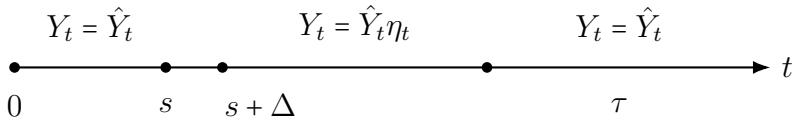

There are several important aspects of the crisis-impact process in Equation (3). First, when the economy enters the crisis state L, a new Brownian shock, $Z_{\eta,t}$, becomes active. This shock affects aggregate output and is therefore priced in equilibrium. Second, both expected output growth and output volatility become state-dependent during crises. Third, output volatility acquires an additional crisis-specific component. Because $\eta_s = 1$ at the onset of the crisis and λ is close to one, this component is initially small, but it grows as η_t moves away from one. Finally, setting $\lambda > 1$ helps ensure that the crisis-impact process returns to one, so the economy exits the crisis state in finite time τ.

Figure 6 shows the crisis driver, x_t (red line), the average path of the crisis impact, η_t (blue line), and output, Y_t (black dotted line with shaded 5-95% confidence bands). Because

[4] Gârleanu and Panageas (2015) use a similar approach with the sum of two negative exponential functions to model the hump-shaped pattern of earnings, while Blanchard (1985) applies the same functional form to capture complex household income paths in Footnote 8.

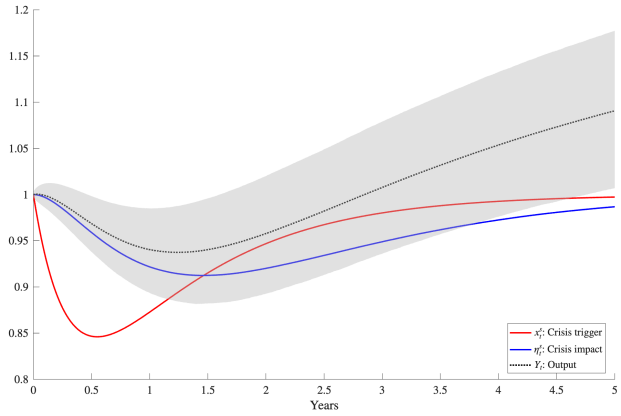

Figure 6. Crisis-driver (x_t) and crisis-impact (η_t) variables. The red line shows the crisis-driver variable x_t that responds to an initial shock $\varepsilon = 0.4$. The blue line describes the output reaction to the crisis, represented by η_t. The shaded area represents the 5-95% confidence bands for the simulated Y_t paths. The parameters are set to $\mu_{\hat{Y}} = 0.02$, $\sigma_{\hat{Y}} = 0.02$, $\kappa_\eta = 0.9$, $\lambda = 1.05$, $\sigma_\eta = 0.3$, $\kappa_1 = 1$ and $\kappa_2 = 3$. Δ is 0.5 years.

η_t mean-reverts toward the U-shaped driver x_t, the crisis impact inherits a delayed decline and gradual recovery.

U-shaped output dynamics can arise when recessions trigger gradual adjustment and propagation, heightened uncertainty and partial irreversibility depress investment and hiring beyond the onset, while financial frictions and balance-sheet constraints amplify and prolong the downturn; as uncertainty resolves and constraints relax, deferred spending and capital deepening generate a rebound with temporarily above-trend growth (e.g., Bloom (2009); Bernanke, Gertler, and Gilchrist (1999); Kiyotaki and Moore (1997)). Empirically, evidence on asymmetric business cycles and bounce-back recoveries is consistent with this delayed-trough and accelerated-recovery pattern (e.g., Neftci (1984b); Sichel (1993); Morley and Piger (2012)).

This delayed reaction of output in our model matches the dynamics of output around crises. In data, we observe that the drop in realized output is not instantaneous. In fact, in Figure 2 we show that output drops in a gradual manner. This gradual drop in output is followed by the 'bounce-back' effect, where the abnormal growth becomes positive. Both features of the aggregate output data are consistent with our model.

The left plot of Figure 7 shows the crisis driver x_t (red line), while the right plot shows the average path of the crisis impact, η_t, (blue line with shaded 5-95% confidence bands) for different initial shocks $\varepsilon \in \{0.05, 0.2, 1\}$. Larger initial shocks lead to more severe and prolonged recessions. The response of x_t to the initial shock is sharper than that of η_t, which is expected since η_t adjusts to x_t over time, depending on κ_η.

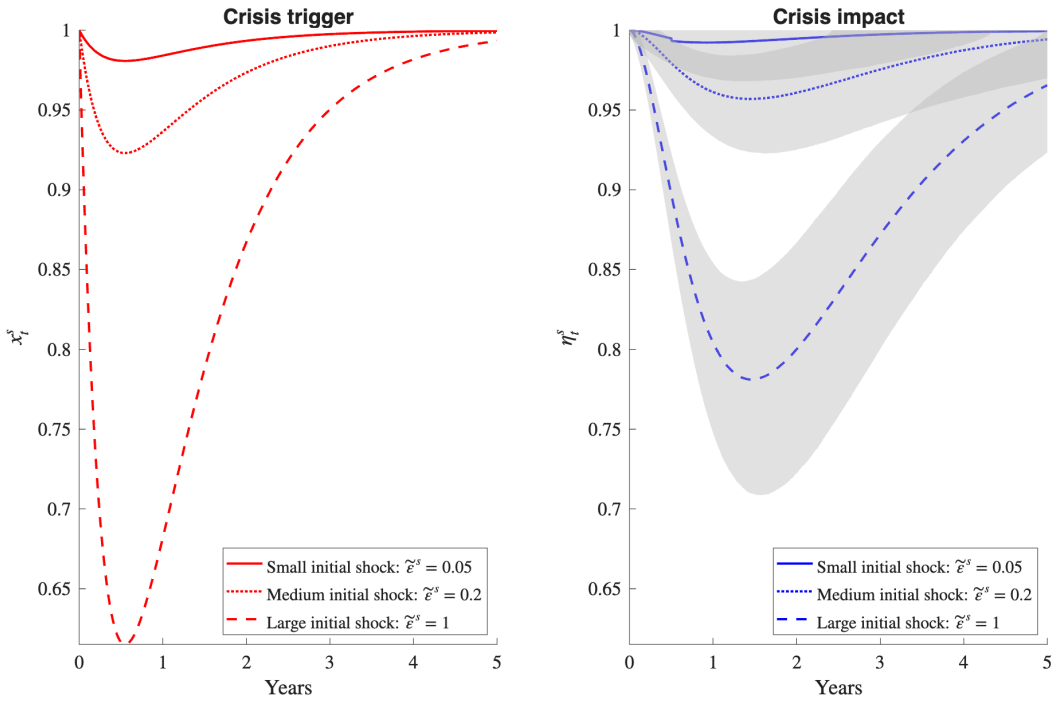

Figure 7. Crisis driver and crisis impact for different initial crisis shocks. The red line shows x_t and the blue line depicts η_t for different levels of the shock ε, that is, 0.05, 0.2, and 1, with shaded areas around η_t representing 5-95% confidence bands. The parameters are set to $\kappa_\eta = 0.9$, $\lambda = 1.05$, $\sigma_\eta = 0.5$. Δ is 0.5 years.

Figure 8 illustrates how the parameter κ_η (in panel (a)) influences the behavior of the crisis impact η_t. The red line shows the crisis driver x_t, the blue solid line shows the average path of the crisis impact, η_t, when $\kappa_\eta = 0.5$ (with shaded 5-95% confidence bands), and the blue dashed line shows the average path of the crisis impact, η_t, (with shaded 5-95% confidence bands) when $\kappa_\eta = 3$, $\kappa_1 = 0.5$ and $\kappa_2 = 3.4$. When κ_η is higher, η_t tracks the path of the crisis driver x_t more closely, leading to faster recoveries and shorter recessions. In contrast, lower values of κ_η result in slower adjustments of η_t to x_t, prolonging the recession

and delaying recovery.

Increasing the levels of κ_1, κ_2, and their ratio (κ_2/κ_1) also impacts the severity and duration of crises, similarly to an increase in κ_η, as shown in panel (b) of Figure 8. This indicates that all three κ parameters are important determinants of crisis characteristics, which can differ significantly across recessions.

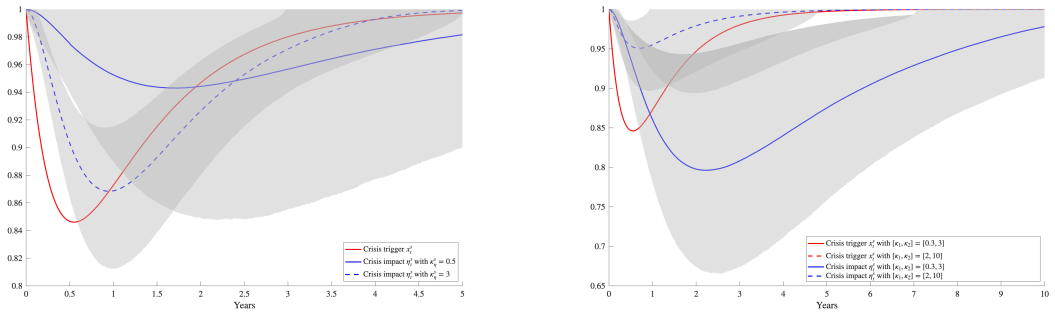

(a) κ_η (b) κ_1 and κ_2

Figure 8. Crisis driver and crisis impact for different κ_i parameters. Panel (a) displays x_t and η_t for different levels of κ_η, that is, 0.5 and 3, with shaded areas around η_t representing 5-95% confidence bands, with κ_1 = 1 and κ_2 = 3. Panel (b) shows x_t and η_t for different levels of $[\kappa_1, \kappa_2]$, that is, [0.3,3] and [2,10], with κ_η = 0.9. The remaining parameters are set to ε = 0.4, λ = 1.05, σ_η = 0.5, in both panels. Δ is 0.5 years.

Deeper and more persistent crises, which cause long-lasting damage, can be generated in this model by a combination of a large initial shock ε and low κ parameters. This aligns with empirical evidence that severe crises often take longer to resolve. By incorporating these dynamics, the model allows us to distinguish between different types of recession events. For example, financial crises often exhibit greater severity and persistence compared to non-financial crises, and these distinctions are reflected in their varying impacts on asset prices (Reinhart and Rogoff, 2009).

B. Learning About Crisis Severity

Our model features incomplete information about the episode-specific severity of a crisis. At the time a crisis begins, investors know that a new source of risk has arrived, but they do

not know how severe or persistent the episode will ultimately be. They observe the realized path of output and use it to update their beliefs about the crisis-severity parameter.

This section formalizes that learning problem. The key object is the parameter ε, which determines the depth of the crisis driver x_t. A larger value of ε pushes the crisis driver farther below one and therefore generates a deeper and more persistent downturn. The representative agent learns about ε from the observed evolution of the crisis impact η_t.

Prior beliefs and the crisis driver. At the crisis onset time s, the crisis-severity parameter is drawn once and remains fixed throughout the episode:

$$\varepsilon \sim \mathcal{N}(\bar{\varepsilon}, \sigma_\varepsilon^2), \qquad \hat{\varepsilon}_s = \bar{\varepsilon}, \qquad \nu_s = \sigma_\varepsilon^2, \tag{6}$$

where

$$\hat{\varepsilon}_t \equiv \mathbb{E}_t(\varepsilon \mid \mathcal{F}_t)$$

is the posterior mean and

$$\nu_t \equiv \mathbb{E}_t[(\varepsilon - \hat{\varepsilon}_t)^2 \mid \mathcal{F}_t]$$

is the posterior variance. The posterior mean captures the agent's current assessment of crisis severity, while the posterior variance captures how uncertain she remains about that assessment. Negative realizations of ε correspond to new-risk episodes that do not develop into recessions because the crisis driver lies above one and pushes η_t back toward its normal value.

During the crisis, for $t \in [s, \tau]$, the crisis driver is linear in the unknown severity parameter:

$$x_t = 1 - g_t\,\varepsilon, \qquad g_t \equiv e^{-\kappa_1(t-s)} - e^{-\kappa_2(t-s)}, \qquad 0 < \kappa_1 < \kappa_2. \tag{7}$$

The deterministic function g_t controls the timing of the crisis. It is zero at the onset of the crisis, rises as the downturn develops, and then declines as the recovery begins. The scalar ε

controls the magnitude of this path. Thus, learning about ε is equivalent to learning about the entire crisis shape.

The agent observes output,

$$Y_t = \widehat{Y}_t \eta_t,$$

and therefore observes η_t after controlling for the normal output component $\widehat{Y}_t$. Because η_t mean-reverts toward the crisis driver x_t, the realized path of η_t reveals information about ε. If η_t falls more than expected, the agent revises her estimate of ε upward. If η_t recovers more quickly than expected, she revises it downward.

Observation equation. To express the learning problem in linear Gaussian form, define the normalized signal

$$ds_t \equiv \frac{d\eta_t}{\sigma_\eta \eta_t (\lambda - \eta_t)}. \tag{8}$$

Using the law of motion for η_t, this signal can be written as

$$ds_t = \Big(A_0(t) + A_1(t)\varepsilon \Big)\, dt + dZ_{\eta,t}, \tag{9}$$

where

$$A_0(t) = \frac{\kappa_\eta (1 - \eta_t)}{\sigma_\eta \eta_t (\lambda - \eta_t)}, \qquad A_1(t) = -\frac{\kappa_\eta g_t}{\sigma_\eta \eta_t (\lambda - \eta_t)}. \tag{10}$$

This is a standard continuous-time filtering problem with a static hidden state. The unknown state is the scalar ε, and the agent observes a noisy signal whose loading on ε is time-varying.

The loading $A_1(t)$ has a useful interpretation. Learning is most informative when the crisis driver is sensitive to ε, that is, when g_t is large. Learning is also stronger when the effective noise in the observation equation,

$$\sigma_\eta \eta_t (\lambda - \eta_t),$$

is small. Hence, the speed of learning depends both on the shape of the crisis and on the

volatility of the observed crisis-impact process.

Innovation and filter dynamics. Define the innovation as the unexpected component of the normalized signal:

$$d\widetilde{Z}_t^{\eta} \;\equiv\; ds_t - \big(A_0(t) + A_1(t)\hat{\varepsilon}_t\big)\,dt. \tag{11}$$

Equivalently, in terms of the observed process η_t,

$$d\widetilde{Z}_t^{\eta} = \frac{d\eta_t - \kappa_\eta\big(1 - g_t\hat{\varepsilon}_t - \eta_t\big)\,dt}{\sigma_\eta\eta_t(\lambda - \eta_t)}. \tag{12}$$

The innovation is the part of realized output dynamics that the agent did not expect given her current belief $\hat{\varepsilon}_t$. It is positive when the observed evolution of η_t is better than expected and negative when the crisis is unfolding worse than expected.

Applying the continuous-time Kalman–Bucy filter for a static parameter gives

$$d\hat{\varepsilon}_t = \nu_t A_1(t)\,d\widetilde{Z}_t^{\eta} = -\nu_t\frac{\kappa_\eta g_t}{\sigma_\eta\eta_t(\lambda - \eta_t)}\,d\widetilde{Z}_t^{\eta}, \tag{13}$$

$$d\nu_t = -\nu_t^2 A_1(t)^2\,dt = -\nu_t^2\left(\frac{\kappa_\eta g_t}{\sigma_\eta\eta_t(\lambda - \eta_t)}\right)^2 dt. \tag{14}$$

These equations are intuitive. The posterior mean changes only when realized output dynamics differ from what the agent expected. The size of the revision is proportional to posterior uncertainty ν_t: when the agent is unsure about crisis severity, she updates strongly; when she is confident, she updates little. Posterior variance declines over time as information accumulates.

