Credit Markets, Uncertainty, and Business Cycles*

Swapnil Sharma†

[Click here for the latest version]

Abstract

The spectrum of U.S. hours worked has a pronounced peak at business-cycle frequencies. This paper investigates the economic forces behind this peak. The hours-based main business-cycle shock identified from a vector autoregression (VAR) that excludes measures of uncertainty and credit conditions does not reproduce the peak. Once macro uncertainty and the credit spread are added, however, the resulting shock does. This shock is nearly identical to one selected to explain the largest possible share of business-cycle variation in macro uncertainty and the credit spread; the two estimated shock series have a correlation of 0.97. Separating this joint movement uncovers two orthogonal shocks that propagate differently. An uncertainty shock produces a broad contraction, followed by a gradual recovery toward steady state. By contrast, a credit-related shock generates boom–bust dynamics: an expansion followed by a contraction. The credit-related shock accounts for the peak in the hours spectrum. It also produces reversals in stock prices, the excess bond premium, and the composition of bond issuance. We develop a model in which firms borrow against collateral to finance their wage bill. Uncertainty about enforcement tightens financing conditions and contracts activity. In the model, the credit-related shock takes the form of a favorable noise shock about future total factor productivity (TFP). It initially relaxes financing conditions; when the signal is revealed to be noise, inherited debt contributes to a subsequent tightening. Nominal wage rigidity allows both shocks to generate comovement in output, consumption, investment, and hours.

*I am grateful to Rosen Valchev, Peter Ireland, Susanto Basu, and other faculty members at Boston College for their guidance and valuable feedback. I also thank Sanjay Moorjani for useful discussions and participants in the Boston College Macroeconomics Lunch Seminar for helpful comments.

†Department of Economics, Boston College. Email: swapnil.sharma@bc.edu.

1 Introduction

Beaudry et al. (2020) document that U.S. macroeconomic activity is concentrated at extended business-cycle frequencies. Hours worked provides a particularly clear example: its spectral density has a pronounced peak within the business-cycle range. This peak is an unconditional feature of the data. It shows where variation is concentrated, but not which shocks account for it or how their effects propagate through the economy. This paper asks which disturbances account for the spectral peak in hours and which propagation mechanisms concentrate their effects at business-cycle frequencies.

We begin with the main business-cycle (MBC) shock of Angeletos et al. (2020), identified as the shock that explains the largest share of hours variation over periodicities of 6 to 40 quarters. We first apply their procedure to a ten-variable VAR containing output, investment, consumption, hours, unemployment, the labor share, the federal funds rate, inflation, labor productivity, and total factor productivity. This system contains no measure of macro uncertainty or credit conditions, and the hours spectrum attributable to the MBC shock does not reproduce the peak in the data. We then add macro uncertainty and the credit spread to the system. In this expanded twelve-variable VAR, the MBC shock does reproduce the spectral peak. The result also holds in a more parsimonious eight-variable specification containing output, consumption, investment, hours, the federal funds rate, utilization-adjusted TFP, the credit spread, and macro uncertainty. Thus, adding uncertainty and credit-spread information changes the dominant hours-based rotation in a way that matters for its frequency composition.

Within the eight-variable system, the shock that jointly maximizes variation in macro uncertainty and the credit spread generates macroeconomic responses close to those of the MBC shock. The two estimated shock series have a correlation of 0.97 (0.96 in the twelve-variable VAR). The movement that accounts for most of the business-cycle variation in hours is therefore closely associated with the dominant joint movement in uncertainty and credit spreads. Because macro uncertainty and the credit spread move together but also vary separately, the max-share criterion may combine distinct shocks that propagate differently through the economy.

We use a sequential penalty-function approach related to Caldara et al. (2016) to distinguish two directions within this joint movement. The first shock maximizes the business-cycle variation in macro uncertainty. An increase in uncertainty causes output, consumption, investment, and hours to fall together and then gradually recover. The shock accounts for nearly 80 percent of uncertainty variation in the target band. We then identify a second shock, orthogonal to the uncertainty shock, that maximizes the remaining joint variation in uncertainty and the credit spread. A decline in the spread produces an expansion lasting roughly three years, followed by a contraction in all four macroeconomic variables. This credit-related shock explains more than half of the business-cycle variation in the spread but little of the variation in uncertainty. Most importantly, this shock generates the spectral peak in hours: the variation in hours attributable to it is concentrated at the same frequencies as the peak observed in the data.

The behavior of broader financial indicators provides further evidence that the second disturbance reflects changes in credit conditions. Stock prices rise during the expansion and subsequently fall below their initial level. The excess bond premium initially declines and later increases, while the share of high-yield issuance rises and then reverses. The shock also accounts for more than half of the variation in the excess bond premium. These responses resemble the periods of easy credit and subsequent retrenchment documented by Greenwood and Hanson (2013) and López-Salido et al. (2017). Together, these patterns motivate our interpretation of the second disturbance as a credit-related shock. This evidence comes from variables that are not used to identify the shock: the credit spread enters the identifying criterion, but stock prices, the excess bond premium, and high-yield issuance do not.

The two shocks place a joint restriction on theory. Both must generate comovement in output, consumption, investment, and hours, but their propagation must differ. Higher uncertainty produces a contraction followed by recovery, whereas the credit-related shock produces an expansion followed by a contraction. The uncertainty response presents an additional difficulty. In a standard flexible-price model, precautionary saving lowers consumption but tends to raise labor supply, output, and investment (Basu and Bundick, 2017). A successful model must overcome this comovement problem while allowing the response to depend on the source of the disturbance.

We develop a rational-expectations model in which credit conditions affect production. Firms pay wages before receiving sales revenue, so they borrow against the value of installed capital. Because firms cannot freely adjust equity payouts, a decline in borrowing capacity leads them to reduce hiring and investment. Investment adjustment costs make the value of installed capital, and therefore collateral, respond immediately even though the current capital stock is predetermined. Nominal wage rigidity then allows lower labor demand to reduce hours together with output, consumption, and investment.

The empirical credit-related shock identifies the dynamics to be explained, but it does not reveal the disturbance that starts the expansion. We model it as a favorable but noisy signal about future TFP. When the signal arrives, firms do not know whether it contains genuine news or only noise. They respond to the prospect of higher productivity by hiring, investing, and borrowing. This follows the noisy-information literature, in which agents initially respond to the signal and learn its true content only as fundamentals are revealed (Lorenzoni, 2009; Forni et al., 2017; Chahrour and Jurado, 2018). In our model, the initial expansion raises collateral values and can relax the borrowing constraint. At maturity, the signal is resolved. If the anticipated increase in TFP does not materialize, collateral values fall while the debt accumulated during the expansion remains. Borrowing capacity tightens, and the expansion turns into a contraction. Thus, financial frictions turn a Lorenzoni-style noise-driven expansion into a boom–bust cycle, while nominal wage rigidity generates comovement across the four macroeconomic quantities.

Macro uncertainty reaches the financing constraint through a different margin. An increase in

uncertainty about enforcement raises the likelihood of future states in which borrowing capacity is scarce. The equilibrium financing wedge rises, causing firms to reduce hiring and investment. Lower investment reduces collateral value and reinforces the tightening. As uncertainty fades, the wedge declines and activity gradually returns toward its steady state. Unlike the noise shock, this response does not depend on a revision in expected productivity or on the constraint moving from binding to slack and back again. The two shocks therefore interact with the same financing constraint but propagate differently: uncertainty produces a precautionary tightening, whereas the noisy signal initially expands collateral and borrowing capacity before these movements reverse when the signal is resolved.

An occasionally binding constraint can also be solved with higher-order perturbation methods, including DynareOBC. We use a Deep Equilibrium Network because it combines several features needed for estimation. It preserves nonlinear uncertainty effects while representing binding and slack outcomes within a single global decision rule, rather than relying on separate local solutions for the two regimes. The estimated parameters enter the network as additional inputs, so one trained network approximates the model's solution across the parameter space and produces model impulse responses without requiring the model to be solved again for every candidate parameter vector. The remaining horizon and magnitude of a pending signal also enter the state vector. This allows us to estimate how far in advance agents receive the noisy signal, together with the other structural parameters, rather than fixing the news horizon before solving the model. We estimate these parameters by matching the model responses to the empirical uncertainty and credit-related shocks. The estimated model reproduces the central distinction in the data: uncertainty generates a hump-shaped contraction in all four macroeconomic quantities, while a favorable signal subsequently revealed as noise generates an expansion followed by a contraction. The model's Tobin's Q also rises and falls, providing a qualitative asset-price implication consistent with the reversal in stock prices, although it has no exact counterpart in the empirical VAR.

The rest of the paper proceeds as follows. The next subsection places the paper in the related literature. Section 2 describes the data and empirical method, compares the MBC shock across information sets, and then sequentially identifies the uncertainty and credit-related shocks. It concludes with the auxiliary financial evidence and the implications for theory. Section 3 develops the model's financing mechanism, introduces the two shocks, and explains how they generate different dynamics through the same channel. It then presents the calibration, estimation, and quantitative results. The final section concludes.

1.1 Literature Review

The paper brings together evidence on the frequency composition of business cycles and work that identifies their principal driving forces. Beaudry et al. (2020) document that aggregate activity is concentrated at extended business-cycle frequencies and develop a stochastic limit-cycle model

to account for this regularity. Their evidence establishes the spectral pattern that motivates our analysis; our question is which disturbances generate the peak in hours and what their propagation reveals about its economic source. To answer it, we build on the max-share method of Uhlig (2003), which identifies the innovation accounting for the largest fraction of a variable's variation. Angeletos et al. (2020) apply this logic across variables and frequency bands and recover a main business-cycle shock with a common macroeconomic footprint. We show that whether the hours-based MBC shock reproduces the spectral peak depends on whether macro uncertainty and the credit spread are included in the VAR. We then separate their dominant joint movement into an uncertainty shock and an orthogonal credit-related shock and find that the latter accounts for the peak.

This result also relates to recent evidence that a single MBC shock can combine distinct economic disturbances. Using a large dynamic factor model, Forni et al. (2025) identify a transitory demand shock and a permanent supply shock. Their demand shock accounts for most cyclical variation in the credit spread, while their supply shock has an important technology-news component. Granese (2026) likewise finds that a data-rich system recovers two dominant shocks and argues that a small VAR can suffer from informational deficiency. We distinguish movements within the broad demand category by using macro uncertainty and the credit spread in a sequential identification. This separates a contractionary uncertainty shock from a credit-related shock that generates the hours-spectrum peak. Our findings concern the disturbances selected by an hours-based criterion in our information set; they do not imply that all low-frequency cyclicality originates in credit markets.

The paper most closely related to our question is Brianti and Cormun (2024). They isolate the component of survey expectations that is orthogonal to observed fundamentals and show that it produces boom–bust dynamics, which they rationalize with self-fulfilling expectations and financial frictions. We instead begin with the hours-based MBC shock and show how its frequency composition changes when uncertainty and credit-spread information are introduced. Our sequential identification uncovers both the credit-related shock that generates the spectral peak and the contractionary uncertainty shock with which it is combined. The model is therefore required to produce two propagation patterns through a common financing environment. It does so with rational noisy information rather than self-fulfilling beliefs.

The paper also contributes to the literature on credit conditions and future economic activity. Gilchrist and Zakrajšek (2012) show that the excess bond premium contains information about subsequent activity and that innovations in this premium generate broad macroeconomic contractions. Greenwood and Hanson (2013) relate the composition of bond issuance to investor sentiment and future corporate-bond returns. Most closely, López-Salido et al. (2017) find that periods in which credit risk is priced aggressively are followed by wider spreads and weaker activity. Evidence based on credit quantities reaches a similar conclusion: rapid credit growth predicts financial crises and deeper subsequent recessions (Schularick and Taylor, 2012; Jordà

et al., 2013), while household-debt expansions forecast lower medium-run growth (Mian et al., 2017). These papers establish that easy credit conditions contain information about later contractions. We recover an orthogonal credit-related shock and trace the full macroeconomic response following the innovation. The initial expansion in output, consumption, investment, and hours is followed by a contraction, and stock prices display a similar reversal. More importantly for our question, the variation in hours attributable to this shock is concentrated around the spectral peak observed in the data. Its explanatory power for the excess bond premium and its relation to the high-yield issuance share connect this cyclical response to the same credit-market evidence without imposing investor sentiment as the primitive disturbance.

A second strand studies uncertainty as a source of recessions. The empirical analysis uses the broad macroeconomic uncertainty measure of Jurado et al. (2015). Bloom (2009) documents a sharp decline and rebound in activity after an uncertainty innovation, and Ludvigson et al. (2021) emphasizes the importance of separating exogenous uncertainty movements from the endogenous response of uncertainty to the business cycle. A central issue for this literature is comovement. Basu and Bundick (2017) show that higher uncertainty lowers output, consumption, investment, and hours in the data, whereas a standard flexible-price model predicts that precautionary labor supply raises hours, output, and investment. Nominal rigidity resolves this conflict by allowing lower demand to reduce production and factor demand. Leduc and Liu (2016) obtain a related contraction through nominal rigidity and labor-market search, while Arellano et al. (2019) use financial frictions to transmit greater firm-level volatility into lower input demand and tighter credit conditions. An alternative resolution comes from Chatterjee et al. (2025). In their flexible-price economy, stochastic volatility in credit conditions raises the expected scarcity of borrowing capacity and the associated financing wedge. The resulting inward shift in labor demand lowers output, consumption, investment, and hours despite the precautionary increase in labor supply. Their empirical measure is uncertainty in credit growth, whereas we identify a shock to the broad macroeconomic uncertainty measure of Jurado et al. (2015). Basu et al. (2025) provide a real mechanism in which higher risk induces firms to substitute toward safer, less productive inputs. Relative to this work, we study uncertainty jointly with credit conditions and show that their dominant movement combines distinct conditional responses. The uncertainty shock produces recessionary comovement, while the orthogonal credit-related shock produces an expansion followed by a contraction.

On the model side, our paper is related to models in which financial conditions originate or amplify aggregate fluctuations. In Jermann and Quadrini (2012), shocks to an enforcement constraint make firms' ability to borrow an independent source of business cycles. Financial constraints instead propagate shocks that originate elsewhere in Kiyotaki and Moore (1997) and Bernanke et al. (1999). Our model combines these roles and adapts the uncertainty mechanism of Chatterjee et al. (2025). Uncertainty about enforcement raises the expected cost of scarce borrowing capacity and produces an immediate contraction. A noisy favorable signal about future

productivity operates differently. It initially supports hiring, investment, and borrowing; when agents learn that the signal was noise, the debt accumulated during the expansion contributes to a subsequent tightening. The main perturbation analysis in Chatterjee et al. (2025) treats the constraint as always binding, although their global analysis allows it to become slack. In our model, regime switching is part of the mechanism: the constraint can become slack during the expansion and bind again when the signal is resolved. This allows one financing environment to produce the contractionary uncertainty response and the boom–bust response to the credit-related shock.

The paper is therefore also related to theories of endogenous boom–bust dynamics. Self-fulfilling fluctuations arise from financial constraints and strategic complementarities in Benhabib and Wang (2013), Liu and Wang (2014), and Azariadis et al. (2016). Together with the stochastic limit-cycle mechanism of Beaudry et al. (2020), these models make recurrent movement a property of the economy's propagation mechanism. Our empirical results instead call for propagation that depends on the disturbance: uncertainty produces contraction and recovery, whereas the credit-related shock produces an expansion and a subsequent contraction. Bordalo et al. (2026) obtain credit reversals from diagnostic expectations: optimism encourages debt and investment, and the later correction in beliefs raises spreads and depresses activity. Our mechanism shares the connection between expectations, leverage, and reversal, but the signal and its eventual resolution are consistent with rational expectations. The noisy-information structure follows the demand-shock mechanism of Lorenzoni (2009) and the business-cycle application of Forni et al. (2017). The equivalence between news and noise representations emphasized by Chahrour and Jurado (2018) provides the information-theoretic basis for our formulation.

Finally, the solution and estimation draw on the computational literature for nonlinear dynamic models. Maliar et al. (2021) use neural networks to solve high- dimensional models by expressing equilibrium restrictions as nonlinear regression equations. Azinovic et al. (2022) develop Deep Equilibrium Nets, which learn policy functions by minimizing equilibrium residuals along simulated paths and can accommodate occasionally binding constraints. For estimation, Kase et al. (2025) integrate a global neural-network solution with the estimation of a nonlinear heterogeneous-agent model, while Wei and Jiang (2025) train a network to infer structural parameters from simulated moments. Our application uses the parameter-as-state strategy described by Scheidegger (2026): one network learns the equilibrium over both the economic state space and the relevant parameter region. We then estimate the model by matching model-implied and empirical impulse responses without resolving it at every candidate parameter vector. This preserves the nonlinear transition between binding and slack credit constraints, allows the signal horizon to be estimated with the other structural parameters, and keeps the estimation criterion explicit.

2 Empirical Analysis

We begin the empirical analysis by documenting evidence of cyclical dynamics in U.S. macroeconomic data. In particular, hours worked has a pronounced spectral peak at business-cycle frequencies, and its medium-frequency autocorrelation changes sign as the horizon increases. To identify the shocks underlying these dynamics, we combine the max-share approach developed by Uhlig (2003) with the penalty-function approach of Caldara et al. (2016). This section first describes the data and then presents the two identification procedures. We next report the empirical findings in three stages. We recover the shock that accounts for the largest share of business-cycle variation in hours, examine its connection with macro uncertainty and credit spreads, and then separate an uncertainty shock from an orthogonal credit-spread shock. Finally, we use stock prices, the excess bond premium, and the high-yield issuance share to clarify the economic content of the latter shock.

2.1 Data

The quarterly sample runs from 1973:I through 2024:IV. We organize the analysis around three VARs. The first is the ten-variable system of Angeletos et al. (2020), containing output, consumption, investment, hours worked, unemployment, the labor share, the federal funds rate, inflation, labor productivity, and total factor productivity (TFP). This system provides the reference specification for the hours-based main business-cycle shock. The second adds the credit spread and macro uncertainty. Comparing the ten- and twelve-variable systems holds the maximum-share criterion fixed while changing whether the VAR contains direct information about uncertainty and credit conditions.

Our baseline results use a more parsimonious eight-variable VAR containing output, consumption, investment, hours worked, the federal funds rate, utilization-adjusted TFP, the credit spread, and macro uncertainty. It retains the four measures of aggregate activity and the two variables needed for the sequential identification while excluding unemployment, the labor share, inflation, and labor productivity. We use the twelve-variable system to verify that the results are not produced by dropping these variables from the larger specification.

Output, consumption, investment, and hours are expressed as 100 times their natural logarithms. Output is real GDP per capita. Consumption consists of nondurable goods and services, while investment combines gross private domestic investment and consumer durables. Both aggregates are converted to real per-capita terms. Hours are total hours in the nonfarm business sector divided by the civilian noninstitutional population aged sixteen and older. These definitions follow Angeletos et al. (2020). Unemployment is the civilian unemployment rate, the labor share is the nonfarm business labor-share index, and labor productivity is nonfarm business output per hour. Inflation is the quarterly growth rate of the GDP implicit price deflator.

The interest rate is the quarterly effective federal funds rate. TFP is the utilization-adjusted

series constructed by Fernald (2014). The credit spread is the Gilchrist–Zakrajšek corporate bond spread, and macro uncertainty is the twelve-month-ahead measure of Jurado et al. (2015). The uncertainty index measures uncertainty about future macroeconomic conditions rather than financial-market volatility. These two series enlarge the information available to the ten-variable system and provide the targets for the joint and sequential identification exercises.

Three additional series help interpret the identified credit shock. The first is the real per-capita S&P index. The second is the excess bond premium of Gilchrist and Zakrajšek (2012), which removes the component of corporate spreads explained by expected defaults. The third is the share of high-yield issuance in total corporate bond issuance constructed by Greenwood and Hanson (2013). This high-yield ratio is a measure of the pricing of credit risk and is closely related to the credit-market sentiment indicators studied by López-Salido et al. (2017). Because these series do not share a common sample, they enter as auxiliary variables rather than enlarging the baseline VAR.

Appendix A gives the sources, transformations, and timing conventions for every series.

2.2 Empirical Method

Let x_t denote the $N \times 1$ vector of observables in a given information set. For each of the ten-, twelve-, and eight-variable systems, the reduced-form VAR is

$$x_t = a + A_1 x_{t-1} + \cdots + A_p x_{t-p} + u_t, \qquad \mathbb{E}(u_t u_t') = \Sigma, \tag{1}$$

where the lag length is $p = 2$. Each VAR is estimated with the same Minnesota prior and Gibbs sampler, following Angeletos et al. (2020). Holding the prior, lag length, sample, frequency band, and identifying objective fixed allows the comparison across systems to isolate the role of the additional information. We use 50,000 draws, discard the first 49,000, and calculate posterior medians and 68 percent posterior intervals from the retained draws.

Write $\Sigma = SS'$ and let q be a unit vector in the space of orthonormal structural innovations. The corresponding reduced-form impact vector is Sq. If B_h is the moving-average coefficient at horizon h, the impulse response is $B_h Sq$. Define

$$C(e^{-i\omega}) = \sum_{h=0}^{\infty} B_h S e^{-i\omega h}. \tag{2}$$

For variable k, the variation attributable to q over a frequency band $\mathcal{B} = [\underline{\omega}, \overline{\omega}]$ is proportional to

$$q' M_k(\mathcal{B}) q, \qquad M_k(\mathcal{B}) = \int_{\mathcal{B}} C_k(e^{-i\omega})' C_k(e^{-i\omega})\, d\omega, \tag{3}$$

where $C_k(\cdot)$ is row k of the transfer function. Dividing $M_k(\mathcal{B})$ by the total spectral variation of variable k over the same band expresses the objective as a variance share. The maximizing

rotation is the leading eigenvector of the normalized matrix. When the target contains more than one variable, we maximize the equally weighted average of their normalized matrices. This normalization prevents a series with larger units or greater unconditional variance from determining the joint objective.

The main analysis sets $\mathcal{B} = [2\pi/40, 2\pi/6]$, corresponding to periodicities between 6 and 40 quarters. We extend the conventional 6–32-quarter window used by Angeletos et al. (2020) because the spectral density of hours in our sample peaks at a periodicity of approximately 36 quarters. Identification remains entirely in the frequency domain. The reported impulse responses are time-domain representations of rotations selected by the frequency-domain criterion; maximizing variation over periodicities from 6 to 40 quarters is not equivalent to maximizing forecast-error variance over horizons 6 through 40.

2.2.1 The main business-cycle and joint shocks

For each information set $s \in \{10, 12, 8\}$, the main business-cycle shock is the unit vector

$$q^{H,s} = \arg\max_{q'q=1} q' M_H^s(\mathcal{B}) q, \tag{4}$$

where H denotes hours worked and the superscript indexes the VAR. We orient each rotation so that hours rise on impact. Applying the same objective to the three systems allows us to ask whether the spectral contribution of the MBC shock changes when the VAR is given direct information about macro uncertainty and the credit spread. The twelve-variable system provides the closest comparison with Angeletos et al. (2020); the eight-variable system is the baseline for the subsequent joint and sequential exercises.

To locate the information underlying q^H, we also identify

$$q^J = \arg\max_{q'q=1} q' \left[\tfrac{1}{2}\widetilde{M}_{GZ}(\mathcal{B}) + \tfrac{1}{2}\widetilde{M}_U(\mathcal{B})\right] q, \tag{5}$$

where tildes denote normalization by each variable's total variation over the band, GZ is the credit spread, and U is macro uncertainty. Comparing the shock series, impulse responses, and hours spectra associated with q^H and q^J establishes whether the main cyclical movement in hours is also the dominant common movement in uncertainty and credit spreads. This comparison is descriptive: neither rotation is assigned a structural label at this stage. We conduct it in both the twelve- and eight-variable systems. In the notation that follows, q^H and q^J refer to the eight-variable baseline unless stated otherwise.

2.2.2 Separating uncertainty and credit conditions

The joint rotation cannot distinguish an uncertainty innovation from a credit innovation because the two observables move together. We therefore apply the sequential identification to the

eight-variable baseline, adapting the logic of Caldara et al. (2016) to the frequency-domain objective. The first rotation is

$$q^U = \arg\max_{q'q=1} q'\widetilde{M}_U(\mathcal{B})q. \tag{6}$$

We call this the uncertainty shock. In the figures, its sign is chosen so that uncertainty falls on impact and real activity expands; the response to an increase in uncertainty is the negative of the reported response. The second rotation is restricted to the orthogonal complement of q^U:

$$q^C = \arg\max_{q'q=1,\ q'q^U=0} q'\left[\tfrac{1}{2}\widetilde{M}_{GZ}(\mathcal{B}) + \tfrac{1}{2}\widetilde{M}_U(\mathcal{B})\right]q. \tag{7}$$

The uncertainty direction has already been removed from the admissible set, so q^C isolates the remaining common credit-spread and uncertainty variation. We orient this shock so that output rises on impact. Its economic interpretation is based on its full set of responses and on the auxiliary evidence below, not on the maximization criterion alone. In the appendix, we also report a rotation orthogonal to q^U that targets hours directly.

Orthogonality in (7) concerns the contemporaneous structural innovations. It does not impose a zero dynamic response of uncertainty to the second shock. The two shocks can therefore have different impact vectors and still generate responses in both the credit spread and uncertainty through the VAR's propagation mechanism.

2.2.3 Auxiliary variables

The S&P index, the excess bond premium, and the high-yield ratio are observed over different samples and are not included in the eight-variable VAR. For each auxiliary series z_t, we estimate

$$z_t = \alpha_z + \Gamma_{z,0}x_t + \Gamma_{z,1}x_{t-1} + \Gamma_{z,2}x_{t-2} + v_{z,t} \tag{8}$$

on the dates for which z_t is available. The response of z_t to the credit shock is then

$$\mathrm{IRF}_z(h) = \sum_{\ell=0}^{\min\{2,h\}} \Gamma_{z,\ell}\,\mathrm{IRF}_x(h-\ell; q^C). \tag{9}$$

This projection method follows the auxiliary-variable exercise in Basu et al. (2025). Posterior bands are computed by applying the projection to each retained VAR draw. Variance shares for an auxiliary variable include the residual variation in $v_{z,t}$ in the denominator, so the reported share is not mechanically inflated by the fitted component of (8).

2.3 The Spectral Density of Hours

Figure 1 provides the empirical starting point. The spectral density of U.S. hours worked has a pronounced peak at a periodicity of approximately 36 quarters. The peak appears in the

level series and remains visible across the high-pass filters summarized by the gray envelope. This pattern is consistent with the concentration of aggregate activity at extended business-cycle frequencies documented by Beaudry et al. (2020). The spectrum does not identify which disturbances generate the peak. We address that question by applying the same hours-based maximum-share criterion across the three information sets described above.