The agent's perceived crisis driver is therefore

$$\hat{x}_t \equiv \mathbb{E}_t[x_t \mid \mathcal{F}_t] = 1 - g_t\hat{\varepsilon}_t, \qquad \mathrm{Var}_t(x_t) = g_t^2\nu_t. \tag{15}$$

Thus, the posterior mean $\hat{\varepsilon}_t$ determines the expected crisis path, while the posterior variance ν_t determines uncertainty about that path.

The full-information special case. The economy without learning is obtained by setting the initial posterior variance to zero:

$$\nu_s = 0.$$

Then equation (14) implies $\nu_t = 0$ for all $t \geq s$, and equation (13) implies that the posterior mean is constant. If the prior is centered on the true value, then

$$\hat{\varepsilon}_t = \varepsilon \qquad \text{and} \qquad \hat{x}_t = x_t$$

throughout the crisis. In this case, the representative agent knows the full crisis path immediately. The no-learning model is therefore the degenerate-belief case of the learning model.

This nesting is important for interpretation. The full-information case isolates the direct pricing effect of a new crisis-risk source. The learning case adds belief revisions about the severity and persistence of that crisis. The two mechanisms are complementary: the new risk source generates hump-shaped risk premia and volatility, while learning generates stronger and more persistent movements in risk aversion and valuation ratios.

C. Asset Pricing

We now derive equilibrium asset prices. The representative agent prices assets using a stochastic discount factor that depends on consumption and external habit. Since consumption equals aggregate output in equilibrium, the pricing kernel depends on the output process and on the agent's beliefs about crisis severity.

In this economy, the representative agent observes that the economy has entered the new-risk regime, but she does not directly observe the episode-specific severity parameter ε. She forms the posterior mean $\hat{\varepsilon}_t$ and posterior variance ν_t using the filtering equations in

Section IV.B. Her perceived crisis driver is

$$\hat{x}_t = 1 - g_t \hat{\varepsilon}_t,$$

whereas the true crisis driver is

$$x_t = 1 - g_t \varepsilon.$$

The full-information economy is nested as the special case in which posterior uncertainty is zero, $\nu_t = 0$, and the prior is centered on the true value, so that $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$.

Throughout this section, tildes denote objects computed under the representative agent's beliefs. These are the relevant pricing objects in the learning economy. Objective dynamics matter because realized belief errors affect the path of risk aversion, but asset prices are formed using the agent's perceived conditional distribution. The propositions below therefore characterize equilibrium prices in the learning economy, with the full-information case stated as a special case.

All proofs are provided in Appendix A.

C.a. Preferences

We consider a representative agent who maximizes

$$\mathbb{E}_0\left[\int_0^\infty u(C_t, H_t, t)\, dt\right] = \mathbb{E}_0\left[\int_0^\infty e^{-\rho t} \log\left(C_t - H_t\right)\, dt\right], \tag{16}$$

where C_t is consumption, ρ is the time discount rate, and

$$H_t = C_t(1 - 1/\mathcal{R}_t)$$

is an external habit. We refer to $\mathcal{R}_t$ as the local curvature of the utility function

$$\mathcal{R}_t \equiv -\frac{u_{CC}(C_t, H_t, t)C_t}{u_C(C_t, H_t, t)} = \frac{C_t}{C_t - H_t}. \tag{17}$$

Equivalently, $1/\mathcal{R}_t$ is the surplus consumption ratio. When consumption is close to the habit level, the surplus consumption ratio is low, $\mathcal{R}_t$ is high, and marginal utility is especially sensitive to consumption shocks.

Rather than modeling the surplus consumption ratio as in Campbell and Cochrane (1999), we follow Menzly, Santos, and Veronesi (2004) and let risk aversion evolve as a mean-reverting process that is locally negatively correlated with unexpected consumption growth. In the learning economy, unexpected consumption growth is measured relative to the agent's perceived conditional mean. During crises, this perceived mean depends on the posterior expected crisis driver $\hat{x}_t$.

Specifically, risk aversion evolves according to

$$d\mathcal{R}_t = \kappa_{\mathcal{R}} \left(\bar{\mathcal{R}} - \mathcal{R}_t\right) dt - \alpha \left(\mathcal{R}_t - \lambda_{\mathcal{R}}\right) \left(\frac{dC_t}{C_t} - \widetilde{\mathbb{E}}_t\left[\frac{dC_t}{C_t}\right]\right), \tag{18}$$

where $\widetilde{\mathbb{E}}_t$ denotes expectation under the representative agent's information set and beliefs. The parameter $\bar{\mathcal{R}} > \lambda_{\mathcal{R}} > 0$ is the long-run mean of risk aversion, and $\lambda_{\mathcal{R}}$ is a lower bound. The parameter $\kappa_{\mathcal{R}} > 0$ controls the speed of mean reversion. A smaller $\kappa_{\mathcal{R}}$ makes risk aversion more persistent. The parameter $\alpha > 0$ controls how strongly negative consumption news raises risk aversion. A larger α implies that risk aversion responds more strongly to unexpected consumption shocks.

During crises, output growth under the objective measure satisfies

$$\frac{dY_t}{Y_t} = \left(\mu_{\widehat{Y}} + \kappa_\eta\left(\frac{x_t}{\eta_t} - 1\right)\right) dt + \sigma_{\widehat{Y}}\, dZ_{\widehat{Y},t} + \sigma_\eta \left(\lambda - \eta_t\right) dZ_{\eta,t}, \qquad (\omega_t = L). \tag{19}$$

The representative agent, however, forms conditional expectations using the perceived crisis driver $\hat{x}_t$:

$$\widetilde{\mathbb{E}}_t\left[\frac{dY_t}{Y_t}\right] = \left(\mu_{\widehat{Y}} + \kappa_\eta\left(\frac{\hat{x}_t}{\eta_t} - 1\right)\right) dt, \qquad (\omega_t = L). \tag{20}$$

Therefore, the realized innovation relative to beliefs is

$$\frac{dY_t}{Y_t} - \widetilde{\mathbb{E}}_t\left[\frac{dY_t}{Y_t}\right] = -\kappa_\eta\left(\frac{g_t(\varepsilon - \hat{\varepsilon}_t)}{\eta_t}\right)dt + \sigma_{\widehat{Y}}\, dZ_{\widehat{Y},t} + \sigma_\eta\,(\lambda - \eta_t)\, dZ_{\eta,t}, \qquad (\omega_t = L). \tag{21}$$

The first term is the belief-error component. If the agent underestimates crisis severity, $\hat{\varepsilon}_t < \varepsilon$, then she overestimates the crisis driver, $\hat{x}_t > x_t$, and expects output growth to be stronger than it actually is. Realized output growth then arrives as worse-than-expected news.

Proposition 1 (Risk-aversion dynamics under learning): *During a crisis, $\omega_t = L$, the objective dynamics of risk aversion are*

$$\begin{aligned} d\mathcal{R}_t = &\left[\kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t) + \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\kappa_\eta\left(\frac{g_t(\varepsilon - \hat{\varepsilon}_t)}{\eta_t}\right)\right]dt \\ &- \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_{\widehat{Y}}\, dZ_{\widehat{Y},t} - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_\eta\,(\lambda - \eta_t)\, dZ_{\eta,t}. \end{aligned} \tag{22}$$

The additional drift term in equation (22) is the key learning channel. When the agent underestimates crisis severity, realized output growth is persistently worse than expected. These negative belief errors raise risk aversion and increase discount rates. As learning progresses, $\hat{\varepsilon}_t$ moves closer to ε, posterior variance falls, and the additional drift effect weakens.

Corollary 1 (Full-information special case): *If posterior uncertainty about crisis severity is zero, $\nu_t = 0$, and the prior is centered on the true value, then $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$. The belief-error term in equation* (22) *disappears, and risk aversion follows the full-information dynamics*

$$d\mathcal{R}_t = \kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t)dt - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_{\widehat{Y}}\, dZ_{\widehat{Y},t} - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_\eta\,(\lambda - \eta_t)\, dZ_{\eta,t}, \qquad (\omega_t = L). \tag{23}$$

The corollary shows that, under full information, risk aversion responds only to realized consumption shocks. It is not affected by systematic belief errors about crisis severity because

the perceived crisis path coincides with the true path.

C.b. The stochastic discount factor

Imposing market clearing, $C_t = Y_t$, and using the marginal utility of the representative agent as the pricing kernel gives

$$M_t = e^{-\rho t} \frac{1}{Y_t - H_t} = e^{-\rho t} \frac{\mathcal{R}_t}{Y_t}. \tag{24}$$

The term $e^{-\rho t}/Y_t$ is the standard log-utility component of the stochastic discount factor. The term $\mathcal{R}_t$ amplifies marginal utility when consumption is close to habit. Hence, bad states are discounted more heavily when they are also states in which risk aversion is high.

The learning channel affects the stochastic discount factor through $\mathcal{R}_t$. When realized output reveals that the crisis is more severe than expected, risk aversion rises through equation (22). This raises marginal utility in crisis states and increases the compensation investors require for bearing aggregate risk.

Because consumption does not jump when the economy enters or exits a crisis, the random times s and τ are not priced under log utility.[5]

Proposition 2 (Stochastic discount factor under learning): *In the learning economy, the representative agent prices assets using the perceived conditional law of motion for consumption. The equilibrium stochastic discount factor satisfies*

$$\frac{dM_t}{M_t} = -\widetilde{r}_t \, dt - \widetilde{\theta}_{\hat{Y},t} \, d\widetilde{Z}_{\hat{Y},t} - \widetilde{\theta}_{\eta,t} \, d\widetilde{Z}_{\eta,t}, \tag{25}$$

where tildes denote objects under the representative agent's beliefs.

[5]While the output level does not jump at these times, the conditional output distribution changes. Such events can be priced under alternative preferences, such as Epstein-Zin-Weil preferences.

The perceived risk-free rate is

$$\widetilde{r}_t = r^{log} + r_t^{habit} + \begin{cases} \widetilde{r}_t^{crisis} & \text{if} \quad \omega_t = L, \\ 0 & \text{if} \quad \omega_t = H, \end{cases} \tag{26}$$

where

$$r^{log} = \rho + \mu_{\hat{Y}} - \sigma_{\hat{Y}}^2 \tag{27}$$

is the constant log-utility component and

$$r_t^{habit} = \kappa_{\mathcal{R}} \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}\right) - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right) \sigma_{\hat{Y}}^2 \tag{28}$$

is the habit-driven component. The crisis component under learning is

$$\widetilde{r}_t^{crisis} = \kappa_\eta \left(\hat{x}_t/\eta_t - 1\right) - \sigma_\eta^2 \left(\lambda - \eta_t\right)^2 - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right) \sigma_\eta^2 \left(\lambda - \eta_t\right)^2. \tag{29}$$

The perceived market price of normal output risk is

$$\widetilde{\theta}_{\hat{Y},t} = \sigma_{\hat{Y}} \left(1 + \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right)\right). \tag{30}$$

The perceived market price of crisis-impact risk is

$$\widetilde{\theta}_{\eta,t} = \begin{cases} \sigma_\eta \left(\lambda - \eta_t\right) \left(1 + \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right)\right) & \text{if} \quad \omega_t = L, \\ 0 & \text{if} \quad \omega_t = H. \end{cases} \tag{31}$$

Proposition 2 shows how the new crisis-risk source is priced in the learning economy. In normal times, only the standard output shock is priced. During crises, the additional shock $Z_{\eta,t}$ becomes priced because it affects aggregate consumption. The market price of this risk is largest when the crisis-impact process is far from its upper boundary and when risk aversion is high.

The risk-free rate also acquires a crisis-specific component. Under learning, this component depends on the perceived crisis driver $\hat{x}_t$. Thus, investors price assets using their current assessment of crisis severity. When realized outcomes reveal that the crisis is more severe than expected, the posterior estimate $\hat{\varepsilon}_t$ rises, the perceived driver $\hat{x}_t$ falls, and risk aversion increases through equation (22). These forces raise risk compensation and strengthen the asset-pricing response.

Corollary 2 (Full-information stochastic discount factor): *If posterior uncertainty about crisis severity is zero, $\nu_t = 0$, and the prior is centered on the true value, then $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$. In this case, Proposition 2 collapses to the full-information economy. In particular, the crisis component of the risk-free rate becomes*

$$r_t^{crisis} = \kappa_\eta \left(x_t/\eta_t - 1\right) - \sigma_\eta^2 \left(\lambda - \eta_t\right)^2 - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right) \sigma_\eta^2 \left(\lambda - \eta_t\right)^2. \tag{32}$$

Crisis-impact risk becomes priced only in the crisis state and generates a hump-shaped market price of risk. The left-hand panels show how learning modifies this mechanism. When persistent bad outcomes reveal that the crisis is more severe than initially assessed, risk aversion rises and the market prices of risk become more sensitive to the evolving crisis path. Thus, learning does not replace the direct pricing of new risk; it amplifies it along severe crisis paths.

C.c. The stock market

The market portfolio is a claim to the aggregate output stream. Its price is

$$S_t = \widetilde{\mathbb{E}}_t \left[\int_t^\infty \frac{M_u}{M_t} Y_u \, du \right] = \phi_t Y_t, \tag{33}$$

where ϕ_t is the price-dividend ratio. Substituting the stochastic discount factor from equation (24) into equation (33) gives

$$\phi_t = \int_t^\infty e^{-\rho(u-t)} \widetilde{\mathbb{E}}_t \left[\frac{\mathcal{R}_u}{\mathcal{R}_t} \right] du. \tag{34}$$

Conditional on the current level of risk aversion, the representative agent expects $\mathcal{R}_t$ to mean-revert toward its long-run mean $\bar{\mathcal{R}}$. Hence,

$$\widetilde{\mathbb{E}}_t \left[\frac{\mathcal{R}_u}{\mathcal{R}_t} \right] = \frac{\bar{\mathcal{R}}}{\mathcal{R}_t} + \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t} \right) e^{-\kappa_{\mathcal{R}}(u-t)}. \tag{35}$$

Evaluating the integral gives the price-dividend ratio. Applying Itô's lemma to $S_t = Y_t \phi_t$ yields the return dynamics of the market portfolio.

This closed-form expression uses the perceived law of motion for risk aversion implied by equation (18), under which consumption innovations are mean zero conditional on the agent's information set. Learning affects the realized path of $\mathcal{R}_t$ through belief errors, but conditional on the current state the pricing formula preserves the tractability of the Menzly–Santos–Veronesi structure.