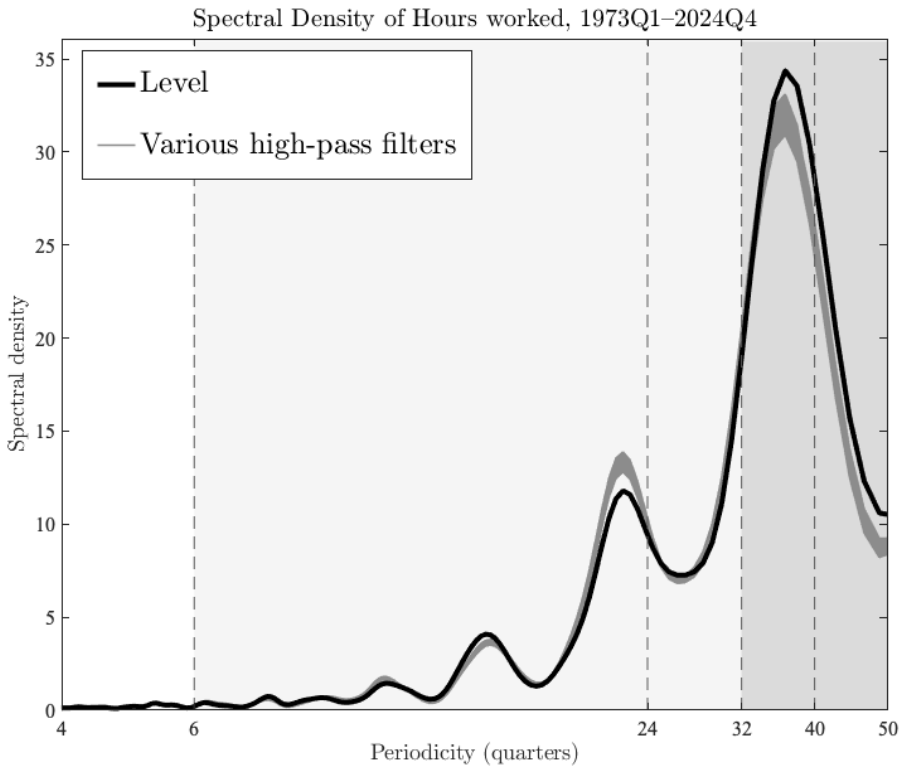


Figure 1: Spectral Density of Hours Worked

Notes. The figure reports the spectral density of log hours per capita from 1973:I to 2024:IV. The black line is the spectrum of the level series, and the gray region spans estimates obtained under alternative high-pass filters. Dashed vertical lines mark periodicities of 6, 24, 32, and 40 quarters.

2.4 Empirical Results

2.4.1 The information set and the hours spectrum

Figures 2 and 3 compare the frequency composition of the hours-based MBC shock across the three VARs. In the ten-variable system, shown in Figure 2, the spectral contribution of the first principal component rises toward lower frequencies but does not reproduce the peak in the data near 36 quarters. The second component adds variation at long periodicities without creating a peak at the observed location.

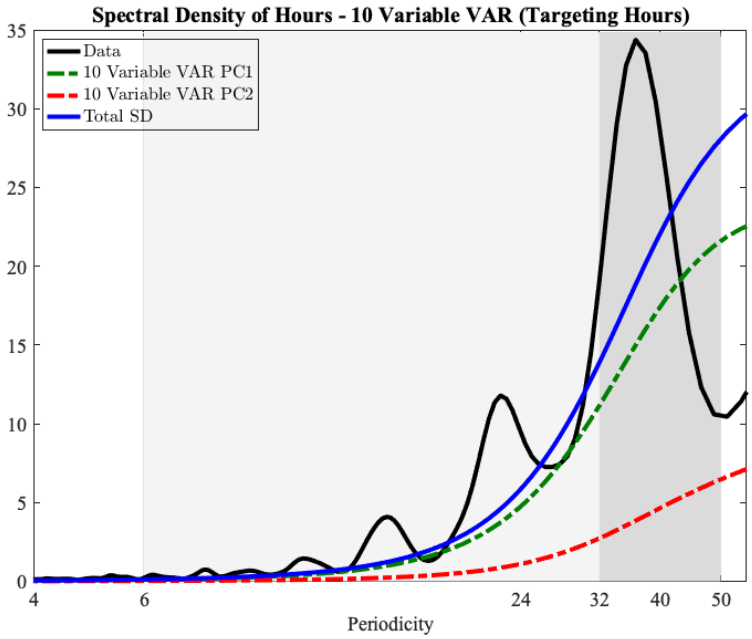


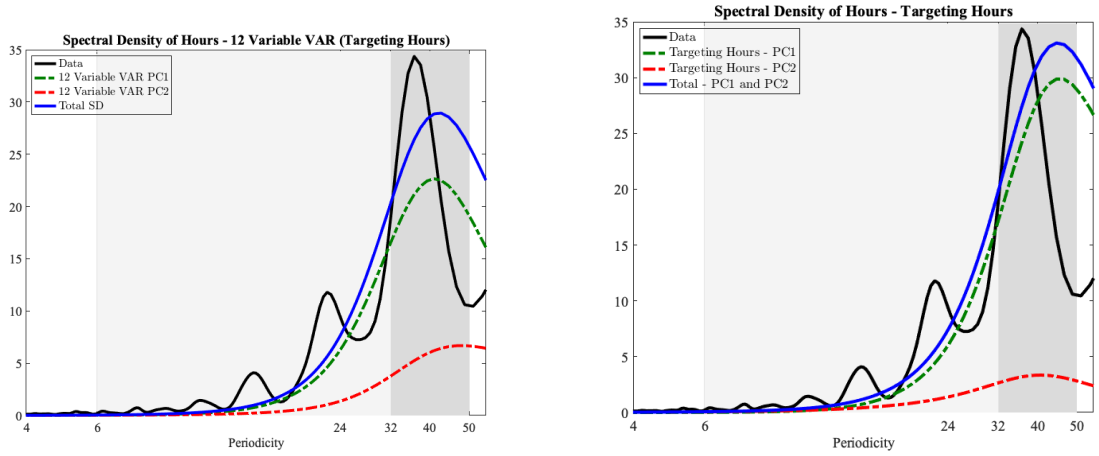
Figure 2: The Main Business-Cycle Shock in the Ten-Variable VAR

Notes. The black line is the spectral density of hours in the data. The green dashed line is the spectrum attributable to the first principal component of the hours objective, which is the MBC shock. The red dashed line is the spectrum attributable to the second component, and the blue line is the sum of the first two components. The shaded region marks periodicities from 32 to 50 quarters.

The result changes when macro uncertainty and the credit spread enter the VAR. Figure 3 shows that the first component develops a pronounced maximum in the same long business-cycle region as the data in both the twelve-variable and the more parsimonious eight-variable systems.

The comparison holds the target variable, sample, frequency band, prior, and maximum-share criterion fixed. What changes is the information included in the VAR. The spectral peak emerges only after the system includes direct measures of macro uncertainty and credit conditions. Its presence in both the twelve- and eight-variable systems also shows that the result does not depend on removing unemployment, the labor share, inflation, or labor productivity.

The variables added to the larger information sets suggest where the relevant movement is located. Figure 4 reports the hours spectrum associated with the rotation that jointly maximizes the normalized 6–40-quarter variation in macro uncertainty and the credit spread in the eight-variable VAR. Its first principal component produces a pronounced maximum in the same region as the hours-based MBC shock. Thus, both an hours-maximizing criterion and a joint uncertainty–credit-spread criterion select movements whose contributions to hours are concentrated around the empirical peak.



(a) Twelve-variable VAR

(b) Eight-variable VAR

Figure 3: The Main Business-Cycle Shock after Adding Macro Uncertainty and the Credit Spread

Notes. The black line is the spectral density of hours in the data. The green dashed line is the spectrum attributable to the first principal component of the hours objective, which is the MBC shock. The red dashed line is the spectrum attributable to the second component, and the blue line is the sum of the first two components. The shaded region marks periodicities from 32 to 50 quarters.

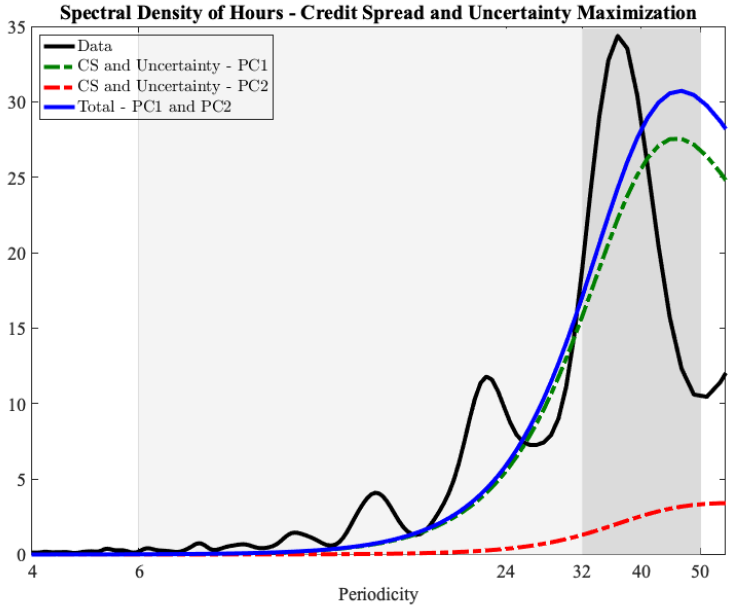


Figure 4: Hours Spectrum under the Joint Uncertainty–Credit-Spread Rotation

Notes. The black line is the spectral density of hours in the data. The green and red dashed lines are the spectra attributable to the first and second principal components of the joint credit-spread and macro-uncertainty objective. The blue line is their sum. The shaded region marks periodicities from 32 to 50 quarters.

2.4.2 Comparing the MBC and joint shocks

The similarity in their spectral contributions raises a natural question: whether the two criteria recover the same underlying movement. The answer is yes. In the eight-variable VAR, the estimated MBC and joint shock series have a correlation of 0.97. The correlation is 0.96 in the twelve-variable system. Their impulse responses are also nearly identical.

Figure 5 reports the responses to the hours-based MBC shock in the eight-variable VAR. Output, consumption, investment, and hours rise together, peak within the first ten quarters, and then return gradually toward zero. The shock accounts for 74 percent of the variation in hours over the target band. Its corresponding shares of output, consumption, and investment are 63, 49, and 62 percent. The credit spread and macro uncertainty fall sharply on impact and later overshoot. The shock explains 75 percent of the business-cycle variation in the spread and 51 percent of the variation in uncertainty.

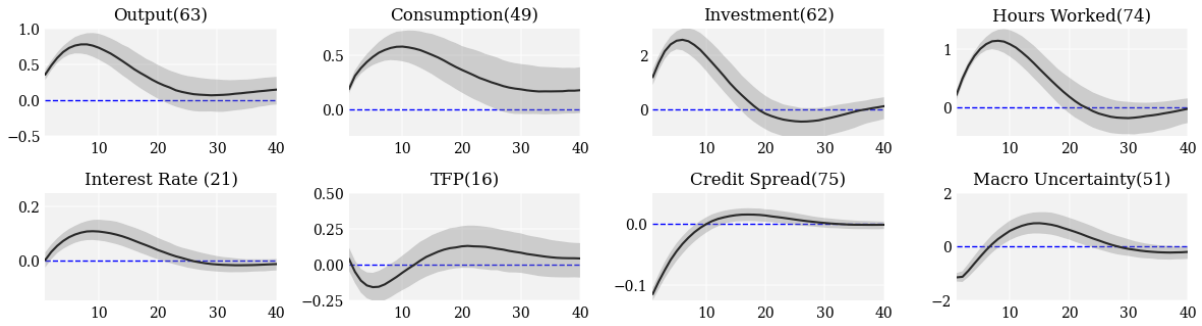


Figure 5: Impulse Responses to the Main Business-Cycle Shock

Notes. Solid lines are posterior median responses and shaded regions are 68 percent posterior intervals. The shock maximizes the variance share of hours over periodicities from 6 to 40 quarters and is normalized to raise hours on impact. Numbers in parentheses report the percentage of each variable's variation over that band attributed to the shock. The horizontal axis is in quarters.

Figure 6 shows the responses to the joint uncertainty–credit-spread shock. The four activity variables again rise together, while the credit spread and macro uncertainty fall on impact and subsequently overshoot. The timing and shape of these responses closely track those generated by the MBC shock. The joint shock also accounts for substantial business-cycle variation in hours, output, investment, the credit spread, and macro uncertainty. The agreement between the spectra, shock series, and impulse responses indicates that the movement responsible for most business-cycle variation in hours is closely associated with the dominant joint movement in uncertainty and credit spreads.

This correspondence does not mean that a single structural innovation drives both variables. Macro uncertainty and the credit spread move together, but each also contains independent variation. A maximum-share criterion can therefore combine distinct disturbances that propagate differently. We use the sequential restrictions in equations (6) and (7) to separate them.

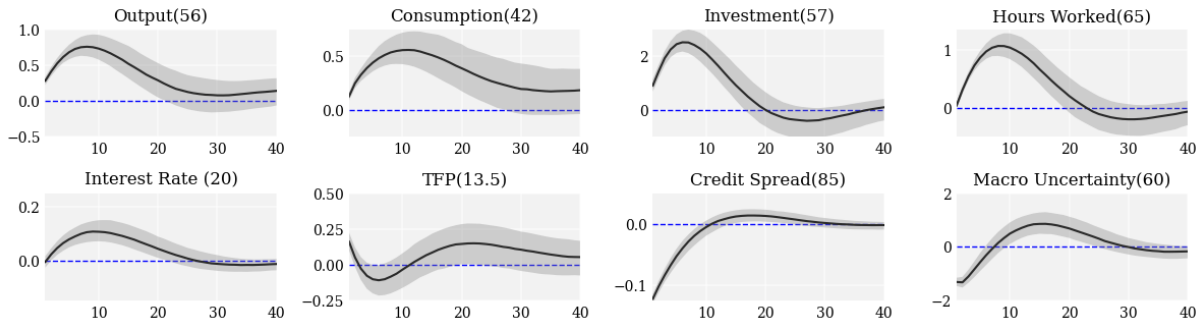


Figure 6: Impulse Responses to the Joint Uncertainty–Credit-Spread Shock

Notes. Solid lines are posterior median responses and shaded regions are 68 percent posterior intervals. The shock jointly maximizes the normalized variance shares of the credit spread and macro uncertainty over periodicities from 6 to 40 quarters. Numbers in parentheses report variance shares over that band. The horizontal axis is in quarters.

2.4.3 Separating uncertainty and credit conditions

Figure 7 reports the first rotation, with the sign normalized as a decline in uncertainty. Output, consumption, investment, and hours all expand. Their responses gradually return toward zero without a subsequent contraction. Reversing the sign gives the uncertainty experiment: an increase in macro uncertainty contracts all four activity variables, after which they recover toward steady state without an economically meaningful boom. The shock accounts for 78 percent of the target-band variation in uncertainty and 40 percent of the variation in the credit spread.

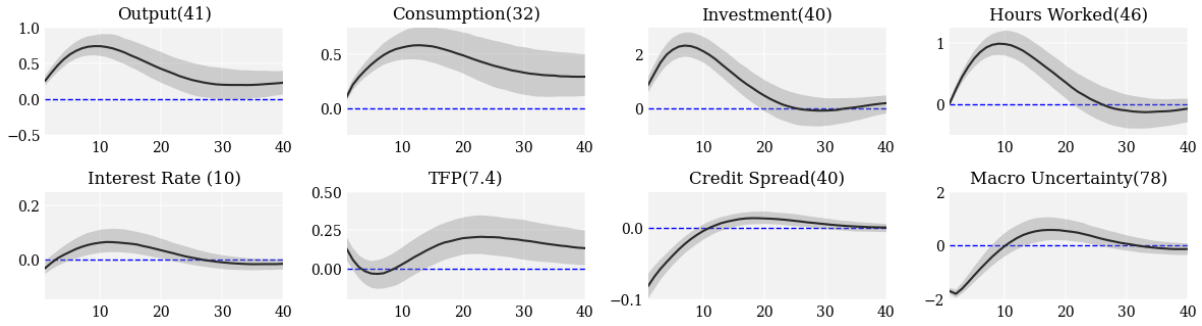


Figure 7: Impulse Responses to the Uncertainty Shock

Notes. Solid lines are posterior median responses and shaded regions are 68 percent posterior intervals. The shock maximizes the variance share of macro uncertainty over periodicities from 6 to 40 quarters. Its sign is normalized so that uncertainty falls on impact. Numbers in parentheses report variance shares over the target band. The horizontal axis is in quarters.

The second rotation, shown in Figure 8, is orthogonal to the uncertainty shock and maximizes the remaining joint variation in the credit spread and macro uncertainty. A sharp decline in the credit spread is followed by a broad expansion. Output, consumption, investment, and hours rise together for roughly 12–15 quarters and then change sign. The contraction is clearest in

output, investment, and hours; the decline in consumption is weaker and less precisely estimated. Activity reaches its trough after approximately 20–25 quarters and then gradually recovers.

The shock explains 53 percent of the business-cycle variation in the credit spread but only 12 percent of the variation in macro uncertainty. Uncertainty can nevertheless respond after impact because the orthogonality restriction applies to the contemporaneous innovations, not to the subsequent propagation of the VAR. We refer to this second rotation as the credit-related shock. The label describes the evidence developed below; it is not imposed by the identifying restriction.

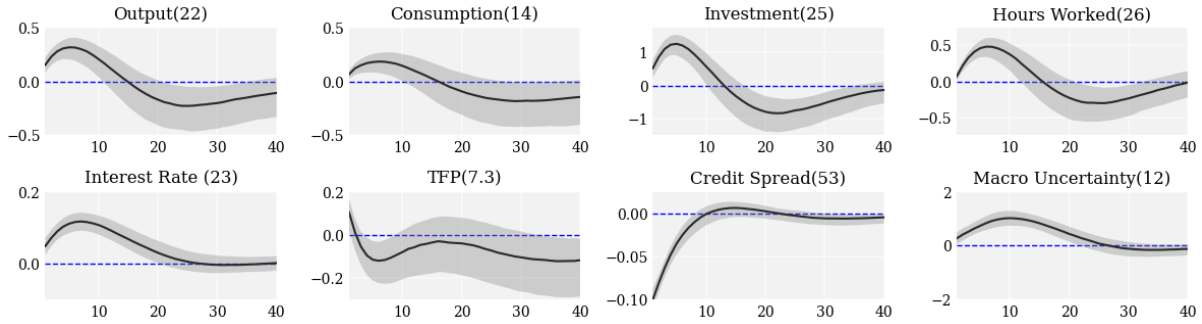


Figure 8: Impulse Responses to the Credit-Related Shock

Notes. Solid lines are posterior median responses and shaded regions are 68 percent posterior intervals. The shock is orthogonal to the uncertainty shock and maximizes the remaining joint variation in the credit spread and macro uncertainty over periodicities from 6 to 40 quarters. It is normalized to lower the credit spread on impact. Numbers in parentheses report variance shares over the target band. The horizontal axis is in quarters.

Figure 9 identifies which component accounts for the peak in hours. The uncertainty shock contributes variation at long periodicities, but it does not reproduce the pronounced maximum in the data. The credit-related component rises sharply around the observed peak and reaches its maximum near 40 quarters. The spectral concentration recovered by the MBC and joint rotations therefore comes primarily from the credit-related shock, whereas their broad contraction-and-recovery component comes from uncertainty.

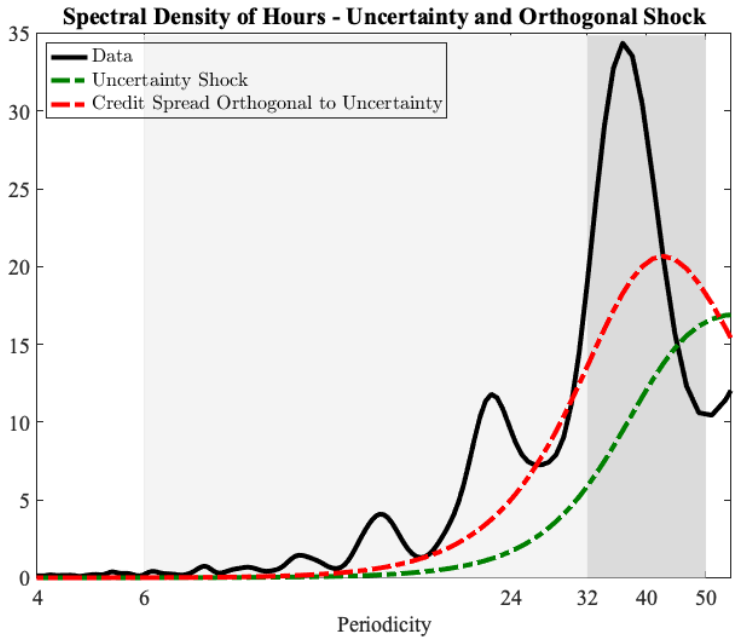


Figure 9: Hours Spectra of the Uncertainty and Credit-Related Shocks

Notes. The black line is the spectral density of hours in the data. The green dashed line is the spectrum attributable to the uncertainty shock. The red dashed line is the spectrum attributable to the credit-related shock orthogonal to uncertainty. The shaded region marks periodicities from 32 to 50 quarters.

2.4.4 Financial evidence on the credit-related shock

Figure 10 uses the auxiliary projections to examine the financial content of the second rotation. The real S&P index rises sharply and then falls below its initial level, reproducing the boom–bust pattern in real activity. The excess bond premium declines on impact, subsequently becomes positive, and then returns toward zero. The shock explains 54 percent of the variation in the excess bond premium captured by the auxiliary projection. The high-yield issuance share also rises during the expansion, later turns negative, and recovers gradually.

These responses support the credit-related interpretation. Compressed risk compensation, greater risky issuance, rising equity values, and their subsequent reversal connect the shock to the credit-market evidence of Greenwood and Hanson (2013) and López-Salido et al. (2017). They are not used to identify the shock, nor do they require investor sentiment to be its primitive source.

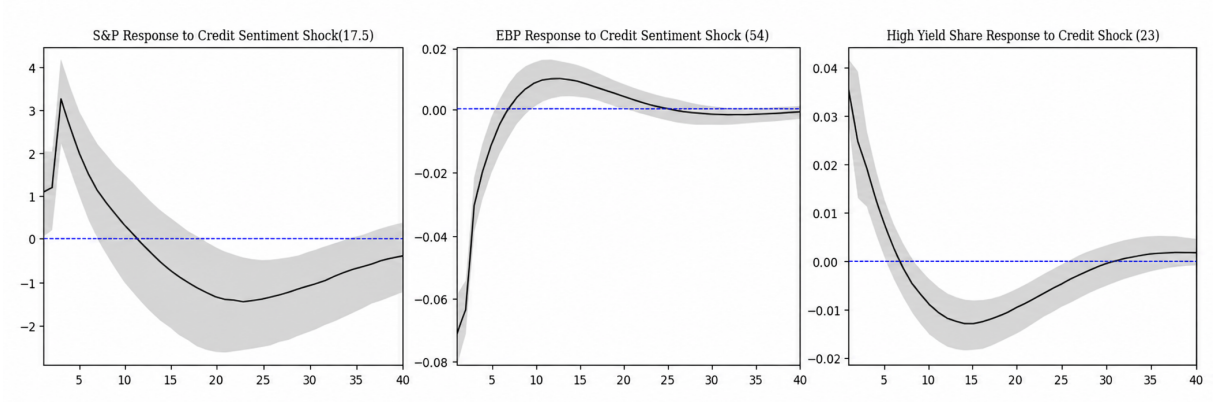


Figure 10: Financial Responses to the Credit-Related Shock

Notes. Panels report the projected responses of the real per-capita S&P index, the excess bond premium, and the high-yield issuance share to the shock in Figure 8. Solid lines are posterior medians and shaded regions are 68 percent posterior intervals. Numbers in parentheses report the variance shares attributed to the shock. The horizontal axis is in quarters.

2.5 Discussion: Implications for the Model

The empirical results place two requirements on the model. The credit-related shock produces a boom followed by a contraction and accounts for the peak in the hours spectrum. An increase in macro uncertainty instead contracts output, consumption, investment, and hours, after which activity gradually returns to steady state. The model must therefore generate comovement among all four quantities after both shocks while allowing their propagation to differ. In particular, boom–bust dynamics cannot be an automatic response to every disturbance.

The empirical exercise identifies these propagation patterns, but it does not identify the primitive disturbance behind the credit-related shock. In the model, we represent that shock as a favorable but noisy signal about future productivity. The signal initially raises hiring, investment, borrowing, and collateral values. When agents later learn that the signal was noise, expected productivity falls while firms retain the debt accumulated during the expansion, causing credit conditions and activity to reverse. Macro uncertainty operates through the same financing environment, but through a different margin: it raises the expected scarcity of borrowing capacity and produces a contraction without a preceding credit boom. Nominal wage rigidity allows the resulting shifts in labor demand to generate the comovement observed in the data. The next section develops this common financial mechanism and evaluates whether it can reproduce both empirical responses.

3 The Model

3.1 Overview

The model is designed to account for two empirical responses documented in the previous section: a drop in credit spread (or excess bond premia) led to expansion that later reversed to contraction of the economy, and a contraction following an increase in macro uncertainty. In both cases, output, consumption, investment, and hours move in the same direction. This co-movement is a demanding restriction for one-sector models (Basu and Bundick, 2017; Angeletos et al., 2020). Since both innovations generate similar responses in the credit spread, the analysis asks whether a common financing mechanism can generate both responses.

The economy contains infinitely lived households, competitive firms, and a central bank. Households consume, supply differentiated labor, and hold one-period nominal bonds. Firms own capital, hire labor, make investment, issue debt, and make equity payouts. The central bank sets the short nominal rate with a Taylor rule. Goods prices are flexible; nominal wages adjust slowly.

Households deal with wage rigidity, while firms navigate financial constraints. The financial setup rests on three key elements. First, firms need to borrow during each period to cover labor wages before they receive any revenue, referred to as working capital loans. However, a constraint limits how much firms can borrow for this purpose, tying it to the value of capital that can be pledged, after accounting for outstanding debt. Second, if firms adjust their equity payouts away from their long-term target, they face convex adjustment costs. Lastly, the model incorporates tax benefits for debt and costs of adjusting payouts, as discussed by Jermann and Quadrini (2012), along with a changing enforcement limit on wage bills and its uncertainty, following Chatterjee et al. (2025). These costs keep firms from easily using equity to make up for drops in borrowing capacity. As a result, credit conditions end up influencing hiring decisions through the value firms place on their internal funds.

One further feature links investment to borrowing capacity. Convex investment adjustment costs make the value of installed capital respond to the current investment rate. Because the borrowing limit depends on this value, investment can affect borrowing capacity immediately even though the capital stock is predetermined. Together with the nominal wage rigidity described above, this link allows changes in credit conditions to move output, consumption, investment, and hours in the same direction.

The model is perturbed by productivity surprises, noisy signals about future productivity, and innovations to enforcement and its volatility. A signal arrives with an observed magnitude and horizon, but agents do not initially know whether it contains genuine news or noise. A favorable signal can relax the enforcement constraint during anticipation. If it is later revealed to be noise, the constraint can bind again as the expected productivity gain disappears while previously issued debt remains. The constraint's status is determined in equilibrium. The uncertainty experiment raises the dispersion of enforcement innovations, following Chatterjee et al. (2025).