Proposition 3 (Market portfolio):

$$\phi_t = \frac{1}{\rho} \left(\frac{\bar{\mathcal{R}}}{\mathcal{R}_t} + \frac{\rho}{\rho + \kappa_{\mathcal{R}}} \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t} \right) \right) = \phi_0 + \phi_{\mathcal{R}} \frac{1}{\mathcal{R}_t}, \tag{36}$$

where

$$\phi_0 = \frac{1}{\rho + \kappa_{\mathcal{R}}}, \qquad \phi_{\mathcal{R}} = \bar{\mathcal{R}} \phi_0 \frac{\kappa_{\mathcal{R}}}{\rho}.$$

The instantaneous return of the market portfolio including dividends satisfies

$$dR_t \equiv \frac{dS_t + Y_t \, dt}{S_t} = \widetilde{\mu}_{R,t} \, dt + \sigma_{R,\hat{Y},t} \, d\widetilde{Z}_{\hat{Y},t} + \sigma_{R,\eta,t} \, d\widetilde{Z}_{\eta,t}.$$

The exposure to normal output risk is

$$\sigma_{R,\hat{Y},t} = \sigma_{\hat{Y}} V_{R,t}, \tag{37}$$

where

$$V_{R,t} = 1 + \left(\frac{\kappa_{\mathcal{R}} \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}}{\rho + \kappa_{\mathcal{R}} \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}} \right) \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right). \tag{38}$$

The exposure to crisis-impact risk is

$$\sigma_{R,\eta,t} = \begin{cases} V_{R,t} \sigma_\eta \left(\lambda - \eta_t \right) & \text{if} \quad \omega_t = L, \\ 0 & \text{if} \quad \omega_t = H. \end{cases} \tag{39}$$

The perceived instantaneous expected return and volatility are

$$\widetilde{\mu}_{R,t} = \widetilde{r}_t + \widetilde{\theta}_{\hat{Y},t} \sigma_{\hat{Y}} V_{R,t} + \widetilde{\theta}_{\eta,t} \sigma_\eta \left(\lambda - \eta_t \right) V_{R,t}, \tag{40}$$

$$\sigma_{R,t} = V_{R,t} \sqrt{\sigma_{\hat{Y}}^2 + \mathbb{1}\{\omega_t = L\} \sigma_\eta^2 (\lambda - \eta_t)^2}. \tag{41}$$

Proposition 3 shows how the new crisis-risk source affects the market portfolio in the learning economy. In a crisis, the market return is exposed to both normal output risk and crisis-impact risk. The term $V_{R,t}$ captures the effect of time-varying risk aversion on the market's exposure to these shocks. When risk aversion is high, the stock market becomes more sensitive to consumption shocks.

The risk premium has two components. The first compensates investors for exposure to normal output risk. The second compensates investors for exposure to crisis-impact risk. The crisis-risk component is active only when $\omega_t = L$. It rises when the crisis-impact process is more volatile and when the market price of crisis risk is high. Therefore, the risk premium naturally follows the shape of the crisis: it rises as the downturn deepens and falls as the economy recovers.

The functional form of the price-dividend ratio is the same as in the full-information

special case because, conditional on current risk aversion, the representative agent expects $\mathcal{R}_t$ to mean-revert toward $\bar{\mathcal{R}}$. Learning matters because it changes the realized path of $\mathcal{R}_t$. When persistent bad output realizations lead investors to revise upward the estimated crisis severity, the learning drift in equation (22) raises risk aversion. Since ϕ_t is decreasing in $\mathcal{R}_t$, these belief revisions generate valuation declines during severe crises.

Thus, learning does not introduce an additional closed-form valuation state in ϕ_t beyond $\mathcal{R}_t$. Its effect on valuations operates through the realized evolution of risk aversion. Negative belief revisions raise $\mathcal{R}_t$, and because ϕ_t is decreasing in $\mathcal{R}_t$, they lower the price-dividend ratio.

Corollary 3 (Full-information market portfolio): *If posterior uncertainty about crisis severity is zero, $\nu_t = 0$, and the prior is centered on the true value, then $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$. In this case, all perceived pricing objects in Proposition 3 coincide with their full-information counterparts:*

$$\widetilde{r}_t = r_t, \qquad \widetilde{\theta}_{\hat{Y},t} = \theta_{\hat{Y},t}, \qquad \widetilde{\theta}_{\eta,t} = \theta_{\eta,t}, \qquad \widetilde{\mu}_{R,t} = \mu_{R,t}.$$

Figure 9 places the learning economy on the left and the full-information special case on the right. The full-information panels isolate the direct effect of the new priced crisis-risk source. The crisis-impact component rises during the downturn, peaks when the crisis impact is strongest, and declines during the recovery. This produces hump-shaped risk premia and return volatility.

The learning panels show how belief revisions modify these dynamics. Along persistent crisis paths, early observations do not immediately reveal the full severity of the downturn. Investors gradually learn that the crisis is worse and more persistent than initially assessed. This learning raises risk aversion through equation (22), increases risk compensation, and lowers the price-dividend ratio. Thus, the learning economy preserves the direct new-risk mechanism from the full-information case but adds an endogenous valuation response.

The risk premium can have an S-shaped response in the learning economy. Early in the

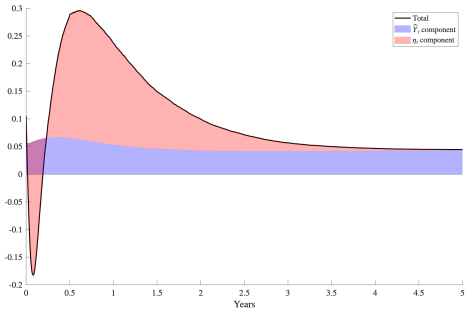

(a) Learning: risk premium decomposition

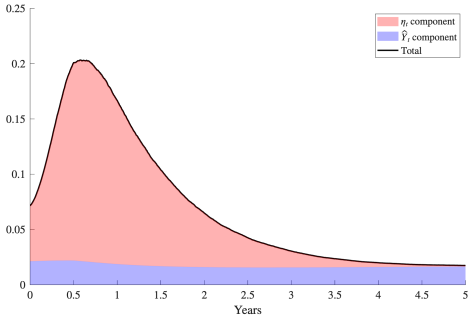

(b) Full information: risk premium decomposition

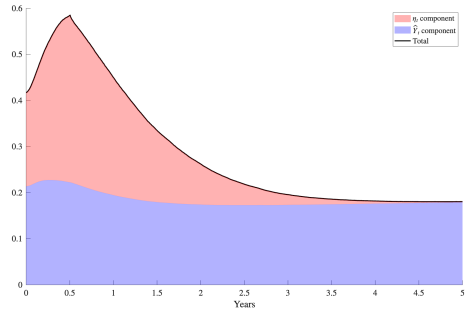

(c) Learning: return volatility decomposition

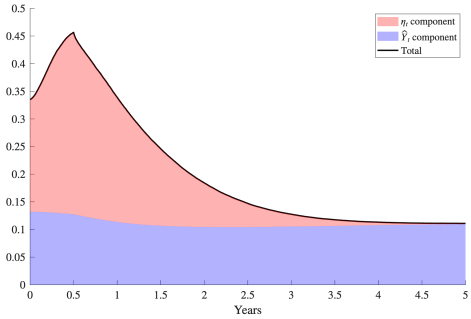

(d) Full information: return volatility decomposition

Figure 9. Equilibrium asset prices under learning and full information. The left-hand panels show the learning economy, and the right-hand panels show the full-information special case. Panels (a) and (b) decompose the equilibrium risk premium into the components driven by normal output risk and crisis-impact risk. Panels (c) and (d) decompose return volatility. The prior mean is centered at the true crisis severity, $\hat{\varepsilon}_0 = \bar{\varepsilon} = 0.0166$, with prior variance $\nu_0 = \sigma_\varepsilon^2 = 0.3886^2$. In the full-information special case, $\nu_0 = 0$ and $\hat{x}_t = x_t$. Parameters used to create the figure are the following: $\mu_{\hat{Y}} = 0.0296$, $\sigma_{\hat{Y}} = 0.0183$, $\rho = 0.04$, $\kappa_{\mathcal{R}} = 0.16$, $\lambda_{\mathcal{R}} = 22$, $\bar{\mathcal{R}} = 34$, $\alpha = 22$, $\kappa_1 = 2.9865$, $\kappa_2 = 5.0960$, $\kappa_\eta = 1.5607$, $\sigma_\eta = 0.5623$, and $\lambda = 1.05$. The minimum crisis duration is $\Delta = 0.5$ years.

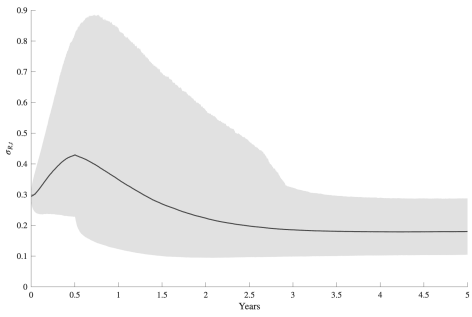

(a) Learning: return volatility

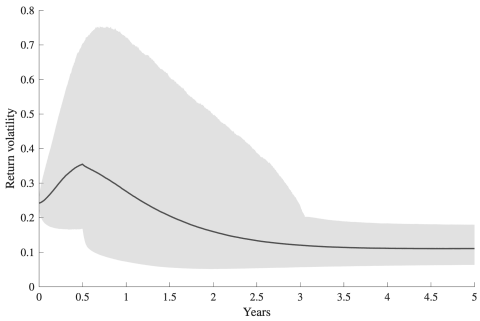

(b) Full information: return volatility

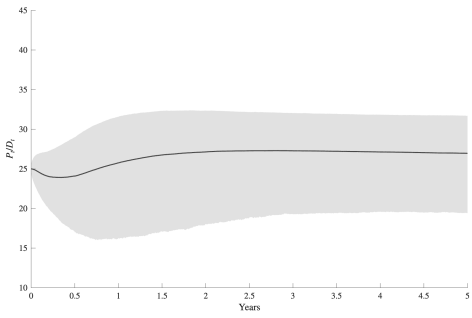

(c) Learning: price-dividend ratio

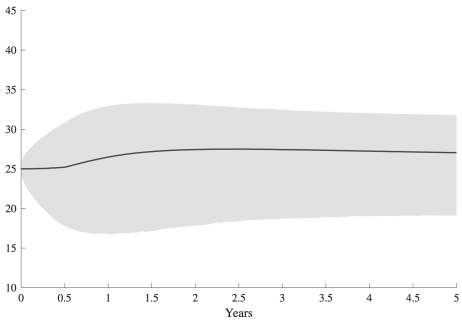

(d) Full information: price-dividend ratio

Figure 9. Equilibrium asset prices under learning and full information, continued. Panels (a) and (b) show return volatility with 5–95% confidence bands. Panels (c) and (d) show the price-dividend ratio with 5–95% confidence bands. In the learning economy, investors infer the crisis-severity parameter ε from realized output dynamics. The prior mean is centered at the true crisis severity, $\hat{\varepsilon}_0 = \bar{\varepsilon} = 0.0166$, with prior variance $\nu_0 = \sigma_\varepsilon^2 = 0.3886^2$. In the full-information special case, $\nu_0 = 0$ and $\hat{x}_t = x_t$. Parameters used to create the figure are the following: $\mu_{\hat{Y}} = 0.0296$, $\sigma_{\hat{Y}} = 0.0183$, $\rho = 0.04$, $\kappa_{\mathcal{R}} = 0.16$, $\lambda_{\mathcal{R}} = 22$, $\bar{\mathcal{R}} = 34$, $\alpha = 22$, $\kappa_1 = 2.9865$, $\kappa_2 = 5.0960$, $\kappa_\eta = 1.5607$, $\sigma_\eta = 0.5623$, and $\lambda = 1.05$. The minimum crisis duration is $\Delta = 0.5$ years.

crisis, a positive shock to η_t is good news: it raises output today and suggests that the crisis may be milder than feared. Assets that pay off when η_t improves provide insurance against states in which investors would otherwise revise beliefs downward. This hedging role can give crisis-impact risk a negative price early in the episode. As the crisis persists, the interpretation of η_t risk changes. Realizations increasingly reveal that the episode is severe and long-lasting. Beliefs about ε are revised upward, risk aversion rises, and exposure to crisis-impact risk commands positive compensation. This transition from hedging demand to risk compensation generates the steep middle part of the S-shaped risk-premium response.

Return volatility follows a more standard hump-shaped pattern. Unlike the risk premium, volatility depends on exposure magnitudes rather than compensation signs, so it does not become negative. The increase in volatility is driven primarily by the crisis-impact component, showing that the new crisis-risk source dominates asset-price uncertainty during persistent downturns.

The price-dividend ratio distinguishes the learning economy from the full-information special case. In the full-information case, the valuation response is muted because expected future risk aversion is pinned down by the current value of $\mathcal{R}_t$ and its mean reversion. In the learning economy, persistent negative belief revisions raise the realized path of risk aversion and generate a decline in valuations. As beliefs stabilize and the crisis begins to unwind, risk compensation falls and valuations recover.

D. The Role of Prior Uncertainty

The strength of the learning channel depends on how uncertain investors are at the beginning of the crisis. To isolate this effect, we keep the prior mean fixed at the true crisis severity and vary only the prior variance, ν_0. This exercise changes the confidence investors place in their initial assessment without changing initial beliefs.

A lower prior variance means that investors are more confident in their initial assessment of crisis severity. As a result, they place less weight on new output realizations and update

their beliefs more slowly. If the crisis turns out to be persistent, this slow updating delays the recognition that the episode is severe, producing a longer sequence of negative belief revisions. By contrast, a higher prior variance means that investors are less confident in their initial assessment. They assign more weight to new information, revise beliefs more quickly, and incorporate adverse realizations sooner. Thus, although higher prior uncertainty raises the instantaneous Kalman gain, it can make asset prices adjust more smoothly because investors respond more quickly to incoming information.

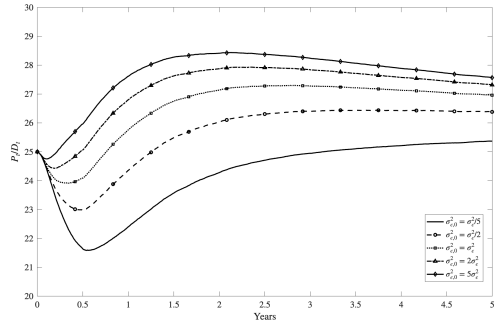

(e) Price-dividend ratio

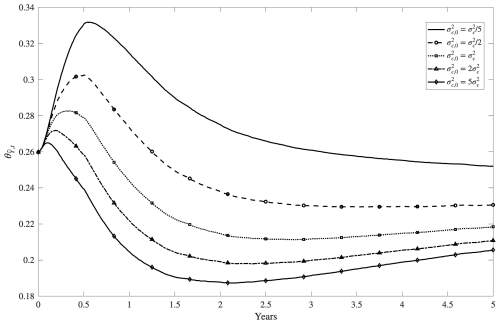

(f) Market price of output risk

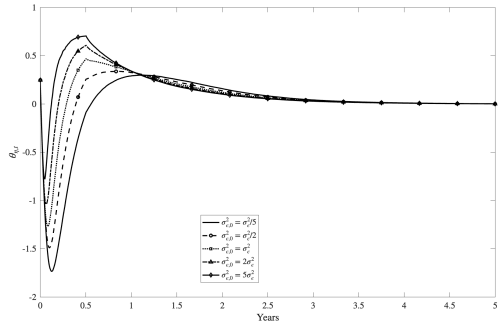

(g) Market price of crisis risk

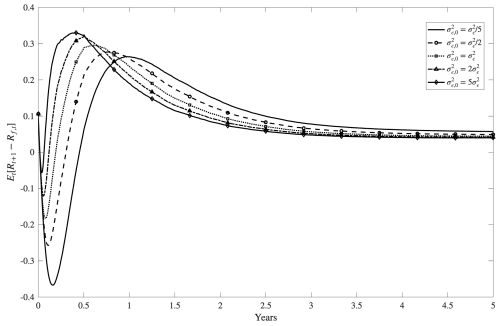

(h) Risk premium

Figure 10. Prior variance and learning dynamics. The figure shows how prior uncertainty about crisis severity affects asset prices along recession paths that remain in the crisis state for at least two quarters. The prior mean is fixed at the true crisis severity, $\hat{\varepsilon}_0 = \varepsilon$, while the prior variance ν_0 varies across scenarios. Lower prior variance means that investors are more confident in their initial assessment and update more slowly, generating more persistent belief corrections. Higher prior variance leads to faster updating and smoother price dynamics. Panel (a) shows the price-dividend ratio, Panels (b) and (c) show the market prices of normal output risk and crisis-impact risk, and Panel (d) shows the equilibrium risk premium.