We solve and estimate the model using a deep equilibrium network (DEQN) (Azinovic et al., 2022; Scheidegger, 2026).

3.2 Households

Preferences and budget constraint. There is a continuum of households indexed by $j \in [0,1]$, each supplying a differentiated variety of labor. Household j chooses consumption, bond holdings and its own nominal wage to maximise

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{(c_t(j) - \iota c_{t-1}(j))^{1-\sigma} - 1}{1-\sigma} - \alpha \frac{n_t(j)^{1+\nu}}{1+\nu} \right], \tag{10}$$

where $\beta \in (0,1)$ is the discount factor, $\sigma > 0$ governs the curvature of the consumption aggregate, $\nu > 0$ is the inverse Frisch elasticity, and $\iota \in [0,1)$ is the degree of habit persistence. Habit is internal: the household recognizes that raising consumption today raises the reference stock against which its own utility is evaluated tomorrow, and internalizes that effect. This differs from the external specification in the estimated model of Jermann and Quadrini (2012) and adds a forward-looking term to the marginal utility of wealth, as recorded below.

Consumption, investment, dividends and taxes are real; bonds and wages are nominal and are deflated by the current price level. The budget constraint is

$$c_t(j) + \frac{B_{t+1}(j)}{P_t(1+i_t)} = \big(1 - \Theta_{w,t}(j)\big) \frac{W_t(j)}{P_t} n_t(j) + \frac{B_t(j)}{P_t} + d_t - T_t, \tag{11}$$

where $B_t(j)$ is the nominal face value contracted previously and due at t, i_t is the net nominal interest rate on a one-period bond, d_t is the real dividend received from firm ownership, T_t is a real lump-sum tax, and $\Theta_{w,t}(j)$ is the wage adjustment cost defined below. There is no inflation-indexed bond. Because all debt is nominal, unanticipated inflation revalues outstanding liabilities in real terms, so the model carries a debt-deflation channel absent from purely real formulations of the same financing block.

labor varieties and wage setting. Household j sells its variety to a competitive labor packer that assembles the aggregate input used by firms. The packer's cost minimization delivers the downward-sloping demand curve

$$n_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\eta_w} n_t, \tag{12}$$

where W_t is the aggregate nominal wage index, n_t is aggregate labor demand, and $\eta_w > 1$ is the elasticity of substitution across varieties, which implies a desired wage markup $\mu_w \equiv \eta_w/(\eta_w - 1)$.

Households are monopolistic suppliers of their own variety and face a real resource cost, borne

as a share of their own wage bill, when they change their posted nominal wage:

$$\Theta_{w,t}(j) = \frac{\phi_w}{2}\big(\Pi_{w,t}(j) - 1\big)^2, \qquad \Pi_{w,t}(j) \equiv \frac{W_t(j)}{W_{t-1}(j)}, \tag{13}$$

with $\phi_w \geq 0$ governing the degree of stickiness. Two features of (13) matter. The cost applies to *nominal* wage growth; a cost in real wage growth would not create a nominal friction. Symmetry and convexity imply a common wage, so wage dispersion is not an additional state. The specification yields a forward-looking wage Phillips curve. Born and Pfeifer (2020) provide the mapping from ϕ_w to the average contract duration in an equivalent staggered-contract economy, which we use to assess the calibration in Section 3.8.

Optimality conditions. Let $\Lambda_t(j)$ denote the multiplier on household j's budget constraint and define its flow marginal felicity as $X_t(j) \equiv [c_t(j) - \iota c_{t-1}(j)]^{-\sigma}$. In a symmetric equilibrium, consumption, bond holdings, and their associated multipliers are common across households. We therefore drop the j index from these objects from this point on. Because habit is internal, the marginal utility of wealth nets off the future cost of a higher habit stock,

$$\Lambda_t = X_t - \iota\beta\,\mathbb{E}_t[X_{t+1}]\,, \tag{14}$$

so that $\Lambda_t < X_t$ whenever $\iota > 0$: the household perceives itself as gaining less from an extra unit of consumption than the flow term alone would suggest. The stochastic discount factor over real payoffs is $m_{t,t+1} \equiv \beta\Lambda_{t+1}/\Lambda_t$, and the pricing condition for the nominal bond is

$$1 = (1 + i_t)\,\mathbb{E}_t\left[\frac{m_{t,t+1}}{\Pi_{t+1}}\right], \qquad \Pi_{t+1} \equiv \frac{P_{t+1}}{P_t}. \tag{15}$$

Equation (15) holds exactly under uncertainty and is the only intertemporal asset-pricing condition in the model.

Imposing symmetry across households and defining the marginal rate of substitution $\text{mrs}_t \equiv \alpha n_t^{\nu}/\Lambda_t$ and the real wage $w_t \equiv W_t/P_t$, the optimal wage satisfies

$$\phi_w\big(\Pi_{w,t} - 1\big)\Pi_{w,t} = (1-\eta_w)\big(1-\Theta_{w,t}\big) + \eta_w\frac{\text{mrs}_t}{w_t} + \mathbb{E}_t\left[m_{t,t+1}\,\phi_w\big(\Pi_{w,t+1} - 1\big)\Pi_{w,t+1}^2\,\frac{n_{t+1}}{n_t}\,\frac{1}{\Pi_{t+1}}\right]. \tag{16}$$

Equation (16) is the nonlinear wage Phillips curve. When $\phi_w = 0$ it collapses to the static markup rule $w_t = \mu_w\,\text{mrs}_t$ of Erceg et al. (2000), which we use as a numerical check on the solution. When $\phi_w > 0$, the real wage need not track the marginal rate of substitution period by period, allowing labor demand to determine hours over the horizons of interest.

3.3 Firms

The firm sector contains the model's financial friction.

Technology and capital accumulation. Firms are perfectly competitive and produce with a constant-returns Cobb–Douglas technology,

$$y_t = z_t k_t^{\theta} n_t^{1-\theta}, \tag{17}$$

where z_t is aggregate productivity, k_t is the predetermined capital stock and θ is the capital share. There is no price rigidity and no markup on the goods side; every nominal rigidity in the model resides in wages. Capital accumulates subject to a convex cost in the level of the investment rate,

$$k_{t+1} = (1-\delta)k_t + I_t - \frac{\phi}{2}\left(\frac{I_t}{k_t} - \delta\right)^2 k_t, \tag{18}$$

following Chatterjee et al. (2025). The cost is specified in terms of I_t/k_t rather than in its growth rate, which keeps the firm's individual state two-dimensional and, as Section 3.7 shows, delivers a transparent closed-form link between the investment rate and the value of collateral.

Financing. Firms use two instruments. Intertemporal debt is nominal and has face value B_{t+1}, and is preferred to equity because interest is tax-deductible: given the policy rate i_t, the effective gross rate paid by the firm is

$$R_t = 1 + i_t(1-\tau), \tag{19}$$

where τ is the deductibility parameter. The resulting subsidy is financed by the lump-sum tax $T_t = B_{t+1}/(P_t R_t) - B_{t+1}/\big(P_t(1+i_t)\big)$ levied on households, so the tax shield is a pure financing wedge and not a transfer of resources into or out of the economy. In addition, firms take an intraperiod working-capital loan to cover the wage bill, which is paid before revenues arrive and repaid within the period at no interest.

The equity payout is costly to adjust. Given a payout d_t, the resource cost borne by the firm is $d_t + \kappa(d_t - \bar{d})^2$, where $\bar{d}$ is the long-run payout target and $\kappa \geq 0$. Following Jermann and Quadrini (2012), this should not be read as a literal pecuniary cost but as a parsimonious way of capturing the speed with which firms can rebalance between debt and equity when financial conditions change; it is consistent both with the dividend-smoothing behaviour documented since Lintner (1956) and with increasing marginal costs of equity issuance. The parameter κ turns out to be the pivotal one: it is what prevents a change in pledgeable value from being absorbed silently by the financial structure. Collecting these elements, the firm's flow of funds in real terms is

$$y_t + \frac{B_{t+1}}{P_t R_t} = d_t + \kappa\left(d_t - \bar{d}\right)^2 + w_t n_t + I_t + \frac{B_t}{P_t}. \tag{20}$$

The enforcement constraint. Debt contracts are imperfectly enforceable. Lenders extend the working-capital loan only up to the value they expect to recover, which is a fraction ζ_t of the collateral remaining after the first claim of intertemporal debt-holders has been met:

$$\zeta_t \left[Q_t k_t - \frac{B_{t+1}}{P_t(1+i_t)} \right] \geq w_t n_t, \tag{21}$$

with associated multiplier $\mu_t \geq 0$ and complementary slackness. Here Q_t is the price at which a unit of installed capital is valued in the contract, measured in goods, and k_t is the stock available when the contract is signed. Three features of (21) deserve comment. First, the constraint is written on the wage bill, not on output, so the credit wedge attaches directly to the hiring margin. Second, it is the predetermined stock k_t that serves as collateral, not the choice k_{t+1}: the firm cannot relax today's borrowing limit by investing more today, and the value of extra capital in loosening future constraints therefore enters the capital Euler equation through the envelope condition rather than as a contemporaneous subsidy to investment. Third, the current market value of newly issued intertemporal debt is deducted, so leverage and pledge-ability compete for the same collateral. The parameter ζ_t is common across firms and stochastic; its innovations are the credit shocks of the model, and we interpret movements in ζ_t as changes in the recover-ability of collateral rather than as shifts in investor beliefs.

Writing the unused borrowing capacity as

$$S_t \equiv \zeta_t \left[Q_t k_t - \frac{B_{t+1}}{P_t(1+i_t)} \right] - w_t n_t, \qquad S_t \geq 0, \quad \mu_t \geq 0, \quad \mu_t S_t = 0, \tag{22}$$

makes the two regimes explicit. When $S_t > 0$, the constraint is slack and $\mu_t = 0$; when $\mu_t > 0$, it binds and $S_t = 0$. Binding in the deterministic steady state does not impose binding along stochastic paths.

Collateral valuation. It remains to say how lenders price the pledged stock. Let q_t be the multiplier on (18), which measures the marginal value of installed capital, and λ_t the multiplier on (20), which measures the marginal value of one good. Their ratio is the marginal value of capital expressed in goods, and we assume that lenders price collateral at this value:

$$Q_t = \frac{q_t}{\lambda_t}. \tag{23}$$

Equation (23) is a valuation rule imposed on the symmetric equilibrium, not an identity derived from an explicit liquidation market, and it does not make Q_t exogenous: the price is determined jointly with investment, payouts and the remaining equilibrium objects. An individual firm takes Q_t as given, and (23) is applied only after the firm's optimality conditions have been derived, so that no firm treats the aggregate valuation rule as a direct dependence of the contract price on its own choices. When payouts adjust freely, $\kappa = 0$ implies $\lambda_t \equiv 1$ and hence $Q_t = q_t$.

The firm's problem. The individual state is (k_t, B_t), with aggregate states suppressed in the time subscript of the value function. The firm chooses $\{d_t, n_t, I_t, k_{t+1}, B_{t+1}\}$ to maximise cum-dividend value

$$V_t(k_t, B_t) = d_t + \mathbb{E}_t[m_{t,t+1} V_{t+1}(k_{t+1}, B_{t+1})] \tag{24}$$

subject to (17), (18), (20) and (21), discounting with the owners' stochastic discount factor and taking w_t, i_t, P_t and Q_t as given.

Optimality conditions. The payout condition is $\lambda_t = \left[1 + 2\kappa(d_t - \bar{d})\right]^{-1}$, so λ_t falls below unity exactly when the firm pays out more than its target. The labor choice yields the central static condition of the model,

$$w_t \left(1 + \frac{\mu_t}{\lambda_t}\right) = (1-\theta)\frac{y_t}{n_t}. \tag{25}$$

labor is hired to the point where its marginal product equals the wage augmented by a financing wedge, μ_t/λ_t, which is the shadow cost of the collateral constraint expressed per unit of internal funds. A higher wedge shifts labor demand inward at a given wage and a lower wedge does the reverse. Equation (25) is the only condition in which the shadow cost of the constraint bears directly on a contemporaneous real allocation, although, as Section 3.7 shows, what moves the wedge is determined elsewhere in the system. It is the model's counterpart to the labor wedge in Jermann and Quadrini (2012) and Chatterjee et al. (2025), generalized here to account for the marginal value of internal funds.

The investment condition is $\lambda_t = q_t\left[1 - \phi(I_t/k_t - \delta)\right]$. Capital and debt satisfy

$$q_t = \mathbb{E}_t\left[m_{t,t+1}\left(q_{t+1}\Xi_{t+1} + \lambda_{t+1}\theta\frac{y_{t+1}}{k_{t+1}} + \underbrace{\mu_{t+1}\zeta_{t+1}Q_{t+1}}_{\text{collateral payoff}}\right)\right], \tag{26}$$

$$\mathbb{E}_t\left[\frac{m_{t,t+1}\lambda_{t+1}}{\Pi_{t+1}}\right] = \frac{\lambda_t}{R_t} - \frac{\mu_t\zeta_t}{1+i_t}, \tag{27}$$

where $\Xi_{t+1} \equiv (1-\delta) + \phi\left(\frac{I_{t+1}}{k_{t+1}} - \delta\right)\frac{I_{t+1}}{k_{t+1}} - \frac{\phi}{2}\left(\frac{I_{t+1}}{k_{t+1}} - \delta\right)^2$ collects the standard user-cost terms and we have used $\varphi_t = \lambda_t$, the optimality condition for output.