Figure 10 shows these comparative statics. Panel (a) shows that lower prior variance

generates a larger and more persistent decline in the price-dividend ratio. The reason is that confident investors are slower to abandon their initial assessment. When the crisis keeps deteriorating, belief revisions keep risk aversion elevated.

Panels (b) and (c) show that prior uncertainty affects both the market price of normal output risk and the market price of crisis-impact risk. The effect is especially strong for crisis-impact risk because this risk directly reflects the evolving assessment of whether the current episode is temporary or persistent. Panel (d) shows the resulting effect on the equilibrium risk premium. Lower prior variance produces a stronger S-shaped response: the risk premium initially falls, then rises sharply as investors learn that the crisis is severe, and finally declines as uncertainty is resolved and the economy recovers.

The prior-variance exercise highlights an important implication of the model. Severe crises are most destabilizing when investors are initially confident that the disruption will be limited. In that case, adverse realizations force a sequence of belief revisions, which raises risk aversion, increases risk compensation, and depresses valuation ratios.

The theoretical contribution can be summarized as follows. The model starts from a crisis-shape process that matches the gradual decline and rebound of output during downturns. This process introduces a new source of aggregate risk. When the economy enters the crisis state, investors require compensation for bearing this risk. Because the importance of the crisis-risk source rises and falls with the crisis impact, the model generates hump-shaped risk premia and return volatility.

Learning about crisis severity is the main information structure of the model. Investors do not know the severity parameter ε when the crisis begins. They infer it from the realized output path. If the crisis proves more severe than initially expected, realized growth is persistently worse than expected. These negative belief revisions raise risk aversion, amplify discount-rate movements, and lower valuation ratios. The full-information model is obtained as the special case in which posterior uncertainty about ε is zero. Thus, the model nests the no-learning benchmark while emphasizing that learning about the severity and persistence

of crises is central for understanding valuation declines during severe recessions.

V. Data

In this section, we describe the datasets used to analyze aggregate output, consumption, and asset prices. Our analysis spans multiple dimensions, including long-term historical trends, crisis versus non-crisis periods, and sub-samples before and after the 1990s.

A. Aggregate output and consumption

We use quarterly real GDP per capita and monthly real industrial production from FRED, together with U.S. personal consumption expenditures from the Bureau of Economic Analysis. Over the full sample (January 1947 to June 2024), we compute year-on-year (Y-o-Y) and quarter-on-quarter (Q-on-Q) log growth rates for real GDP per capita and real personal consumption expenditures per capita. We report results for the full sample and for key subsamples: pre-1990, post-1990, crisis periods, and normal times. Table II summarizes these statistics.

We date crises using an output-recovery definition aligned with the model's notion of a crisis "impact" that lasts until the economy returns to its pre-crisis path. Specifically, a crisis begins in the NBER peak quarter (the onset of the subsequent contraction). It ends in the first quarter in which the level of output (real GDP per capita) has recovered to at least its pre-crisis peak level. Thus, the crisis window spans the downturn and the recovery phase, and "normal times" are all quarters outside these NBER-peak-to-recovery intervals. Under our crisis definition, 30.3% of all quarters in our full sample (January 1947 to June 2024) are classified as crisis periods.

Real GDP grew at an average annualized rate of 1.973% (Y-o-Y) and 1.954% (Q-on-Q) over the entire sample period, while real consumption increased by 2.09%. In crisis periods, output declines by -0.254% when measures Y-o-Y or increases only modestly by 0.466%,

	GDP growth		Consumption growth	
	Y-on-Y	Q-on-Q	Y-on-Y	Q-on-Q
Entire sample: 1947 - 2024	1.973%	1.954%	2.090%	2.089%
	(2.539%)	(2.234%)	(2.206%)	(2.171%)
Pre-1990s	2.304%	2.257%	2.325%	2.313%
	(2.868%)	(2.204%)	(2.197%)	(1.911%)
Post 1990s	1.570%	1.578%	1.804%	1.812%
	(2.008%)	(2.259%)	(2.200%)	(2.447%)
During crises	-0.254%	0.321%	0.469%	0.982%
	(2.450%)	(3.302%)	(2.176%)	(3.149%)
During normal times	2.945%	2.964%	2.728%	2.553%
	(1.850%)	(1.422%)	(1.539%)	(1.504%)
Crisis-to-non-crisis volatility ratio	1.3241	2.3216	1.4144	2.0931

Table II Quarterly GDP and Consumption Growth Data. This table shows annualized growth rates and standard deviations for quarterly real gross domestic product (GDP) and quarterly consumption per capita. Y-on-Y indicates year-on-year, Q-on-Q means quarter-on-quarter, both annualized, using data from January 1947 to June 2024. Real gross domestic product per capita [A939RX0Q048SBEA] and real personal consumption expenditures per capita [A794RX0Q048SBEA] were retrieved from FRED, Federal Reserve Bank of St. Louis, on October 3, 2024. Crises span all NBER-peak-to-recovery periods.

Q-on-Q, and consumption growth slows to 0.471% (Y-o-Y) and 1.067% (Q-on-Q). During normal times, output grows at a robust rate of 2.964% (Y-o-Y), with similar consumption growth.

Volatility rises significantly during crises. The volatility ratios, calculated as the ratio of standard deviations during crises to those during normal times, range from 1.32 (Y-o-Y) to 2.32 (Q-on-Q) for GDP growth and from 1.41 (Y-on-Y) to 2.09 (Q-on-Q) for consumption growth.

B. Cross-crisis variation

Historical evidence also reveals substantial heterogeneity across crisis episodes. Using the NBER chronology, we identify twelve U.S. recessions between 1947 and 2024. Unlike the NBER dating, we define the end of a crisis as the quarter when real GDP per capita first returns to its pre-crisis level. Table III reports the NBER peak (crisis onset), the trough

NBER recession	NBER Peak	Trough	Recovery	Peak→Trough (quarters)	Peak→Recovery (quarters)
Recession of 1949	1948Q4	1949Q4	1950Q1	5	6
Recession of 1953	1953Q2	1954Q2	1955Q1	5	8
Recession of 1958	1957Q3	1958Q1	1959Q1	3	7
Recession of 1960–61	1960Q2	1960Q4	1961Q2	3	5
Recession of 1969–70	1969Q4	1970Q4	1971Q1	5	6
Recession of 1973–75	1973Q4	1975Q1	1976Q1	6	10
Recession of 1980	1980Q1	1980Q3	1981Q1	3	5
Recession of 1981–82	1981Q3	1982Q4	1983Q3	6	9
Early 1990s Recession	1990Q3	1991Q1	1992Q2	3	8
Early 2000s Recession	2000Q1	2000Q1	2000Q2	1	2
Global Financial Crisis	2007Q4	2009Q2	2013Q1	7	22
COVID-19 Recession	2019Q4	2020Q2	2021Q1	3	6
				mean: 4.17	7.83
				std: (1.75)	(4.93)

Table III Crisis dating based on real GDP per capita. NBER Peak is the peak quarter declared by the NBER. Trough is the quarter within the crisis window (from Start through Recovery) in which real GDP per capita attains its minimum. Recovery is the quarter when real GDP per capita first reaches or exceeds its pre-crisis level. The duration columns are counted inclusively, so Peak→Trough equals 1 if the trough occurs in the start quarter.

quarter defined as the minimum of real GDP per capita within the crisis window, and the recovery quarter when real GDP per capita first returns to its pre-crisis level. On average, crises reach their trough after 4.17 quarters (standard deviation 1.75) and recover after 7.83 quarters (standard deviation 4.93).

The dispersion in these durations is economically meaningful. For example, the COVID-19 recession features one of the shortest downturn–recovery windows in the sample but is associated with an unusually sharp contraction in output, whereas the Global Financial Crisis involves a prolonged recovery that extends well beyond the official NBER recession end date. Matching both the average duration of crises and the cross-crisis variation in their dynamics is therefore essential for capturing the evolution of expected growth, macroeconomic uncertainty, and risk premia during downturns.

Here is a revised version that is consistent with the current empirical implementation and emphasizes the Goyal-Welch-Zafirov predictor data.

C. Asset prices

We measure aggregate stock-market returns, valuation ratios, volatility, and predictive variables primarily using the monthly return-predictor data set of Goyal et al. (2024). This data set provides a postwar monthly panel of market returns, risk-free rates, dividend-price ratios, stock-market variance, and standard predictors used in the return-predictability literature. Our event-study return panels use future realized total excess market returns constructed from these monthly data. Specifically, for horizons $h \in \{3, 6, 9, 12\}$, we use future h-month total market returns in excess of the corresponding risk-free return.

We use the same data source to measure stock-market volatility and valuations. Return volatility is measured as the square root of the stock-market variance variable, `svar`, from Goyal et al. (2024). The price–dividend ratio is measured as the inverse of the dividend-price ratio, $1/dp$. In the event-study figures, both volatility and the price–dividend ratio are normalized to one at the recession peak so that the plots emphasize their dynamics around business-cycle turning points.

In some robustness exercises, we also use external asset-pricing series to provide additional empirical benchmarks. These include the risk-premium measure of Marfè and Pénasse (2024) and, where relevant, the lower bound on expected market returns developed by Martin (2017). These series are not used to calibrate the model. They serve only as complementary empirical measures of aggregate risk compensation.

VI. Calibrating crisis dynamics

This section estimates the crisis-dynamics parameters that govern the joint evolution of the crisis driver x_t and the crisis impact η_t, and hence the output gap dynamics around

crises. We deliberately restrict the estimation to macro data and do not target asset-pricing moments, so the model's asset-pricing implications remain an out-of-sample test.

A. Matching crisis responses in data and model

To connect the model to data, we connect the crisis impact variable η_t with the empirical output gap, defined as $(Y_t - Y_t^*)/Y_t^*$, where Y_t is real GDP and Y_t^* is potential output. In the model, realized output decomposes multiplicatively as $Y_t = \widehat{Y}_t \eta_t$, where $\widehat{Y}_t$ captures the smooth, capacity-like component of production and η_t captures cyclical deviations from that component. This structure implies

$$\frac{Y_t - \widehat{Y}_t}{\widehat{Y}_t} = \eta_t - 1,$$

so $\eta_t - 1$ is the model analogue of the output gap. We therefore interpret $\widehat{Y}_t$ as a proxy for potential output Y_t^* and map the output gap $(Y_t - Y_t^*)/Y_t^*$ to $\eta_t - 1$ when estimating the dynamics of η and the model's response to recessions.

Let y_t denote the quarterly output gap (measured as log real output relative to potential). To characterize the typical dynamic response of the output gap to the onset of recessions, we estimate local projections at horizons $h = 0, \ldots, H$:

$$y_{t+h} = \alpha_h + \beta_h \, \mathbb{1}\{\text{Crisis}_t\} + \Gamma_h' X_{t-1} + \varepsilon_{t+h}, \tag{42}$$

where $\mathbb{1}\{\text{Crisis}_t\}$ equals one at the beginning of an NBER recession (business-cycle peak), and X_{t-1} contains lags of y_t and standard macro-finance controls. Each horizon is estimated by OLS, and we compute HAC standard errors (Bartlett kernel). The coefficient $\{\beta_h\}_{h=0}^{H}$ is directly interpretable as the impulse response of the output gap to a crisis shock.

We summarize output dynamics during crises using seven moments that capture the average crisis shape, cross-crisis heterogeneity, unconditional crisis frequency, and the relative volatility of output growth in crises versus normal times. Moments M1–M3 are based on

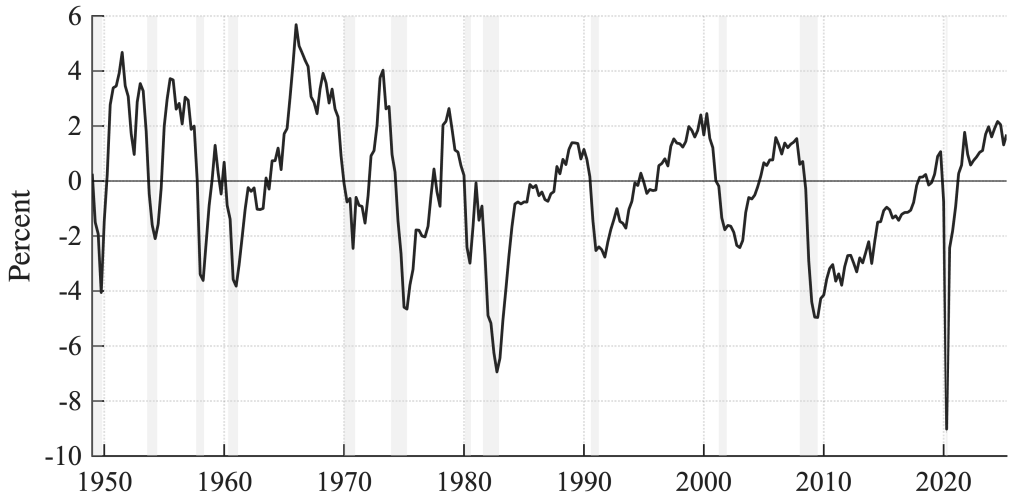

(a) Output gap over time

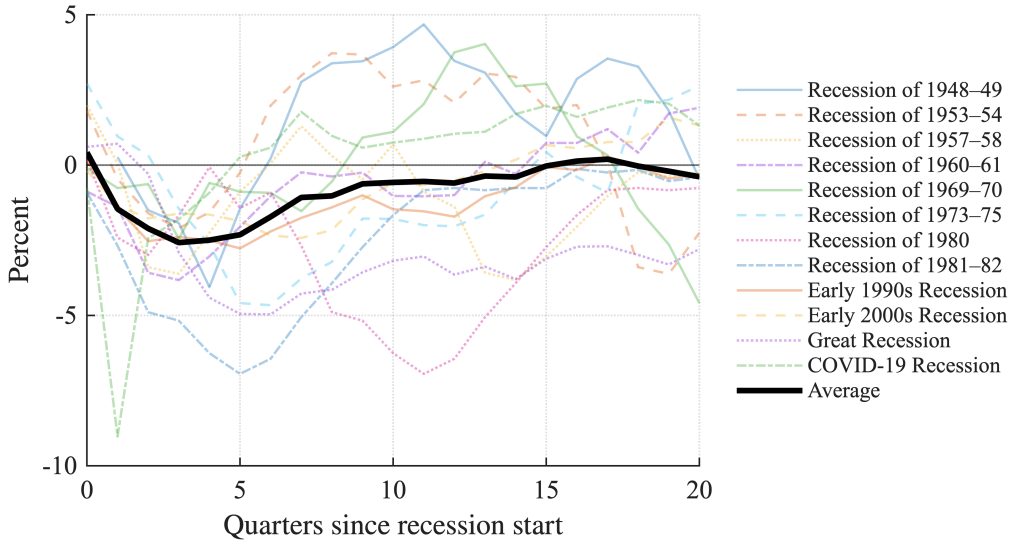

(b) Event study of the output gap following the start of an NBER recession

Figure 11. The output gap over the business cycle and its dynamics around recessions. Panel (a) plots the quarterly output gap, computed as $100 \times (\text{GDPC1}/\text{GDPPOT} - 1)$, where GDPC1 is real GDP and GDPPOT is potential output from FRED. Shaded regions mark NBER recession periods, constructed using the NBER peak–trough chronology, with the recession start defined as the month following the NBER peak and shading extending through the trough month. The horizontal line at zero indicates output at potential. Panel (b) aligns the same output-gap series by recession start and plots the evolution from quarter 0 through quarter +20, where quarter 0 is the first recession quarter (the quarter containing the month after the NBER peak). Each thin colored line corresponds to one NBER recession in the postwar sample, and the thick black line shows the cross-recession average path.

the empirical impulse response $\{\hat{\beta}_h\}_{h=0}^{H}$ of output following a crisis onset and summarize the average depth, timing, and recovery of the output decline. Moments M4 and M5 capture cross-crisis heterogeneity: M4 is the average time from the NBER peak to the output trough, while M5 is the cross-crisis standard deviation of crisis depth, measured as the peak-to-trough decline in real output. Moment M6 captures the unconditional frequency of crisis periods. Moment M7 compares year-on-year output growth volatility during crisis periods with year-on-year output growth volatility in non-crisis periods.

(M1) **Depth (size of the dip).** The trough magnitude of the output-gap response (in percent),

$$m_1^{\text{data}} \equiv \min_{0 \le h \le H} \hat{\beta}_h. \tag{43}$$

(M2) **Time to trough.** The horizon (in quarters) at which the trough occurs,

$$m_2^{\text{data}} \equiv \arg\min_{0 \le h \le H} \hat{\beta}_h. \tag{44}$$

(M3) **Time to recovery.** The first horizon after the trough at which the response recovers to its pre-crisis level,

$$m_3^{\text{data}} \equiv \min\left\{h > m_2^{\text{data}} : \hat{\beta}_h \ge 0\right\}. \tag{45}$$

(M4) **Variation in time to recovery.** The cross-crisis standard deviation of the number of quarters from crisis onset (NBER peak) to recovery back to the pre-crisis output level,

$$m_4^{\text{data}} \equiv \text{Std}_{c \in \mathcal{C}}(T_c^{\text{rec}}), \qquad T_c^{\text{rec}} \equiv t_c^{\text{rec}} - t_c^{\text{peak}}, \tag{46}$$

where t_c^{rec} is the first quarter after the trough in which real GDP per capita reaches (or exceeds) its pre-crisis peak level.

(M5) **Variation in crisis severity.** The cross-crisis standard deviation of crisis depth, where crisis depth is measured by the size of the output decline from the NBER peak

to the crisis trough:

$$m_5^{\text{data}} \equiv \text{Std}_{c\in\mathcal{C}}(D_c)\,, \qquad D_c \equiv 1 - \min_{t\in\mathcal{T}_c}\left(\frac{Y_{c,t}}{Y_{c,t^{\text{peak}}}}\right). \tag{47}$$

Here, t_c^{peak} is the NBER peak quarter for crisis $c \in \mathcal{C}$, $Y_{c,t}$ denotes real GDP per capita normalized to one at the NBER peak, and $\mathcal{T}_c$ is the crisis window over which the trough is measured. Thus, D_c is the maximum peak-to-trough output decline during crisis c, and m_5^{data} captures how much crisis depth varies across recession episodes.

(M6) **Crisis frequency.** The share of periods classified as crises,

$$m_6^{\text{data}} \equiv \frac{N^{\text{crisis}}}{N^{\text{total}}}, \tag{48}$$

where N^{crisis} is the number of quarters that fall within crisis periods and N^{total} is the total number of quarters in the sample.

(M7) **Volatility ratio.** The ratio of the volatility of year-on-year output growth during crisis periods to that during non-crisis periods,

$$m_7^{\text{data}} \equiv \frac{\text{Std}(g_t^{yoy} \mid t \in \mathcal{T}^{\text{crisis}})}{\text{Std}(g_t^{yoy} \mid t \in \mathcal{T}^{\text{normal}})}, \qquad g_t^{yoy} \equiv \frac{Y_t}{Y_{t-4}} - 1, \tag{49}$$

where $\mathcal{T}^{\text{crisis}}$ denotes quarters that fall within crisis periods and $\mathcal{T}^{\text{normal}}$ denotes quarters outside crisis periods.

These seven moments jointly discipline the depth, timing, recovery, cross-crisis variation, overall frequency of crisis downturns, and the amplification of output-growth volatility during crises in the data.

We match the targeted moments closely overall. The model matches depth (0.9682 *vs.* 0.9681, in data vs. in the model), and both the time to trough (M2) and the time to recovery (M3) are closely aligned (3 *vs.* 5 quarters and 16 *vs.* 17 quarters, respectively). The model also reproduces the cross-crisis dispersion in the time to trough (M4: 3.8834 *vs.* 4.9329), indicating

Table IV Parameters describing the Aggregate Economy and Preferences. This table presents the parameters governing aggregate output dynamics during normal times (Panel A) and agent preferences (Panel B). The parameters in Panel A are based on real GDP growth observed during all normal (non-crisis) periods between January 1947 to June 2024 (full sample), as detailed in Table II. Preference parameters in Panel B are based on Menzly, Santos, and Veronesi (2004), except for α, which is set to a lower value of 25.125 to account for the newly introduced crisis component that increases output volatility during crises.

Panel A: Real GDP growth during normal times (1947-2024)	
$\mu_{\hat{Y}}$	2.945%
$\sigma_{\hat{Y}}$	1.850%

Panel B: Parameters from Menzly, Santos, and Veronesi (2004)	
ρ	0.04
$\lambda_{\mathcal{R}}$	20
$\kappa_{\mathcal{R}}$	0.16
$\bar{\mathcal{R}}$	34
α	25.125

Table V Target moments: data vs. model. The table reports the seven moments used to discipline crisis dynamics. Moments M1–M3 summarize the local-projection impulse response of the output gap to crisis onset; M4 summarizes cross-crisis dispersion in time to recovery; M5 measures cross-crisis dispersion in severity; M6 is crisis frequency; and M7 is the ratio of year-on-year output growth volatility in crises relative to normal times.

Moment	Data	Model
(M1) Depth (trough magnitude)	0.9681	0.9682
(M2) Time to trough (quarters)	5	3
(M3) Time to recovery (quarters)	17	16
(M4) Std. time to trough (quarters)	4.9329	3.8834
(M5) Std. crisis depth	0.0227	0.0226
(M6) Crisis frequency (share of quarters)	0.3030	0.3044
(M7) Output volatility ratio (crisis/normal)	1.3241	1.8893

that it captures heterogeneity in downturn timing reasonably well. The main discrepancy arises in the last moment matched: the model implies higher volatility in crises relative to normal times (M7: 1.8893 *vs.* 1.3241).

Table VI reports the estimated parameters. The combination of the estimated κ parameters, particularly the distance between $\kappa_2 - \kappa_1$, the level of κ_η and the average size of the ε shock generates a severity and duration of crises that are quantitatively consistent with observed data, see Figure 12.

Table VI Estimated Parameters. This table presents the estimated parameters used to match the model with output-gap impulse responses. These parameters from the crisis driver x_t and crisis impact η_t processes determine how output evolves during crises.

Parameter θ_i	Estimated Value
κ_η	1.5607
σ_η	0.5623
κ_1	2.9865
κ_2	5.0960
$\bar{\varepsilon}$	0.0166
σ_ε	0.3886
ν	0.5313

Figure 12 shows that the model-implied local-projection response closely matches the empirical response. The simulated impulse response reproduces the sharp initial decline in the output gap, the depth and timing of the trough (around year 1), and the subsequent recovery path back toward zero over the medium run. At longer horizons, the empirical estimate exhibits some additional oscillations and a modest positive overshoot, whereas the simulated response is slightly smoother and reverts more steadily toward zero. Overall, the visual fit indicates that the estimated crisis process captures both the magnitude and the dynamics of the average output-gap response to crisis onset in the data reasonably well.

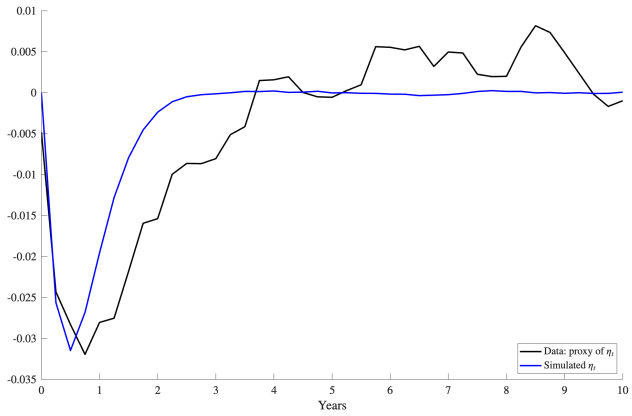

Figure 12. Matching output-gap crisis dynamics. This figure compares the empirical local-projection impulse response of the output gap to the model-implied impulse response estimated from simulated data. The black line plots the empirical response (data proxy for η_t), obtained by aligning the output gap at crisis onset and tracing its average evolution over the subsequent years. The blue line plots the corresponding simulated response for η_t, constructed using the same stacked-panel local-projection procedure on Monte Carlo simulations. The simulated response is generated using the estimated parameter vector reported in Table VI, which is chosen to match the seven target moments reported in Table V.

B. Asset-pricing implications

We next examine the model's asset-pricing implications during crises using the parameter estimates disciplined solely by output-gap dynamics and crisis-duration moments. Importantly, we do not match any parameters to asset-pricing data. The return, volatility, and valuation dynamics discussed below are therefore out-of-sample implications of the crisis-output mechanism. We compare the data with two versions of the model. The learning economy is our baseline specification, in which investors infer the crisis-severity parameter ε from realized output dynamics. The no-uncertainty benchmark sets posterior uncertainty to zero, so that investors know the crisis path immediately.

Before turning to the comparison with the data, it is useful to clarify how the event-study design in this section differs from the model-mechanism figure. Figure 9 shows the evolution of asset-pricing moments from time s, the date at which the economy enters the new-risk state. It is therefore designed to isolate the model mechanism: how risk premia, volatility, and valuation ratios evolve after the arrival of a new crisis-risk state. Figure 13, in contrast, is constructed to match the empirical event-study design. In this figure, the model is simulated

unconditionally: the economy starts in normal times and recessions arrive through Poisson jumps with the estimated intensity. We then align simulated episodes around their model-implied output peaks, which are the simulated counterparts of NBER recession peak dates in the data. We further restrict the simulated event sample to paths that have entered a recession state and have not exited that state within 0.5 years. This conditioning mirrors the idea that NBER recessions correspond to persistent downturns rather than short-lived shocks. Thus, Figure 9 is intended to show the model mechanism, while Figure 13 provides an apples-to-apples comparison between the model and the data.

Figure 13 summarizes the asset-pricing implications of the estimated crisis dynamics. The empirical averages display a clear non-monotonic pattern around recession peaks. Future excess returns are low around the output peak and then rise with a delay, remaining elevated during the early phase of the downturn before gradually reverting. This pattern is especially visible at longer horizons, where realized excess returns increase after the recession has already begun. The learning model closely reproduces this S-shaped behavior: returns fall around the peak and then rise as the downturn develops. The no-uncertainty benchmark generates the same broad direction but with a weaker response, especially at longer horizons.[6]

Around the peak, asset prices fall as the economy enters the crisis state and investors observe deteriorating output dynamics. These price declines lower realized returns measured

[6]For the return panels, we compute comparable future realized excess returns in the data and in the model. In the data, we use monthly total market returns and construct future h-month total excess returns, for $h \in \{3, 6, 9, 12\}$, by compounding the market return over the future horizon and subtracting the corresponding compounded risk-free return.

In the model, we construct the same object from simulated equilibrium prices. Let D_t denote simulated aggregate dividends, proxied by total output, and let

$$P_t = \left(\frac{P_t}{D_t}\right) D_t$$

be the model-implied equity price. The future h-month realized total excess return is computed as

$$R^e_{t,t+h} = \frac{P_{t+h} - P_t + \sum_{s=t}^{t+h-1} D_s \Delta t}{P_t} - \left(\exp\left\{\sum_{s=t}^{t+h-1} r_s \Delta t\right\} - 1\right).$$

where r_s is the model-implied instantaneous risk-free rate. Thus, the model return includes both capital gains and dividends and subtracts the risk-free return over the same horizon. This construction mirrors the empirical total excess return measure and makes the model and data panels directly comparable.

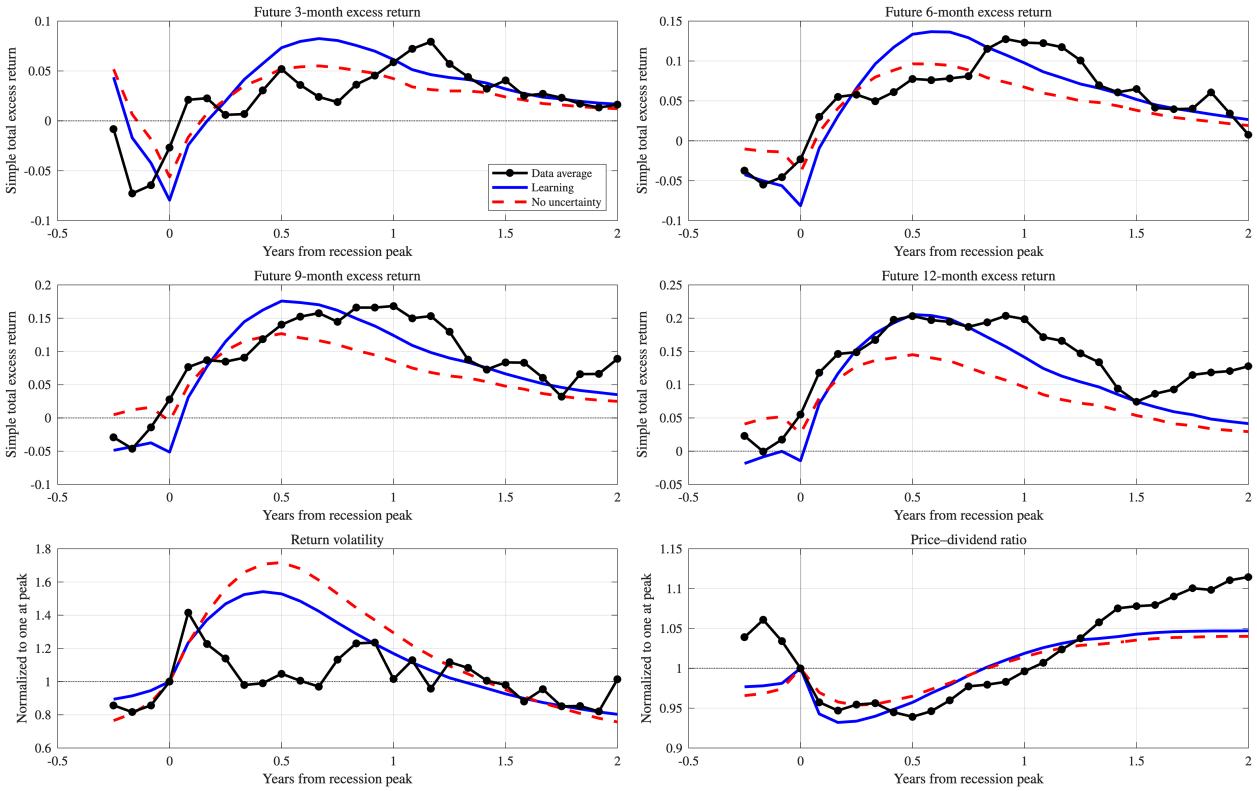

Figure 13. Asset prices around crises: data, learning model, and no-uncertainty benchmark. This figure compares empirical asset-price dynamics around NBER recession peaks with simulated dynamics from the learning economy and the no-uncertainty benchmark. The black line with markers reports the empirical average across postwar recessions. The blue solid line reports the learning economy, with $\hat{\varepsilon}_0 = \bar{\varepsilon} = 0.0166$ and prior variance $\nu_0 = \sigma_\varepsilon^2 = 0.3886^2$. The red dashed line reports the no-uncertainty benchmark, in which posterior uncertainty about crisis severity is set to zero. The model is simulated unconditionally: the economy starts in normal times, recessions arrive through Poisson jumps with the estimated intensity, and simulated episodes are aligned at the model-implied output peak. The simulated event sample is restricted to episodes that remain in the crisis state for at least 0.5 years after entering recession. Panels (a)–(d) report realized future total excess returns over 3, 6, 9, and 12 months. Panel (e) reports return volatility, normalized to one at the recession peak. Panel (f) reports the price–dividend ratio, normalized to one at the recession peak. The empirical return panels use future excess returns from the monthly predictor data. Empirical volatility is measured using the square root of the `svar` variable, and the empirical price–dividend ratio is measured as the inverse dividend-price ratio, $1/dp$. The remaining model parameters are described in the calibration section.

over short forward windows. As the crisis unfolds, however, the economy becomes more exposed to crisis-impact risk, and required compensation rises. This produces the delayed hump-shaped increase in future excess returns. The learning economy amplifies this pattern because bad output realizations lead investors to revise beliefs toward a more severe and persistent crisis.