Two terms in (26)–(27) are specific to this environment. In (26) the marginal unit of capital pays not only its marginal product and its resale value but also a collateral dividend $\mu_{t+1}\zeta_{t+1}Q_{t+1}$: capital carried into the next period has an additional payoff in states where it relaxes a binding financing constraint. This payoff is zero in slack states. In (27) the marginal cost of debt is raised by $\mu_t\zeta_t/(1+i_t)$, because a unit of new borrowing consumes pledgeable collateral one for one. The tax shield pushes the firm towards more leverage and the enforcement constraint pushes back; Section 3.7 shows that this tension makes the constraint bind at the deterministic steady state.

3.4 The Nominal Block and Monetary Policy

Goods prices are flexible, so there is no price Phillips curve. Price inflation is instead recovered from the definition of the real wage: since $w_t \equiv W_t/P_t$,

$$\Pi_t = \Pi_{w,t} \cdot \frac{w_{t-1}}{w_t}. \tag{28}$$

Equation (28) is an accounting identity and holds regardless of any stickiness assumption. It links the real wage determined by (25) and nominal wage inflation determined by (16) to the inflation rate in the policy rule. Price inflation falls when the real wage rises faster than the nominal wage.

The central bank follows a rule in current inflation and the level of output,

$$1 + i_t = (1 + i_{ss}) \left(\frac{\Pi_t}{\Pi^*} \right)^{\phi_\pi} \left(\frac{y_t}{y_{ss}} \right)^{\phi_y}, \tag{29}$$

with target $\Pi^* = 1$ and $1+i_{ss} = 1/\beta$. Output is measured relative to its deterministic steady state rather than to a time-varying efficient benchmark. We depart from the simplest inflation-forecast rule of Christiano et al. (2008) by responding to current rather than expected inflation and by including an output term; interest-rate smoothing and a response to credit growth are natural extensions that we do not pursue here. Choosing $\phi_\pi > 1$ satisfies the Taylor principle, but determinacy is a property of the full system and is verified numerically at the calibration used below.

The distinction between the policy target and the sticky nominal variable is substantive. Stabilising wage inflation would move the economy toward its flexible-wage allocation (Erceg et al., 2000). By targeting price inflation instead, rule (29) leaves room for wage rigidity to affect real allocations.

3.5 Exogenous Processes

Technology and noisy signals. Following Lorenzoni (2009), agents receive an imperfect signal about future productivity. The signal may contain genuine news, in which case future TFP changes, or noise, in which case it does not. Agents observe the signal's magnitude and horizon when it arrives, but they cannot initially distinguish between these two possibilities. They therefore respond to a favorable signal before knowing whether the anticipated productivity improvement will materialize.

Chahrour and Jurado (2018) show that news and noise can have equivalent representations once the underlying information structure is specified. We use this equivalence to implement the signal with a pending productivity announcement. Confirmation at maturity delivers the announced change; nonconfirmation reveals that the initial signal was noise. Equivalently, the noise realization can be represented as favorable news followed at maturity by an unforeseen innovation that offsets

it. The implementation below uses the former representation, so productivity never changes when the signal proves to be noise.

Aggregate productivity follows

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_t^0 + D_t, \qquad \varepsilon_t^0 \sim N(0, \sigma_0^2), \tag{30}$$

where D_t is the productivity change associated with a previously observed signal. At most one signal is pending. When no signal is pending, a new signal arrives with probability λ_{news}. Its magnitude $a \sim N(0, \sigma_0^2)$ and horizon $H \sim 8 + \text{Binomial}(8, p_H)$ are drawn once and observed on arrival. Signal magnitudes and unanticipated productivity innovations therefore share the variance σ_0^2. The signal state consists of its remaining horizon χ_t and magnitude a_t. A zero countdown means that no signal is pending. At maturity, an independent confirmation draw F_t determines delivery:

$$D_t = \mathbf{1}\{\chi_{t-1} = 1\} F_t a_{t-1}, \qquad F_t \sim \text{Bernoulli}(\pi_{\text{confirm}}). \tag{31}$$

Agents know the confirmation probability, but not the outcome before maturity. Thus, their expectations are rational even when a signal is ultimately revealed to be noise. After resolution, a new signal may arrive in the same period. Otherwise, the magnitude is retained and the countdown falls by one. Appendix B.8 states the transition and timing convention.

This process retains the distinction between anticipated and unanticipated productivity movements in Christiano et al. (2008, 2010), while allowing an observed signal to contain either genuine news or noise. The arrival rate, horizon distribution, and confirmation probability are estimated rather than fixing the timing or likelihood of resolution.

The noise-shock experiment. The boom–bust experiment introduces a favorable signal with magnitude $\bar{a}$ and horizon H at $t = 1$ and conditions on nonconfirmation at $t = 1 + H$. During anticipation, agents assign probability π_{confirm} to a genuine future productivity increase. At maturity, $F_{1+H} = 0$ reveals that the signal was noise, sets $D_{1+H} = 0$, and removes the pending signal. With all other technology innovations set to zero and no further signals, productivity remains constant throughout the experiment. The genuine-news benchmark conditions on $F_{1+H} = 1$ for the same initial signal. Its anticipation path is identical because agents do not know the eventual outcome before maturity.

The empirical credit-related shock generates a boom followed by a bust. The noise experiment provides a rational-expectations interpretation of that sequence: favorable information initially supports borrowing and activity, but the expansion reverses when agents learn that the signal contained no genuine news. The signal and its resolution are exogenous; the movements in borrowing capacity and the occasionally binding constraint are equilibrium outcomes.

Credit conditions and macro uncertainty. The enforcement parameter follows a level autoregression whose innovation variance is itself stochastic,

$$\zeta_t - \zeta_{ss} = \rho_\zeta \left(\zeta_{t-1} - \zeta_{ss}\right) + \exp(h_t/2)\, \varepsilon_{\zeta,t}, \tag{32}$$

$$h_t = (1-\rho_h)\bar{h} + \rho_h h_{t-1} + \sigma_h\, \varepsilon_{h,t}, \tag{33}$$

The untruncated specification follows Chatterjee et al. (2025). To keep enforcement positive in the global solution, define $\bar{\zeta}_t = (1-\rho_\zeta)\zeta_{ss} + \rho_\zeta \zeta_{t-1}$ and draw $\varepsilon_{\zeta,t}$ from a standard normal conditional on $\varepsilon_{\zeta,t} > -\bar{\zeta}_t/\exp(h_t/2)$. The volatility innovation satisfies $\varepsilon_{h,t} \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$. The two innovations are generated from independent underlying draws, also independent of the technology and signal draws. The unconditional mean of h_t is $\bar{h}$, so $\exp(\bar{h}/2)$ is the baseline scale of the enforcement innovation. An innovation to h_t changes that scale without changing the autoregressive location $\bar{\zeta}_t$. Positive-support truncation nevertheless raises the conditional mean above $\bar{\zeta}_t$ by a scale-dependent amount. We therefore interpret $\varepsilon_{h,t}$ as a macro-uncertainty shock, while retaining this truncation effect in both expectations and simulations. Uncertainty in the enforcement parameter represents macro uncertainty, measured empirically using the index of Jurado et al. (2015).

3.6 Equilibrium and Solution

A symmetric competitive equilibrium is a set of processes for the quantities, prices and multipliers

$$\{c_t, n_t, y_t, I_t, k_{t+1}, d_t, B_{t+1}\}, \qquad \{w_t, \Pi_t, \Pi_{w,t}, i_t, Q_t\}, \qquad \{\Lambda_t, \lambda_t, q_t, \mu_t\},$$

such that households and firms optimise, the enforcement constraint holds with complementary slackness, the policy rule (29) is satisfied, and markets clear, given the exogenous processes (30)–(33) and an initial state.

Goods market clearing is not an independent restriction. Substituting the tax rule and (20) into the household budget constraint (11), all nominal debt flows and the financing subsidy cancel and Walras' law delivers

$$y_t = c_t + I_t + \kappa \left(d_t - \bar{d}\right)^2 + \Theta_{w,t} w_t n_t. \tag{34}$$

The cancellation of nominal terms in (34) is an accounting result and does not make the nominal block irrelevant: inflation still revalues outstanding debt in (20), and the policy rate still enters the pledgeability of collateral in (21) and the cost of debt in (27).

We approximate the state-dependent decision rules with a DEQN (Azinovic et al., 2022), following the implementation described by Scheidegger (2026, Chapters 2 and 10). The state is $s_t = (k_t, b_t, c_{t-1}, w_{t-1}, \log z_t, \zeta_t, h_t, \chi_t, a_t)$, where $b_t = B_t/P_{t-1}$, so current repayment is b_t/Π_t. The network returns six coordinates for hours, the real wage, inflation, household marginal utility,

total payout outlay, and collateral slackness. The remaining allocation is recovered from the model identities. In particular, a signed coordinate ψ_t determines

$$S_t = \bar{S}\max\{\psi_t, 0\}, \qquad \mu_t = \bar{\mu}\max\{-\psi_t, 0\}, \tag{35}$$

with positive numerical scales $\bar{S}$ and $\bar{\mu}$. Thus (22) holds exactly, and the network can switch between regimes without imposing a binding calendar. Proposition 1 applies only to the deterministic steady state.

Training minimises the scaled residuals of internal habit, the household Euler equation, wage setting, labor demand, and the capital and debt Euler equations. Expectations use the same network at endogenous next-period states, with discrete news outcomes enumerated and continuous innovations integrated by quadrature. Perturbation and DynareOBC calculations serve as starting calculations and benchmarks; final estimation uses the DEQN decision rules, not a pruned third-order state-space system. Appendix B.9 gives the numerical procedure.

3.7 Inspecting the Mechanism

Three results organise the mechanism. The tax advantage of debt makes the enforcement constraint bind at the deterministic steady state. The investment adjustment cost makes collateral value vary with the investment rate. The payout adjustment cost allows that variation to affect the financing wedge. We establish these results first and then use them to trace the two experiments.

When does the constraint bind? At the deterministic steady state the payout equals its target, so $\lambda_{ss} = 1$, and the investment rate equals δ, so $q_{ss} = 1$ as well. The following result is the counterpart of Proposition 1 in Jermann and Quadrini (2012).

Proposition 1. *At the deterministic steady state with inflation target $\Pi^* = 1$, the multiplier on the enforcement constraint is strictly positive if and only if $\tau > 0$. Hence the constraint is strictly binding whenever $\tau > 0$, and the steady-state financing wedge is*

$$\frac{\mu_{ss}}{\lambda_{ss}} = \frac{\tau(1-\beta)}{\zeta_{ss}\,(1-\tau+\tau\beta)}. \tag{36}$$

Proof. With $\Pi^* = 1$, (15) gives $1 + i_{ss} = 1/\beta$, so $R_{ss} = 1 + (1-\tau)(1-\beta)/\beta = (1-\tau+\tau\beta)/\beta$ and $\beta R_{ss} = 1 - \tau + \tau\beta$. The debt condition (27) is a Kuhn–Tucker condition and holds whether or not the constraint binds, with $\mu_{ss} \geq 0$ and $\mu_{ss} = 0$ if the constraint is slack. Evaluated at the steady state with $\lambda_{ss} = 1$ it reads $\beta = 1/R_{ss} - \mu_{ss}\zeta_{ss}\beta$, so that

$$\mu_{ss}\,\zeta_{ss}\,\beta R_{ss} = 1 - \beta R_{ss} = \tau(1-\beta).$$

If $\tau > 0$ the right-hand side is strictly positive, so $\mu_{ss} > 0$ and the constraint must bind.

Conversely, $\tau = 0$ implies $\mu_{ss} = 0$. In that case complementary slackness does not determine whether the constraint holds with equality or is slack; the answer depends on the equilibrium selection for debt. Solving for μ_{ss} delivers (36). □

The tax advantage of debt is therefore the reason the constraint binds in steady state: firms lever up to exhaust the shield until enforcement stops them. Equation (36) also shows that the average wedge is decreasing in ζ_{ss}, so the steady-state severity of the friction is disciplined by a single parameter, and it is increasing in τ and in the steady-state real rate.

Binding in steady state is necessary but not sufficient for credit conditions to matter. The next result identifies the case in which they do not.

Proposition 2. *If $\tau = 0$, $\kappa = 0$, and $\zeta_t > 0$, then $\mu_t = 0$ at every date. Innovations to ζ_t, whether to its level or its volatility, have no effect on real allocations or on the associated price system. Debt and payout remain financing variables whose values can depend on equilibrium selection.*

Proof. With $\kappa = 0$ the payout condition gives $\lambda_t \equiv 1$, so (27) reads $\mathbb{E}_t[m_{t,t+1}/\Pi_{t+1}] = 1/R_t - \mu_t\zeta_t/(1+i_t)$. The bond pricing condition (15) gives $\mathbb{E}_t[m_{t,t+1}/\Pi_{t+1}] = 1/(1+i_t)$, and $\tau = 0$ implies $R_t = 1 + i_t$. Combining, $\mu_t\zeta_t = 0$, so $\mu_t = 0$ for $\zeta_t > 0$. With $\mu_t = 0$ the parameter ζ_t drops out of the conditions that determine real allocations and prices. It remains in the enforcement constraint, which restricts the admissible combinations of debt and payout. Because payouts are costless when $\kappa = 0$, those financing choices need not be unique and do not feed back into the real allocation. □

Proposition 2 is the analogue of the neutrality result in Jermann and Quadrini (2012), and it makes clear that the financing block is not a free amplification device: with frictionless payouts and no tax shield, the real side of the model collapses to a sticky-wage economy with flexible prices. Credit conditions may select a balance-sheet configuration, but they do not alter real allocations or prices.

Remark 1. The same algebra with $\kappa = 0$ but $\tau > 0$ yields $\mu_t = \tau i_t/(\zeta_t R_t)$ at every date. The multiplier is then pinned down statically by the policy rate and the enforcement parameter, and is *invariant to the value of collateral* Q_t. Credit conditions still move the labor wedge through ζ_t, but a revaluation of pledged capital does not. A positive payout adjustment cost is therefore necessary for the collateral-price channel studied here: it makes λ_t endogenous and time-varying, and it is only through the path of λ_t that the value of collateral can reach the wedge, as (38) below makes explicit.