Return volatility also rises after the recession peak and then mean reverts. This pattern appears in both the data and the model. The mechanism is not driven only by learning. Even in the no-uncertainty benchmark, the crisis state introduces a new source of aggregate risk. As the crisis-impact process moves away from its normal value, exposure to this risk increases and return volatility rises. As the economy recovers and the crisis-impact process moves back toward normal times, volatility declines. The no-uncertainty economy therefore generates a pronounced volatility response, while the learning economy produces a somewhat different timing because posterior beliefs and risk aversion adjust gradually during the downturn.

In the data, the pd ratio falls after the recession peak and recovers gradually. The learning economy generates a sharper and more persistent valuation decline than the no-uncertainty benchmark. This decline reflects belief updating: when realized output is worse than expected, investors infer that the crisis is more severe and persistent than initially believed. These negative belief revisions raise risk aversion and discount rates, lowering valuations.

The no-uncertainty benchmark also exhibits a modest decline in the price–dividend ratio, but the source of this decline is different from the learning economy. Because posterior uncertainty is set to zero, the decline cannot reflect belief revisions about crisis severity. Instead, it reflects the direct effect of realized crisis shocks on risk aversion along the persistent-crisis paths used in the event study.

The event-study sample conditions on simulated crises that remain in the crisis state for at least 0.5 years after entry. This selection matters. Paths that satisfy this condition are precisely those in which early realizations of the crisis-impact process are sufficiently

adverse that η_t stays below one rather than quickly recovering. In the model, these paths are associated with negative shocks to η_t. Even without learning, such shocks increase effective risk aversion through the external habit/risk-aversion state. Intuitively, a negative η shock means that the crisis impact is worsening: output is moving farther below its normal path, the economy is becoming more exposed to the new crisis-risk source, and marginal utility becomes more sensitive to additional bad shocks. This raises discount rates and lowers the price–dividend ratio.

Thus, the temporary valuation decline in the no-uncertainty benchmark is not a learning effect. It is the direct equilibrium pricing effect of adverse realized crisis shocks on risk aversion and discount rates, amplified by the fact that the plotted event sample selects persistent crises. The difference is that, in the no-uncertainty economy, investors already know the crisis-severity parameter, so valuations respond only to the realized movement in the crisis state and the associated risk exposure. In the learning economy, the same adverse output realizations also lead investors to revise beliefs toward a more severe and persistent crisis. These negative belief revisions further raise risk aversion and expected discount rates, making the decline in the price–dividend ratio larger and more persistent.

Overall, Figure 13 supports three conclusions. First, realized excess returns around recessions are S-shaped: they fall around the peak and rise only after the downturn intensifies. Second, return volatility is hump-shaped, increasing during the crisis and declining during the recovery. Third, valuation ratios fall around recessions, with the learning economy producing the stronger and more persistent decline. Since the model is estimated using macroeconomic moments rather than asset-pricing moments, these asset-price dynamics are not mechanically imposed by the calibration.

C. Non-Recession Risk Episodes

We study episodes in which a salient new macroeconomic, financial, geopolitical, or health-related risk emerged but did not ultimately develop into an NBER recession. These

episodes provide an important empirical counterpart to the model because they isolate periods in which agents became aware of elevated risks without experiencing a full macroeconomic collapse.

The central implication of the model is that awareness of a new risk alone is insufficient to generate recession-like valuation dynamics. In the model, large declines in valuation ratios arise primarily when investors update toward sufficiently severe and persistent crisis realizations. When risks fail to propagate into a broad macroeconomic contraction, the model predicts substantially weaker movements in risk premia and valuation ratios.

Table VII lists the major non-recession risk episodes used in the event-study analysis. The selected episodes correspond to economically salient shocks that generated substantial uncertainty, financial stress, or public concern, but that ultimately did not coincide with NBER recessions.

Conditional Dynamics Around Non-Recession Risk Events

Figure 14 compares the empirical dynamics around non-recession risk episodes with the corresponding model simulations. The empirical panels are aligned around the onset dates listed in Table VII. The model panels are constructed from simulated new-risk episodes that arrive through Poisson jumps but resolve quickly: we condition on paths in which the crisis impact variable η_t reverts back to its normal value within 0.5 years. These simulated episodes are therefore the model counterpart of salient risk events that raise uncertainty but do not develop into persistent recessions.

The comparison is informative because the model parameters are not chosen to match these non-recession episodes. Instead, the same calibrated model used to explain recession dynamics is used to generate the simulated non-recession risk events. The exercise therefore asks whether the model can also distinguish between persistent crisis realizations and short-lived risk scares.

The data and the model display similar qualitative patterns. Future excess returns do

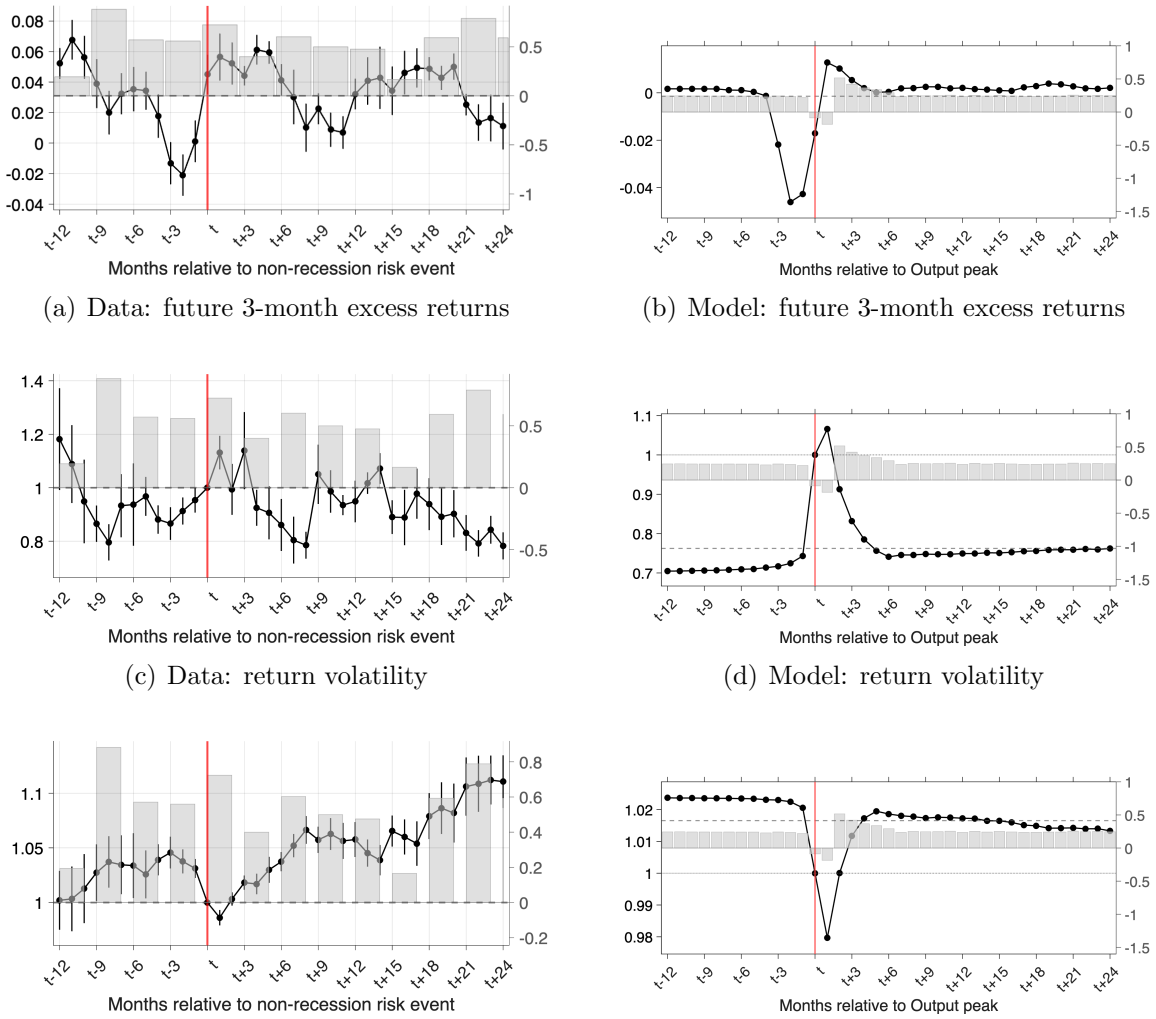

(a) Data: future 3-month excess returns

(b) Model: future 3-month excess returns

(c) Data: return volatility

(d) Model: return volatility

(e) Data: price-dividend ratio

(f) Model: price-dividend ratio

Figure 14. Asset prices around non-recession risk episodes: Data and model. The figure compares empirical asset-price dynamics around non-recession risk episodes with model simulations of short-lived risk episodes. Panels (a), (c), and (e) show the data, aligned around the onset dates in Table VII. Panels (b), (d), and (f) show the model, aligned around simulated risk-event dates. In the model, new-risk episodes arrive through Poisson jumps, but the sample is restricted to paths in which the crisis impact reverts to its normal value within 0.5 years, so the shock does not become a persistent recession. The same calibrated model parameters used in the main recession analysis are used here. Return volatility and the price–dividend ratio are normalized to one at the event month. Black dots report averages across episodes, vertical bars show HAC standard errors, and gray bars show output growth on the secondary right axis. Output growth is quarterly real GDP growth in the data and monthly simulated output growth in the model.

Table VII: Major Non-Recession Risk Episodes

Event	Start date	Description
Korean War inflation	1950:06	Inflation surge and wartime mobilization fears.
Suez Crisis	1956:10	Geopolitical conflict and oil supply concerns.
Sputnik shock	1957:10	Cold War technology and defense uncertainty.
Berlin Crisis	1961:08	Escalation of Cold War geopolitical tensions.
Cuban Missile Crisis	1962:10	Extreme geopolitical and nuclear-risk fears.
1966 Credit Crunch	1966:06	Sharp tightening in financial conditions and bank lending.
Vietnam escalation	1968:01	Rising fiscal and geopolitical uncertainty.
Penn Central crisis	1970:06	Large corporate default and financial stress episode.
Oil embargo	1973:10	Oil-price shock and energy uncertainty.
Second oil shock	1979:01	Iranian Revolution and renewed oil-price spike.
Black Monday	1987:10	Large equity-market crash without subsequent recession.
Asian Financial Crisis	1997:07	Emerging-market and global financial contagion fears.
Russia/LTCM crisis	1998:08	Systemic hedge-fund and liquidity-risk episode.
Y2K	1999:07	Technology and operational-risk concerns.
SARS	2003:02	Global public-health scare.
Euro debt crisis	2010:04	European sovereign default concerns and financial spillovers.
Debt ceiling crisis	2011:05	Fiscal brinkmanship and U.S. downgrade fears.
Fiscal cliff	2012:11	Sharp policy uncertainty regarding fiscal tightening.
Taper tantrum	2013:05	Monetary-policy normalization and rate volatility.
China devaluation	2015:08	Global growth and external-demand concerns.
Brexit referendum	2016:02	Political and policy-regime uncertainty.
Inflation fears	2021:04	Post-pandemic inflation surge and hard-landing concerns.
SVB stress	2023:03	Regional-bank stress and financial-stability fears.

not show the pronounced delayed recessionary hump documented in the main text. Return volatility rises around the event date, consistent with the idea that these episodes are periods of heightened uncertainty, but the increase is less persistent than around recessions. Most importantly, the price–dividend ratio does not exhibit the sharp and sustained decline observed in recession episodes. In both the data and the model, valuations are broadly stable or recover after the risk event rather than collapsing.

This evidence supports the paper's mechanism. Awareness of a new risk can raise uncertainty and temporarily affect discount rates, but awareness alone is not sufficient to generate recession-like valuation dynamics. Large and persistent declines in valuation ratios arise when agents learn that the new-risk episode is severe and persistent. When the shock resolves quickly and the economy avoids a sustained contraction, the model predicts much

weaker asset-price effects, matching the non-recession event evidence.

The contrast between Figure 3 and Figure 14 therefore provides an additional validation of the learning channel. What matters for asset prices is not only that a new risk becomes salient, but whether subsequent output realizations reveal that the risk is associated with a persistent deterioration in macroeconomic conditions.

VII. Conclusion

Our main message is that asset markets price the shape of a crisis. Across NBER recessions, output does not fall instantly at the onset of a downturn. Instead, economic activity typically deteriorates with a delay, reaches a trough, and then rebounds with temporarily elevated growth. Financial markets display a parallel pattern: expected returns and return volatility rise in the early stages of a crisis, peak around the trough, and decline as the economy recovers. Valuation ratios decline during severe downturns and recover only gradually. These dynamics are a robust, cross-crisis feature of the data.

We develop a tractable equilibrium model in which crises are episodes of a new priced risk that gradually intensifies and then dissipates. A crisis-impact process generates a U-shaped path for expected output growth and a temporary elevation in macroeconomic uncertainty. When this crisis block is embedded in habit-based preferences, equilibrium market prices of risk vary over the crisis episode, producing hump-shaped responses of expected returns and return volatility. The model is disciplined using macroeconomic moments only. We do not target equity returns, return volatility, or valuation ratios, so the asset-pricing implications provide an out-of-sample test of the crisis-output mechanism.

The main model also shows why severe crises generate valuation declines. Investors know when a new-risk regime begins, but they do not know its severity. They infer the severity parameter from realized macroeconomic outcomes. Persistent bad realizations lead investors to revise upward the perceived severity of the crisis. These negative belief revisions raise risk

aversion, increase discount rates, and depress valuation ratios. The full-information economy is nested as a special case with zero posterior uncertainty, and it isolates the direct effect of the new priced risk on premia and volatility.

Overall, the framework provides a simple mechanism linking crisis-shaped macro dynamics, learning about crisis severity, and crisis-shaped asset prices. By emphasizing the timing of deterioration and recovery, rather than treating crises as one-time jumps, the model clarifies why risk compensation and volatility can build during downturns, why valuations decline in severe persistent crises, and why asset prices recover as uncertainty about the crisis path is resolved.