The collateral price is an investment-rate object. Combining the investment condition with the valuation rule (23) eliminates both multipliers and leaves a single relationship.

Lemma 1. *On any equilibrium path with* $\lambda_t > 0$ *and* $1 - \phi\left(I_t/k_t - \delta\right) > 0$,

$$Q_t = \left[1 - \phi\left(\frac{I_t}{k_t} - \delta\right)\right]^{-1}, \tag{37}$$

so that $Q_{ss} = 1$ *and* $\partial Q_t / \partial (I_t/k_t) = \phi\, Q_t^2 > 0$ *whenever* $\phi > 0$. *If* $\phi = 0$ *then* $Q_t \equiv 1$.

Lemma 1 makes collateral value a strictly increasing function of the investment rate. Higher desired investment therefore raises the price applied to a predetermined collateral stock and expands borrowing capacity. If $\phi = 0$, Q_t is constant and this revaluation channel disappears. Combined with Remark 1, the result assigns distinct roles to the two adjustment costs: ϕ makes pledgeable value move, while κ transmits that movement to the financing wedge.

What pins down the financing wedge. Lemma 1 describes how the value of collateral moves; it does not by itself say how the labor wedge responds. Rearranging (27),

$$\frac{\mu_t}{\lambda_t} = \frac{1 + i_t}{\zeta_t}\left[\frac{1}{R_t} - \frac{\mathbb{E}_t[m_{t,t+1}\lambda_{t+1}/\Pi_{t+1}]}{\lambda_t}\right]. \tag{38}$$

The collateral price does not appear in (38). Conditional on the policy rate, enforcement, and the expected path of the marginal value of internal funds, the debt margin determines the wedge. If the constraint binds, (21) determines debt residually, $b_{t+1} = (1 + i_t)\left[Q_t k_t - w_t n_t/\zeta_t\right]$ with $b_{t+1} \equiv B_{t+1}/P_t$. A revaluation of collateral therefore reaches the wedge only indirectly. It changes how much the firm can borrow; the additional borrowing is reconciled with investment and payouts in the flow of funds; and the resulting movement in λ_t relative to its expected future value moves the wedge through (38). The wedge falls when the marginal value of internal funds is expected to rise, when the firm is flush today relative to tomorrow, and rises in the opposite case. Because μ_t and λ_t both respond, the direction of the movement is not signed by the structure of the model; it is a general-equilibrium outcome. When $\kappa = 0$, $\lambda_t \equiv 1$ and (15) reduces (38) to the expression in Remark 1, which is why the payout friction is necessary for this route to operate.

In a slack state, debt instead satisfies $b_{t+1} < (1 + i_t)[Q_t k_t - w_t n_t/\zeta_t]$ and $\mu_t = 0$. The contemporaneous financing wedge vanishes, and labor demand reduces to $w_t = (1 - \theta) y_t / n_t$. Borrowing is then determined jointly by the flow of funds and the intertemporal conditions, not by exhausting collateral. Future binding states can still affect current choices through the Euler equations. Equation (38) holds in both regimes; its right-hand side is zero in a slack equilibrium.

How a favorable noise signal can generate a boom and then a bust. Consider a favorable signal at $t = 1$ with maturity at $t = 1 + H$. The possibility that it contains genuine news raises the expected marginal product of capital and, because capital also serves as collateral, it can raise the expected collateral payoff $\mu_{t+1}\zeta_{t+1}Q_{t+1}$ in (26). To the extent that these raise q_t relative to λ_t, the collateral price rises and, by Lemma 1, so does the investment rate. These are

two expressions of the same equilibrium movement. The higher collateral price raises borrowing capacity against a predetermined capital stock. The firm's use of that capacity and its allocation between investment and payouts determine the path of λ_t. Through (38) that path moves the financing wedge, and through (25) the wedge moves labor demand. A boom arises if the wedge falls, so that employment and output rise before any productivity gain has materialised. If capacity expands sufficiently, the constraint becomes slack and the wedge reaches zero. Further movements in hours then reflect wages, productive capacity and intertemporal choices rather than a further decline in the financing wedge.

If the signal is revealed to be noise at $t = 1 + H$, the expected productivity gain is removed. Investment and collateral value can fall while the debt accumulated during anticipation remains due. A lower value of collateral against that debt can make the constraint bind again, restoring the hiring wedge and producing a contraction. The favorable signal supports borrowing during anticipation, and its resolution leaves the economy with the resulting repayment obligations. The noise realization contributes no productivity innovation. The boom–bust response, the duration of slackness, and rebinding after the signal is resolved are quantitative outcomes to be checked in the DEQN solution, not restrictions imposed on the decision rules.

Why nominal wage rigidity is required. The need for wage rigidity arises from the noise-shock experiment. Following an uncertainty shock, a higher financing wedge can generate a contraction in output, consumption, investment and hours even with flexible wages. During anticipation of a favorable signal, however, a lower wedge is insufficient to produce an expansion in all four variables. In the flexible-wage benchmark, consumption rises while output, investment and hours fall. The obstacle resembles the co-movement problem in the real model discussed by Basu and Bundick (2017), although here it arises in response to a signal about future productivity.

When $\phi_w = 0$, (16) requires $w_t = \mu_w \, \alpha n_t^{\nu} / \Lambda_t$. The prospect of a productivity gain raise consumption and reduce the household's marginal value of wealth, increasing the wage required to supply a given amount of labor. This wage increase offsets the expansionary effect of the lower financing wedge in (25), so hours fall despite easier credit. With $\phi_w > 0$, the real wage need not satisfy the static markup condition each period. Slower nominal wage adjustment limits the wage response during anticipation, allowing labor demand and hours to rise alongside consumption, output and investment. Wage rigidity is therefore introduced to obtain co-movement in the signal-driven expansion; it is not needed solely to obtain co-movement in the uncertainty-driven contraction.

The policy rule then determines how much of that shift survives. Through (28), a boom in which the real wage rises faster than nominal wages adjust is a boom with falling price inflation, and (29) responds by cutting the policy rate. This lowers the effective cost of debt in (19). It may also raise the present value of the collateral payoff in (26) by lowering the real discount rate, although the firm discounts with the household's stochastic discount factor rather than directly with i_t. The output term pushes in the opposite direction. The net response depends on

ϕ_π relative to ϕ_y. Christiano et al. (2008) find, in a related setting, that the boom following a signal that does not materialize is substantially larger than the efficient response to the same signal, and attribute the difference to a policy rule that fails to track the rise in the natural rate. Section 3.8 estimates ϕ_π and ϕ_y rather than fixing them, precisely because they are part of the mechanism.

Macro uncertainty. The second experiment raises the scale of enforcement innovations through $\varepsilon_{h,t}$ in (33), retaining the positive-support truncation described above. In the debt margin (38), greater dispersion of future enforcement changes the expected path of the marginal value of internal funds, because firms anticipate states in which borrowing capacity contracts and payouts must be cut. In the capital Euler equation (26) it changes the expected collateral payoff $\mathbb{E}_t[\mu_{t+1}\zeta_{t+1}Q_{t+1}]$, which depends on the covariance between future enforcement, the multiplier and the collateral price as well as on their means. In the environment of Chatterjee et al. (2025) the net effect is a precautionary tightening in which output, consumption, investment and hours fall together. Here the same forces operate through (38) and (26), and any fall in investment lowers the collateral price through Lemma 1 and contracts borrowing capacity further. Unlike the anticipation experiment, precautionary tightening can raise a positive multiplier without a switch to slackness. Whether the constraint remains binding throughout this response must be checked in the solution. Both experiments act through the financing margin, but they need not generate the same temporal pattern.

This distinction corresponds to the empirical separation of a predominantly contractionary uncertainty shock from a credit-spread shock with boom–bust dynamics. The model is asked to reproduce both responses with one financing block. It does not identify the two shocks by the sign of the spread response alone. Because this environment adds payout costs, nominal debt and a monetary block to that of Chatterjee et al. (2025), neither four-way co-movement nor a subsequent reversal follows analytically. These are restrictions against which the DEQN responses are evaluated.

3.8 Calibration and Estimation

Following Christiano et al. (2005), we split the parameters into two groups. The first is fixed ex ante, either from steady-state relationships or from values that are standard in the literature. The second is estimated by minimising the distance between the model's impulse responses, computed from the DEQN solution, and the two sets of empirical responses identified in Section 2: the responses to a credit shock, which the noise-shock experiment is asked to match, and the responses to an uncertainty shock, which the volatility experiment is asked to match. Table 1 lists both groups. Its numerical values are benchmark calibrations or starting values, not reported estimates.

For estimation, the estimated parameters enter the network as additional inputs, or pseudo-

states, as in Scheidegger (2026, Chapters 9–10). The network is trained over their admissible ranges with each parameter vector held fixed along a simulated path. This approximates a family of equilibrium decision rules $g(s_t;\vartheta)$ without re-solving the model at every objective evaluation. For a candidate ϑ, stack the model generalised impulse responses over forty quarters into $m(\vartheta)$ and the corresponding empirical responses into $\hat{m}$. The estimate minimises

$$J(\vartheta) = [m(\vartheta) - \hat{m}]' W [m(\vartheta) - \hat{m}], \tag{39}$$

where W is constructed from the covariance matrix of the empirical responses. Common random numbers are retained across candidates, and the search uses derivative-free methods because signal resolution is discrete.

The credit-spread target is the empirical shock identified after separating the uncertainty component; its model counterpart is a favorable signal about future productivity that is subsequently revealed to be noise, not an enforcement-level innovation. The uncertainty target is matched to an innovation in h_t. Output, consumption, investment, hours, the policy rate and inflation provide direct model counterparts. The empirical credit spread of Gilchrist and Zakrajšek (2012), excess bond premium, and high-yield issuance responses motivate the financing interpretation, but this model has no default-risk or bond-pricing block that reproduces those series directly. They are therefore not equated to μ_t/λ_t or included as model observables without an additional measurement relation.

The calibrated block is standard. The discount factor implies an annualised steady-state real rate of about four percent; the capital share, depreciation rate and technology persistence take conventional values. The elasticity of substitution across labor varieties implies a steady-state wage markup of twenty percent, and the tax parameter follows Jermann and Quadrini (2012). Signal horizons range from eight to sixteen quarters, with mean $8 + 8p_H$. Their distribution is estimated using the timing of the empirical credit-shock response. The genuine-news and noise experiments use the same signal magnitude and horizon for comparability.

The data discipline the two parameters that govern the collateral-price channel. Lemma 1 makes ϕ the elasticity of collateral value to the investment rate, while Remark 1 shows that κ controls whether collateral revaluation reaches the labor wedge. The estimate of ϕ_w is reported together with the contract duration implied by the mapping in Born and Pfeifer (2020).

The steady-state enforcement parameter admits a similar external check. Appendix B.5 shows that, holding the calibrated block fixed, the deterministic steady state maps ζ_{ss} monotonically into the ratio of corporate debt to quarterly output, so the estimate can be read against observed leverage even though it is not calibrated to it. We report that implied ratio alongside the estimate, together with the two admissibility restrictions derived in the same appendix.

Every candidate parameter vector must satisfy two restrictions. First, the deterministic steady state must exist, with positive intertemporal debt and a binding enforcement constraint. At the values in Table 1, this requires ζ_{ss} to exceed approximately 0.06; there is no upper bound.

Appendix B.5 derives the restrictions and explains why the direction of the bound differs from Chatterjee et al. (2025). Second, the DEQN must satisfy the equilibrium residual, feasibility and stability checks in Appendix B.9 throughout the parameter region used for estimation. Complementarity is imposed by (35); both regimes are admissible. Positive consumption surplus, capital, prices and marginal values, and the investment branch in (B.25), are checked separately. Enforcement remains positive under the same truncated transition in expectations and simulations. Failed paths and domain exits are reported, not hidden by averaging the surviving paths. At the estimate, a dedicated fixed-parameter network is solved with stricter accuracy requirements and its responses are compared with those of the parameterised network. Artificial-data parameter recovery and local identification checks assess the estimator.

Table 1: Parameters

Parameter	Description	Value	Treatment
Preferences and technology			
β	Discount factor	0.9900	Calibrated
σ	Consumption curvature	4.9831	Estimated
ι	Habit persistence	0.5318	Estimated
ν	Inverse Frisch elasticity	0.3903	Estimated
θ	Capital share	0.3300	Calibrated
δ	Depreciation rate	0.0250	Calibrated
ϕ	Investment adjustment cost	12.8704	Estimated
labor market and nominal rigidity			
η_w	Elasticity of substitution across varieties	9.4619	Estimated
ϕ_w	Rotemberg wage adjustment cost	63.9835	Estimated
Financial structure			
κ	Payout adjustment cost	0.0917	Estimated
τ	Tax advantage of debt	0.350	Calibrated
ζ_{ss}	Steady-state enforcement parameter	0.1516	Estimated
Monetary policy			
ϕ_π	Response to inflation	2.1254	Estimated
ϕ_y	Response to output	0.2500	Calibrated
Exogenous processes			
ρ_z	Persistence of technology	0.950	Calibrated
σ_0	S.d. of productivity surprises and signals	0.0080	Estimated
λ_{news}	Signal arrival probability	0.0312	Estimated
p_H	Signal horizon distribution parameter	0.7480	Estimated
π_{confirm}	Signal confirmation probability	0.9354	Estimated
ρ_ζ	Persistence of enforcement parameter	0.950	Calibrated
$\bar{h}$	Mean log variance of enforcement innovation	-10.27	Calibrated
ρ_h	Persistence of log volatility	0.8836	Estimated
σ_h	S.d. of volatility shock	0.1378	Estimated

Notes. The model is calibrated at quarterly frequency. Calibrated parameters are fixed prior to estimation; estimated parameters are chosen to minimise the distance between model-implied and empirical impulse responses. Values shown for estimated parameters are starting values; dashes indicate signal-process parameters for which no final values are reported here. The value of $\bar{h}$ is the posterior mean reported for the baseline stochastic-volatility specification in Chatterjee et al. (2025).