REFERENCES

Ai, H., and R. Bansal, 2018, "Risk Preferences and the Macroeconomic Announcement Premium," *Econometrica*, 86(4), 1383–1430.

Ai, H., and A. Bhandari, 2021, "Asset pricing with endogenously uninsurable tail risk," *Econometrica*, 89(3), 1471–1505.

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., and J. Ursúa, 2008, "Consumption disasters in the twentieth century," *American Economic Review*, 98(2), 58–63.

Basu, S., G. Candian, R. Chahrour, and R. Valchev, 2021, "Risky business cycles," Working paper, National Bureau of Economic Research.

Beaudry, P., and F. Portier, 2006, "Stock prices, news, and economic fluctuations," *American Economic Review*, 96(4), 1293–1307.

Beeler, J., and J. Y. Campbell, 2012, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1(1), 141–182.

Bernanke, B. S., M. Gertler, and S. Gilchrist, 1999, "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, 1, 1341–1393.

Blanchard, O. J., 1985, "Debt, deficits, and finite horizons," *Journal of Political Economy*, 93(2), 223–247.

Bloom, N., 2009, "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

Brunnermeier, M., D. Palia, K. A. Sastry, and C. A. Sims, 2021, "Feedbacks: financial markets and economic activity," *American Economic Review*, 111(6), 1845–1879.

Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Christiano, L. J., M. Eichenbaum, and C. L. Evans, 2005, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, 113(1), 1–45.

Gârleanu, N., and S. Panageas, 2015, "Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing," *Journal of Political Economy*, 123(3), 670–685.

Ghaderi, M., M. Kilic, and S. B. Seo, 2022, "Learning, slowly unfolding disasters, and asset prices," *Journal of Financial Economics*, 143(1), 527–549.

Gómez-Cram, R., 2022, "Late to recessions: Stocks and the business cycle," *The Journal of Finance*, 77(2), 923–966.

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

Goyal, A., I. Welch, and A. Zafirov, 2024, "A Comprehensive 2022 Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 37(11), November, 3490–3557.

Hasler, M., and R. Marfè, 2016, "Disaster recovery and the term structure of dividend strips," *Journal of Financial Economics*, 122(1), 116–134.

Jordà, Ò., M. Schularick, and A. M. Taylor, 2011, "Financial crises, credit booms, and external imbalances: 140 years of lessons," *IMF Economic Review*, 59(2), 340–378.

Jordà, Ò., and A. M. Taylor, 2025, "Local projections," *Journal of Economic Literature*, 63(1), 59–110.

Kim, C.-J., J. Morley, and J. Piger, 2005, "Nonlinearity and the permanent effects of recessions," *Journal of Applied Econometrics*, 20(2), 291–309.

Kiyotaki, N., and J. Moore, 1997, "Credit Cycles," *Journal of Political Economy*, 105(2), 211–248.

Kroencke, T. A., 2022, "Recessions and the stock market," *Journal of Monetary Economics*, 131, 61–77.

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

Marfè, R., and J. Pénasse, 2024, "Measuring macroeconomic tail risk," *Journal of Financial Economics*, 156, 103838.

Martin, I., 2017, "What is the Expected Return on the Market?" *The Quarterly Journal of Economics*, 132(1), 367–433.

Menzly, L., T. Santos, and P. Veronesi, 2004, "Understanding predictability," *Journal of Political Economy*, 112(1), 1–47.

Moench, E., and T. Stein, 2025, "Equity premium predictability over the business cycle," *Journal of Financial and Quantitative Analysis*, 2025-09, 1–31.

Morley, J., and J. Piger, 2012, "The Asymmetric Business Cycle," *Review of Economics and Statistics*, 94(1), 208–221.

Muir, T., 2017, "Financial crises and risk premia," *The Quarterly Journal of Economics*, 132(2), 765–809.

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(3), 35–74.

Neftci, S. N., 1984a, "Are economic time series asymmetric over the business cycle?" *Journal of Political Economy*, 92(2), 307–328.

Neftci, S. N., 1984b, "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy*, 92(2), 307–328.

Reinhart, C. M., and K. S. Rogoff, 2009, *This time is different: Eight centuries of financial folly*, Princeton University press.

Sichel, D. E., 1993, "Business Cycle Asymmetry: A Deeper Look," *Economic Inquiry*, 31(2), 224–236.

Smets, F., and R. Wouters, 2007, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 97(3), 586–606.

Tsai, J., and J. A. Wachter, 2015, "Disaster risk and its implications for asset pricing," *Annual Review of Financial Economics*, 7(1), 219–252.

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

Appendix A. Proofs

This appendix proves the propositions in Section IV. Throughout, market clearing implies $C_t = Y_t$. Tildes denote objects evaluated under the representative agent's perceived conditional law of motion. In the learning economy, the agent replaces the true crisis driver $x_t = 1 - g_t\varepsilon$ with the perceived driver $\hat{x}_t = 1 - g_t\hat{\varepsilon}_t$ when forming conditional expectations. The full-information economy is the special case $\nu_t = 0$, so that $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$.

A. Output dynamics

Normal output satisfies

$$d\widehat{Y}_t = \mu_{\widehat{Y}}\widehat{Y}_t dt + \sigma_{\widehat{Y}}\widehat{Y}_t dZ_{\widehat{Y},t}.$$

During crises, $Y_t = \widehat{Y}_t\eta_t$ and

$$d\eta_t = \kappa_\eta(x_t - \eta_t)dt + \sigma_\eta\eta_t(\lambda - \eta_t)dZ_{\eta,t}.$$

Since $d\widehat{Y}_t\, d\eta_t = 0$, Itô's product rule gives

$$\begin{aligned}\frac{dY_t}{Y_t} &= \frac{d\widehat{Y}_t}{\widehat{Y}_t} + \frac{d\eta_t}{\eta_t} \\ &= \left[\mu_{\widehat{Y}} + \kappa_\eta\left(\frac{x_t}{\eta_t} - 1\right)\right]dt + \sigma_{\widehat{Y}}dZ_{\widehat{Y},t} + \sigma_\eta(\lambda - \eta_t)dZ_{\eta,t}, \qquad \omega_t = L. \end{aligned} \tag{A1}$$

In normal times, $\eta_t = 1$, so

$$\frac{dY_t}{Y_t} = \mu_{\widehat{Y}}dt + \sigma_{\widehat{Y}}dZ_{\widehat{Y},t}, \qquad \omega_t = H. \tag{A2}$$

Under the representative agent's perceived law of motion during crises, x_t is replaced by $\hat{x}_t$. Thus,

$$\frac{dY_t}{Y_t} = \left[\mu_{\widehat{Y}} + \kappa_\eta\left(\frac{\hat{x}_t}{\eta_t} - 1\right)\right]dt + \sigma_{\widehat{Y}}d\widetilde{Z}_{\widehat{Y},t} + \sigma_\eta(\lambda - \eta_t)d\widetilde{Z}_{\eta,t}. \tag{A3}$$

Subtracting the perceived conditional mean from the true law of motion gives

$$\begin{aligned}\frac{dY_t}{Y_t} - \widetilde{\mathbb{E}}_t\left[\frac{dY_t}{Y_t}\right] &= \kappa_\eta\left(\frac{x_t - \hat{x}_t}{\eta_t}\right)dt + \sigma_{\widehat{Y}}dZ_{\widehat{Y},t} + \sigma_\eta(\lambda - \eta_t)dZ_{\eta,t} \\ &= -\kappa_\eta\left(\frac{g_t(\varepsilon - \hat{\varepsilon}_t)}{\eta_t}\right)dt + \sigma_{\widehat{Y}}dZ_{\widehat{Y},t} + \sigma_\eta(\lambda - \eta_t)dZ_{\eta,t}. \end{aligned} \tag{A4}$$

Proof of Proposition 1. The law of motion for risk aversion is

$$d\mathcal{R}_t = \kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t)dt - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\left(\frac{dY_t}{Y_t} - \widetilde{\mathbb{E}}_t\left[\frac{dY_t}{Y_t}\right]\right).$$

Using (A4), during crises we obtain

$$\begin{aligned} d\mathcal{R}_t &= \kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t)dt - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\left[-\kappa_\eta\left(\frac{g_t(\varepsilon - \hat{\varepsilon}_t)}{\eta_t}\right)dt + \sigma_{\widehat{Y}}dZ_{\widehat{Y},t} + \sigma_\eta(\lambda - \eta_t)dZ_{\eta,t}\right] \\ &= \left[\kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t) + \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\kappa_\eta\left(\frac{g_t(\varepsilon - \hat{\varepsilon}_t)}{\eta_t}\right)\right]dt \\ &\quad - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_{\widehat{Y}}dZ_{\widehat{Y},t} - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}})\sigma_\eta(\lambda - \eta_t)dZ_{\eta,t}. \end{aligned} \tag{A5}$$

This is equation (22).

For the full-information special case, $\nu_t = 0$ implies $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$. Therefore the drift term involving $g_t(\varepsilon - \hat{\varepsilon}_t)$ is zero, which gives Corollary 1. □

B. Stochastic discount factor

The stochastic discount factor is

$$M_t = e^{-\rho t}\frac{\mathcal{R}_t}{Y_t}. \tag{A6}$$

It is useful to derive dM_t/M_t from Itô's lemma applied to the product $e^{-\rho t}\mathcal{R}_t Y_t^{-1}$. For any perceived consumption process

$$\frac{dY_t}{Y_t} = \widetilde{\mu}_{Y,t}dt + \sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\}\sigma_\eta(\lambda - \eta_t)d\widetilde{Z}_{\eta,t},$$

define

$$\widetilde{\sigma}^2_{Y,t} = \sigma^2_{\widehat{Y}} + \mathbb{1}\{\omega_t = L\}\sigma^2_\eta(\lambda - \eta_t)^2.$$

Under the perceived law of motion, the risk-aversion process has diffusion

$$\frac{d\mathcal{R}_t}{\mathcal{R}_t} = \frac{\kappa_\mathcal{R}(\bar{\mathcal{R}} - \mathcal{R}_t)}{\mathcal{R}_t}dt - \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\left[\sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\}\sigma_\eta(\lambda - \eta_t)d\widetilde{Z}_{\eta,t}\right].$$

Also,

$$\frac{d(Y_t^{-1})}{Y_t^{-1}} = -\frac{dY_t}{Y_t} + \left(\frac{dY_t}{Y_t}\right)^2 = (-\widetilde{\mu}_{Y,t} + \widetilde{\sigma}^2_{Y,t})dt - \sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} - \mathbb{1}\{\omega_t = L\}\sigma_\eta(\lambda - \eta_t)d\widetilde{Z}_{\eta,t}.$$

The quadratic covariation term is

$$\begin{aligned}\frac{d\mathcal{R}_t}{\mathcal{R}_t}\frac{d(Y_t^{-1})}{Y_t^{-1}} &= \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\left[\sigma^2_{\widehat{Y}} + \mathbb{1}\{\omega_t = L\}\sigma^2_\eta(\lambda - \eta_t)^2\right]dt \\ &= \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\widetilde{\sigma}^2_{Y,t}dt. \end{aligned} \tag{A7}$$

Therefore,

$$\begin{aligned}\frac{dM_t}{M_t} &= -\rho dt + \frac{d\mathcal{R}_t}{\mathcal{R}_t} + \frac{d(Y_t^{-1})}{Y_t^{-1}} + \frac{d\mathcal{R}_t}{\mathcal{R}_t}\frac{d(Y_t^{-1})}{Y_t^{-1}} \\ &= -\left[\rho + \widetilde{\mu}_{Y,t} - \widetilde{\sigma}^2_{Y,t} - \kappa_\mathcal{R}\left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}\right) - \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\widetilde{\sigma}^2_{Y,t}\right]dt \\ &\quad - \sigma_{\widehat{Y}}\left[1 + \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\right]d\widetilde{Z}_{\widehat{Y},t} \\ &\quad - \mathbb{1}\{\omega_t = L\}\sigma_\eta(\lambda - \eta_t)\left[1 + \alpha\left(1 - \frac{\lambda_\mathcal{R}}{\mathcal{R}_t}\right)\right]d\widetilde{Z}_{\eta,t}. \end{aligned} \tag{A8}$$

Proof of Proposition 2. Equation (A8) has the form

$$\frac{dM_t}{M_t} = -\widetilde{r}_t dt - \widetilde{\theta}_{\widehat{Y},t} d\widetilde{Z}_{\widehat{Y},t} - \widetilde{\theta}_{\eta,t} d\widetilde{Z}_{\eta,t}.$$

Thus, the market price of normal output risk is

$$\widetilde{\theta}_{\widehat{Y},t} = \sigma_{\widehat{Y}} \left[1 + \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \right],$$

and the market price of crisis-impact risk is

$$\widetilde{\theta}_{\eta,t} = \mathbb{1}\{\omega_t = L\} \sigma_\eta (\lambda - \eta_t) \left[1 + \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \right].$$

The perceived risk-free rate follows from the negative of the drift in (A8):

$$\widetilde{r}_t = \rho + \widetilde{\mu}_{Y,t} - \widetilde{\sigma}^2_{Y,t} + \kappa_{\mathcal{R}} \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t} \right) - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \widetilde{\sigma}^2_{Y,t}. \tag{A9}$$

In normal times,

$$\widetilde{\mu}_{Y,t} = \mu_{\widehat{Y}}, \qquad \widetilde{\sigma}^2_{Y,t} = \sigma^2_{\widehat{Y}},$$

so

$$\widetilde{r}_t = \rho + \mu_{\widehat{Y}} - \sigma^2_{\widehat{Y}} + \kappa_{\mathcal{R}} \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t} \right) - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \sigma^2_{\widehat{Y}}.$$

This gives $r^{log} + r^{habit}_t$.

During crises,

$$\widetilde{\mu}_{Y,t} = \mu_{\widehat{Y}} + \kappa_\eta \left(\frac{\hat{x}_t}{\eta_t} - 1 \right), \qquad \widetilde{\sigma}^2_{Y,t} = \sigma^2_{\widehat{Y}} + \sigma^2_\eta (\lambda - \eta_t)^2.$$

Substituting these terms into (A9) yields

$$\widetilde{r}_t = r^{log} + r^{habit}_t + \kappa_\eta \left(\frac{\hat{x}_t}{\eta_t} - 1 \right) - \sigma^2_\eta (\lambda - \eta_t)^2 - \alpha \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \sigma^2_\eta (\lambda - \eta_t)^2.$$

This is the decomposition in Proposition 2.

If $\nu_t = 0$, then $\hat{x}_t = x_t$. Substituting x_t for $\hat{x}_t$ in the crisis component of the risk-free rate gives Corollary 2. □

C. Market portfolio

The market portfolio is a claim to the aggregate output stream:

$$S_t = \widetilde{\mathbb{E}}_t\left[\int_t^\infty \frac{M_u}{M_t} Y_u \, du\right].$$

Using $M_t = e^{-\rho t}\mathcal{R}_t / Y_t$,

$$\begin{aligned} S_t &= \widetilde{\mathbb{E}}_t\left[\int_t^\infty e^{-\rho(u-t)} \frac{\mathcal{R}_u}{\mathcal{R}_t} Y_t \, du\right] \\ &= Y_t \int_t^\infty e^{-\rho(u-t)} \widetilde{\mathbb{E}}_t\left[\frac{\mathcal{R}_u}{\mathcal{R}_t}\right] du. \end{aligned} \tag{A10}$$

Therefore $S_t = \phi_t Y_t$, where

$$\phi_t = \int_t^\infty e^{-\rho(u-t)} \widetilde{\mathbb{E}}_t\left[\frac{\mathcal{R}_u}{\mathcal{R}_t}\right] du. \tag{A11}$$

Under the perceived conditional law, the drift of $\mathcal{R}_t$ is $\kappa_{\mathcal{R}}(\bar{\mathcal{R}} - \mathcal{R}_t)$, so

$$\widetilde{\mathbb{E}}_t[\mathcal{R}_u] = \bar{\mathcal{R}} + (\mathcal{R}_t - \bar{\mathcal{R}}) e^{-\kappa_{\mathcal{R}}(u-t)}.$$

Thus,

$$\begin{aligned} \phi_t &= \int_t^\infty e^{-\rho(u-t)} \left[\frac{\bar{\mathcal{R}}}{\mathcal{R}_t} + \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}\right) e^{-\kappa_{\mathcal{R}}(u-t)}\right] du \\ &= \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}\frac{1}{\rho} + \left(1 - \frac{\bar{\mathcal{R}}}{\mathcal{R}_t}\right)\frac{1}{\rho + \kappa_{\mathcal{R}}} \\ &= \phi_0 + \phi_{\mathcal{R}} \frac{1}{\mathcal{R}_t}, \end{aligned} \tag{A12}$$

where

$$\phi_0 = \frac{1}{\rho + \kappa_{\mathcal{R}}}, \qquad \phi_{\mathcal{R}} = \bar{\mathcal{R}} \frac{\kappa_{\mathcal{R}}}{\rho(\rho + \kappa_{\mathcal{R}})} = \bar{\mathcal{R}} \phi_0 \frac{\kappa_{\mathcal{R}}}{\rho}.$$

Proof of Proposition 3. The price-dividend ratio follows from (A11)–(A12). It remains to derive the return exposures, expected return, and volatility.

Since

$$\phi_t = \phi_0 + \phi_{\mathcal{R}} \mathcal{R}_t^{-1},$$

Itô's lemma gives

$$d\phi_t = -\phi_{\mathcal{R}} \mathcal{R}_t^{-2} d\mathcal{R}_t + \phi_{\mathcal{R}} \mathcal{R}_t^{-3} (d\mathcal{R}_t)^2. \tag{A13}$$

Only the diffusion terms are needed for return exposures. Under the perceived law,

$$d\mathcal{R}_t = \ldots dt - \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}}) \left[\sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\} \sigma_\eta (\lambda - \eta_t) d\widetilde{Z}_{\eta,t} \right].$$

Therefore,

$$\begin{aligned} \frac{d\phi_t}{\phi_t} &= \ldots dt + \frac{\phi_{\mathcal{R}}}{\phi_t \mathcal{R}_t^2} \alpha(\mathcal{R}_t - \lambda_{\mathcal{R}}) \left[\sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\} \sigma_\eta (\lambda - \eta_t) d\widetilde{Z}_{\eta,t} \right] \\ &= \ldots dt + \alpha \frac{\phi_{\mathcal{R}} / \mathcal{R}_t}{\phi_t} \left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t} \right) \left[\sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\} \sigma_\eta (\lambda - \eta_t) d\widetilde{Z}_{\eta,t} \right]. \end{aligned} \tag{A14}$$

Because $S_t = Y_t \phi_t$,

$$\frac{dS_t}{S_t} = \frac{dY_t}{Y_t} + \frac{d\phi_t}{\phi_t} + \frac{dY_t}{Y_t} \frac{d\phi_t}{\phi_t}.$$

The quadratic covariation term affects only the drift, not the diffusion. Combining the diffusion terms in dY_t/Y_t and $d\phi_t/\phi_t$ gives

$$\frac{dS_t}{S_t} = \ldots dt + V_{R,t} \sigma_{\widehat{Y}} d\widetilde{Z}_{\widehat{Y},t} + \mathbb{1}\{\omega_t = L\} V_{R,t} \sigma_\eta (\lambda - \eta_t) d\widetilde{Z}_{\eta,t},$$

where

$$V_{R,t} = 1 + \alpha \frac{\phi_{\mathcal{R}}/\mathcal{R}_t}{\phi_t}\left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right)$$
$$= 1 + \left(\frac{\kappa_{\mathcal{R}}\frac{\bar{\mathcal{R}}}{\mathcal{R}_t}}{\rho + \kappa_{\mathcal{R}}\frac{\bar{\mathcal{R}}}{\mathcal{R}_t}}\right)\alpha\left(1 - \frac{\lambda_{\mathcal{R}}}{\mathcal{R}_t}\right). \tag{A15}$$

Hence,

$$\sigma_{R,\widehat{Y},t} = \sigma_{\widehat{Y}} V_{R,t},$$

and

$$\sigma_{R,\eta,t} = \begin{cases} V_{R,t}\sigma_\eta(\lambda - \eta_t) & \text{if} \quad \omega_t = L, \\ 0 & \text{if} \quad \omega_t = H. \end{cases}$$

Independence of the Brownian shocks gives

$$\sigma_{R,t} = V_{R,t}\sqrt{\sigma_{\widehat{Y}}^2 + \mathbb{1}\{\omega_t = L\}\sigma_\eta^2(\lambda - \eta_t)^2}.$$

Finally, the expected return follows from the pricing identity

$$\widetilde{\mathbb{E}}_t\left[\frac{dM_t}{M_t}dR_t\right] + \widetilde{\mathbb{E}}_t[dR_t] - \widetilde{r}_t dt = 0,$$

or equivalently,

$$\widetilde{\mu}_{R,t} - \widetilde{r}_t = \widetilde{\theta}_{\widehat{Y},t}\sigma_{R,\widehat{Y},t} + \widetilde{\theta}_{\eta,t}\sigma_{R,\eta,t}.$$

Substituting the exposures gives

$$\widetilde{\mu}_{R,t} = \widetilde{r}_t + \widetilde{\theta}_{\widehat{Y},t}\sigma_{\widehat{Y}}V_{R,t} + \widetilde{\theta}_{\eta,t}\sigma_\eta(\lambda - \eta_t)V_{R,t}.$$

This proves Proposition 3.

If $\nu_t = 0$, then $\hat{\varepsilon}_t = \varepsilon$ and $\hat{x}_t = x_t$. Therefore all perceived pricing objects coincide with their full-information counterparts, which proves Corollary 3. □

Appendix B. Calibration details

Unconditional model simulation

We simulate 10,000 simulations spanning 100 years of quarterly data. The economy alternates between a normal state and a crisis state. We start from the normal state, the economy switches to crisis state through a Poisson jump arrival process with intensity ν, which is one of the estimated parameters.

Upon a crisis start, the model activates the crisis driver x_τ (where τ counts time since crisis onset) and the crisis impact η_t begins to mean-revert toward x_τ according to (3). When η_t reaches 1, the crisis ends and the economy returns to the normal state.

Model-implied local projections from simulated data

We recover the model's impulse response using the same local-projection procedure as in the data. Specifically, in each simulated path we construct a crisis-onset indicator $\mathbb{1}\{\text{CrisisStart}_t^{(j)}\}$ that equals one in the first period of a simulated crisis episode (and zero otherwise). To reduce Monte Carlo noise and to mirror an event-study design, we pool the simulated data across Monte Carlo replications and estimate the local projections on a single stacked panel.

Specifically, for each Monte Carlo replication $j = 1, \ldots, MC$, we simulate a time series $\{y_t^{(j)}, \mathbb{1}\{\text{CrisisStart}_t^{(j)}\}, X_{t-1}^{(j)}\}_{t=1}^{T}$, where $\mathbb{1}\{\text{CrisisStart}_t^{(j)}\}$ equals one in the first period of a simulated crisis episode (and zero otherwise) and $y_t^{(j)}$ is the simulated crisis impact η variable. We then stack all replications into one pooled dataset,

$$\mathcal{D} \equiv \left\{ (y_t^{(j)}, \mathbb{1}\{\text{CrisisStart}_t^{(j)}\}, X_{t-1}^{(j)}) : j = 1, \ldots, MC;\ t = 1, \ldots, T \right\},$$

treating each simulation as a separate "unit" and time as the within-unit dimension. For

each horizon $h = 0, \ldots, H$, we run a pooled OLS regression on $\mathcal{D}$:

$$y_{t+h}^{(j)} = \alpha_h + \beta_h^{\text{model}}(\theta)\, \mathbb{1}\{\text{CrisisStart}_t^{(j)}\} + \Gamma_h' X_{t-1}^{(j)} + \varepsilon_{t+h}^{(j)}. \tag{B1}$$

Thus, the model-implied local-projection coefficients $\{\beta_h^{\text{model}}(\theta)\}_{h=0}^{H}$ are estimated directly from the stacked Monte Carlo panel, rather than by estimating separate regressions within each replication and averaging coefficients ex post. In practice, stacking delivers the same population estimates but substantially reduces simulation noise.

To avoid spurious dynamics across simulation boundaries when constructing lags and leads, we insert missing-value separators between replications (equivalently, we treat each replication as a distinct panel unit). This ensures that lagged controls $X_{t-1}^{(j)}$ and forward outcomes $y_{t+h}^{(j)}$ are formed within a simulation path and never across paths.

This construction ensures an apples-to-apples comparison: both in the data and in the model, the impulse response is identified as the conditional mean path of the output gap following crisis onset, as recovered by local projections, proxied with $\eta - 1$ in the model.

We summarize the model-implied dynamics using the same seven moments as in the data. Moments M1–M3 are computed from the model-implied impulse response $\{\beta_h^{\text{model}}(\theta)\}_{h=0}^{H}$. Moments M4–M5 are computed from simulated crisis episodes by measuring the time from crisis onset to the output trough and the time to recovery back to the pre-crisis output level. Moment M6 is the simulated crisis frequency, i.e., the share of simulated periods classified as crises. Moment M7 is the ratio of the volatility of year-on-year output growth in simulated crisis periods relative to simulated normal times.

Formally,

$$m_1^{\text{model}}(\theta) \equiv \min_{0\le h\le H} \beta_h^{\text{model}}(\theta), \tag{B2}$$

$$m_2^{\text{model}}(\theta) \equiv \arg\min_{0\le h\le H} \beta_h^{\text{model}}(\theta), \tag{B3}$$

$$m_3^{\text{model}}(\theta) \equiv \min\left\{h > m_2^{\text{model}}(\theta) : \beta_h^{\text{model}}(\theta) \ge 0\right\}, \tag{B4}$$

$$m_4^{\text{model}}(\theta) \equiv \text{Std}_{c\in\mathcal{C}^{\text{model}}}\left(T_c^{\text{trough,model}}\right), \qquad T_c^{\text{trough,model}} \equiv t_c^{\text{trough,model}} - t_c^{\text{start,model}}, \tag{B5}$$

$$m_5^{\text{model}}(\theta) \equiv \text{Std}_{c\in\mathcal{C}^{\text{model}}}\left(T_c^{\text{rec,model}}\right), \qquad T_c^{\text{rec,model}} \equiv t_c^{\text{rec,model}} - t_c^{\text{start,model}}, \tag{B6}$$

$$m_6^{\text{model}}(\theta) \equiv \frac{N^{\text{crisis,model}}}{N^{\text{total,model}}}, \tag{B7}$$

$$m_7^{\text{model}}(\theta) \equiv \frac{\text{Std}\left(g_t^{yoy,(j)} \mid t \in \mathcal{T}^{\text{crisis,model}}\right)}{\text{Std}\left(g_t^{yoy,(j)} \mid t \in \mathcal{T}^{\text{normal,model}}\right)}, \qquad g_t^{yoy,(j)} \equiv \frac{Y_t^{(j)}}{Y_{t-4}^{(j)}} - 1, \tag{B8}$$

where $\mathcal{C}^{\text{model}}$ is the set of simulated crisis episodes, $t_c^{\text{start,model}}$ denotes the crisis-onset date in the simulation, $t_c^{\text{trough,model}}$ is the date at which output reaches its minimum within episode c, $t_c^{\text{rec,model}}$ is the first date after the trough at which output returns to (or exceeds) its pre-crisis level, and $\mathcal{T}^{\text{crisis,model}}$ and $\mathcal{T}^{\text{normal,model}}$ denote the sets of simulated crisis and non-crisis periods, respectively.

Appendix A. Parameter space and objective function

We estimate the parameters governing new-risk episodes by simulated method of moments. The estimated parameter vector is

$$\theta = \left(\kappa_\eta, \sigma_\eta, \kappa_1, \kappa_2, \bar{\varepsilon}, \sigma_\varepsilon, \nu\right)',$$

where κ_η and σ_η govern the stochastic dynamics of the crisis-impact process, κ_1 and κ_2 govern the deterministic shape of the episode, $\bar{\varepsilon}$ and σ_ε describe cross-episode severity, and

ν is the arrival intensity of new-risk episodes. We impose

$$\kappa_\eta > 0, \qquad \sigma_\eta > 0, \qquad \kappa_2 > \kappa_1 > 0, \qquad \sigma_\varepsilon > 0, \qquad \nu > 0.$$

The estimated arrival intensity ν governs the frequency of new-risk episodes, not mechanically the frequency of NBER-style recessions. In the simulations, a new-risk episode is classified as a recession only if its realized path deteriorates sufficiently and persists long enough to satisfy the recession-classification rule. Thus, the model allows for false positives: episodes in which the new-risk process becomes active but the realized path is too mild or too short-lived to be classified as a recession. This distinction is important for mapping the model to the data, because the empirical event studies condition on NBER recessions, which correspond to the subset of new-risk episodes that ultimately become realized recessions.

The estimation targets seven macroeconomic moments: average recession depth, average time from peak to trough, average time from peak to recovery, the dispersion of peak-to-trough and peak-to-recovery times, the share of periods classified as recessions, and the ratio of output volatility in recessions to normal times. Equity returns, valuation ratios, and stock-market volatility are not used in the estimation.

Let

$$m^{data} = (m_1^{data}, \ldots, m_7^{data})', \qquad m^{model}(\theta) = (m_1^{model}(\theta), \ldots, m_7^{model}(\theta))'.$$

The SMM estimator solves

$$\widehat{\theta} = \arg\min_{\theta \in \Theta} \left[m^{data} - m^{model}(\theta) \right]' W \left[m^{data} - m^{model}(\theta) \right], \tag{B9}$$

where W is a diagonal weighting matrix with entries

$$W_{jj} = \frac{1}{(m_j^{data})^2}, \qquad j = 1, \ldots, 7.$$

This weighting expresses the objective in proportional deviations from the data moments.

We define the following lower and upper bounds for θ:

$$\theta_{min} = [0.01, 0.01, 0.1, 1.00, 0.01, 0.01, 0.01],$$

$$\theta_{max} = [2.00, 1.00, 4.00, 8.00, 1.00, 1.50, 1.00].$$

The remaining model parameters $(\mu_{\hat{Y}}, \sigma_{\hat{Y}})$ are set to match the GDP growth rate (Y-on-Y) and its volatility measuring during normal times over our sample period (see Panel A in Table II). Preference parameters $(\rho, \lambda_{\mathcal{R}}, \kappa_{\mathcal{R}}, \bar{\mathcal{R}})$ are the same as in Menzly, Santos, and Veronesi (2004) (see Panel B in Table II).

To minimize the objective function (B9), we use particle swarm optimization (`particleswarm`) to explore the parameter space under bounds and the constraint $\kappa_2 > \kappa_1$. To reduce Monte Carlo noise in objective comparisons, we use common random numbers (fixed RNG seed) across candidate parameters. Starting from the best particle-swarm solution, we refine the estimate using a local constrained solver (`fmincon`). This step improves precision once the algorithm has located a reasonable parameter space.