3.9 Model Results

Figures 11 and 12 compare the responses to macro uncertainty and a noise shock about future productivity with their empirical targets. The comparison concerns two restrictions: co-movement among output, consumption, investment and hours, and the distinction between a contraction followed mainly by recovery and an expansion followed by a contraction. The latter is the feature that distinguishes the empirical credit-spread shock after separating the uncertainty component. Since these responses are estimation targets, their fit assesses the model's ability to account for the conditional dynamics rather than providing an independent test of the mechanism.

Macro uncertainty. Figure 11 shows that an increase in macro uncertainty lowers all four quantities. The responses are hump shaped: activity falls further after impact and then recovers gradually. Investment and hours respond most strongly, with approximate troughs of -1.6 and -0.65, respectively, in the figure's impact-normalised units. Consumption reaches a trough of about -0.3. Most of the decline occurs within the first ten quarters. Output, consumption and investment remain below baseline at the end of the forty-quarter window, while hours exhibit a small late overshoot. Thus the dominant response is a contraction and recovery, although the hours response is not entirely one sided.

The model follows the empirical decline and recovery in output and hours reasonably closely. Its investment contraction is somewhat larger, and its consumption recovery is faster than the empirical median. There is a more substantial difference in the uncertainty panel. The model's uncertainty response decays toward zero from above, whereas the empirical uncertainty response becomes negative before returning toward baseline. The figure therefore supports the contractionary quantity responses more closely than the subsequent path of the uncertainty measure itself.

The financing mechanism differs from the demand-uncertainty channel of Basu and Bundick (2017). Here greater dispersion in enforcement exposes firms to future shortages of borrowing capacity and changes the value of internal funds. A higher financing wedge shifts labor demand inward through (25), a hiring channel related to that of Jermann and Quadrini (2012). As discussed in Section 3.7, this channel can generate four-way contraction even with flexible wages. The quantity responses are consistent with that mechanism, but the figure alone does not establish the multiplier's path or whether the ICC remains binding at every date.

Noise about future productivity. Figure 12 displays a different temporal pattern. A favorable signal that is subsequently revealed to be noise produces an initial expansion in output, consumption, investment and hours, followed by responses below baseline. The expansion peaks within roughly the first ten quarters. Output rises by about 0.3 percent, consumption by 0.25 percent, investment by 1.2 percent and hours by 0.6 percent. Investment and hours subsequently turn negative, followed by output and consumption. The contraction is persistent, with the

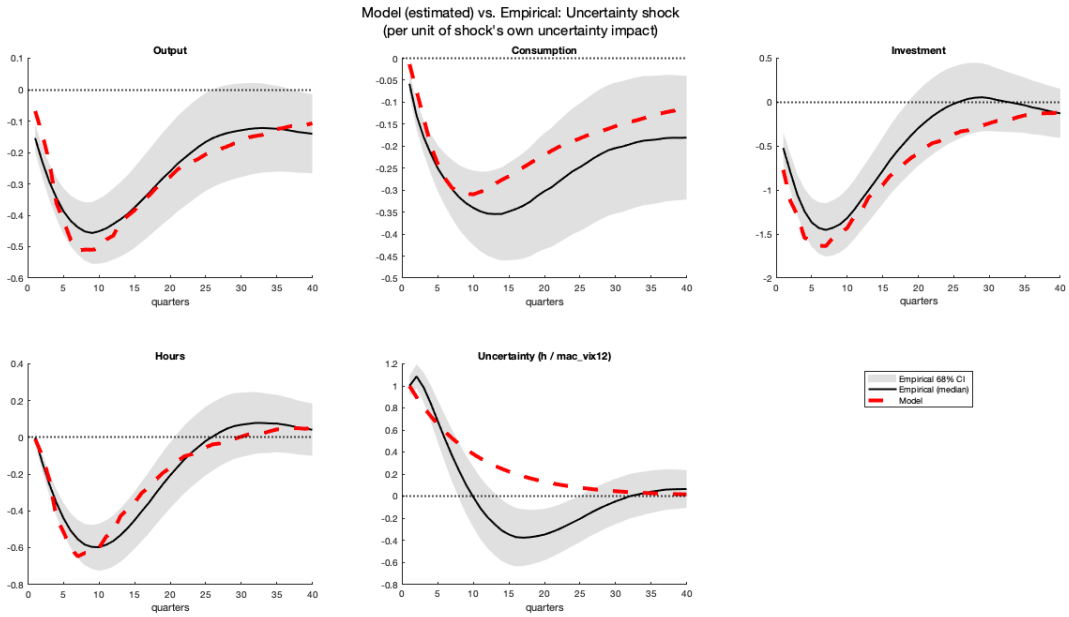


Figure 11: Responses to a macro-uncertainty shock

Notes. Red dashed lines show the supplied model responses, black lines the empirical medians, and shaded areas the empirical 68 percent confidence intervals. Each set of responses is normalised by the impact response of uncertainty in the corresponding model or empirical system. The horizon is forty quarters. Shading describes empirical uncertainty, not uncertainty about the model's parameters or numerical solution.

largest negative responses occurring broadly between quarters twenty-five and thirty-five. All four quantities remain below baseline at quarter forty.

The model captures the empirical ordering of expansion and contraction, including the relatively large movements in investment and hours. The fit is less close for consumption: its model expansion lasts longer and its subsequent contraction is larger than the empirical median. The model also generates deeper and more persistent declines in output, investment and hours. These differences matter because the empirical consumption cycle is weaker than the cycles in the other quantities. Broad empirical confidence bands do not by themselves establish that the discrepancies are quantitatively unimportant.

The expansion is consistent with the collateral-revaluation channel in Lemma 1. The prospect of higher future productivity encourages investment and raises the value attached to predetermined capital. Easier financing can reduce the hiring wedge, while slow nominal wage adjustment allows hours to rise alongside consumption. When the signal is revealed to be noise, the anticipated gain is removed, but the financing and investment choices made during anticipation carry forward. A decline in collateral value against inherited debt can restore a binding constraint and reverse the expansion. This interpretation requires the debt, slack and multiplier paths described in Section 3.7; the four quantity panels cannot independently verify the proposed switch from

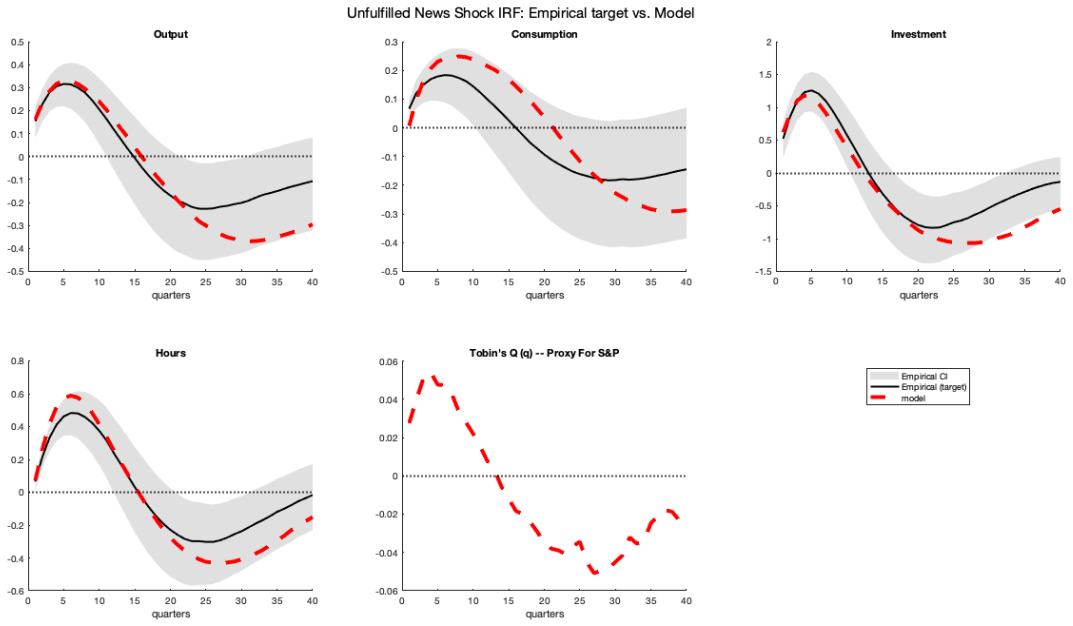


Figure 12: Responses to a Noise Shock about Future Productivity

Notes. Red dashed lines show the supplied model responses, black lines the empirical credit-shock targets, and shaded areas the empirical 68 percent confidence intervals. Quantity responses are percentage deviations from baseline over forty quarters. The fifth panel shows model-implied Tobin's Q, for which there is no exact counterpart in the empirical IRFs. Its boom–bust pattern is compared qualitatively with the empirical S&P response; no empirical confidence band is plotted for Q.

binding to slack and back to binding.

The fifth panel shows model-implied Tobin's Q. Its initial rise and subsequent fall below baseline accompany the expansion and contraction in real activity. This valuation cycle is consistent with the link between investment and collateral value in Lemma 1. It also parallels the boom–bust response of S&P prices to the empirical credit-spread shock. Tobin's Q values installed capital in units of goods, whereas the S&P index measures equity prices. The comparison therefore concerns the direction and sequence of valuation movements, not an exact observable counterpart or a quantitative match of their magnitudes.

Connection to the empirical credit cycle. The empirical credit-related shock produces boom–bust responses in real activity and stock prices and accounts for a substantial share of excess bond premium fluctuations. These findings motivate the model's link between a favorable signal, an expansion in borrowing, and a subsequent reversal, consistent with the sequence of buoyant credit conditions and later weakness emphasised by López-Salido et al. (2017). In the model, that sequence results from rational responses to an imperfect signal and its resolution. The model does not reproduce an investor-belief or risk-appetite mechanism, nor does it directly

generate excess bond premia, the high-yield issuance ratio or the S&P index. Nevertheless, the boom–bust response of model-implied Q provides a valuation analogue to the empirical stock-price cycle alongside the co-movement in real activity.

The two experiments account for distinct conditional patterns within the same financing environment. Uncertainty mainly generates contraction and recovery; the productivity noise shock generates an expansion and a subsequent contraction. This distinction is consistent with the empirical finding that separating uncertainty reveals a credit-related cycle that is less apparent in the jointly identified shock. It does not establish the relative contributions of the two model shocks to aggregate volatility or their spectral densities. Those implications require simulated shock contributions and spectra in addition to the impulse responses reported here.

4 Conclusion

Hours worked has a pronounced spectral peak and sign-changing autocorrelation, yet the shock that explains most of its business-cycle variation produces no clear bust. We show that the MBC shock masks two disturbances. Higher macro uncertainty generates a broad recession followed by recovery. An orthogonal credit-related shock instead raises output, consumption, investment, hours, and stock prices before these variables reverse. It also explains substantial variation in the excess bond premium and is closely related to the high-yield issuance share. The uncertainty component dominates within the MBC shock, muting the reversal in its impulse response even though it remains visible in the spectrum.

These facts discipline a model in which firms borrow against collateral to finance their wage bill. A favorable signal about future productivity raises investment and collateral value and relaxes the borrowing constraint. If the signal is subsequently revealed to be noise, inherited debt and lower collateral value tighten financing conditions and reverse the expansion. Uncertainty about enforcement operates through the same financial environment but immediately raises the financing wedge and contracts activity. Slow nominal wage adjustment generates comovement after the signal, while the financial wedge is sufficient for the contractionary uncertainty response. A global neural-network solution allows the borrowing constraint to switch endogenously between binding and slack regimes.

The quantitative model reproduces the central conditional distinction and a boom–bust movement in Tobin's Q consistent with the stock-price response. The exercise also has limits: the empirical responses are estimation targets, the model does not directly generate the excess bond premium or the high-yield issuance share, and some quantity responses are too persistent. Matching unconditional variance shares and shock-specific spectra is the natural next test of how much business-cycle variation can be attributed to uncertainty and credit-market reversals.

References

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas, "Business-Cycle Anatomy," *American Economic Review*, 2020, *110* (10), 3030–3070.

Arellano, Cristina, Yan Bai, and Patrick J. Kehoe, "Financial Frictions and Fluctuations in Volatility," *Journal of Political Economy*, 2019, *127* (5), 2049–2103.

Azariadis, Costas, Leo Kaas, and Yi Wen, "Self-Fulfilling Credit Cycles," *Review of Economic Studies*, 2016, *83* (4), 1364–1405.

Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger, "Deep Equilibrium Nets," *International Economic Review*, 2022, *63* (4), 1471–1525.

Basu, Susanto and Brent Bundick, "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 2017, *85* (3), 937–958.

_ , Giacomo Candian, Ryan Chahrour, and Rosen Valchev, "Risky Business Cycles," NBER Working Paper 28693, National Bureau of Economic Research 2025.

Beaudry, Paul, Dana Galizia, and Franck Portier, "Putting the Cycle Back into Business Cycle Analysis," *American Economic Review*, 2020, *110* (1), 1–47.

Benhabib, Jess and Pengfei Wang, "Financial Constraints, Endogenous Markups, and Self-Fulfilling Equilibria," *Journal of Monetary Economics*, 2013, *60* (7), 789–805.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist, "The Financial Accelerator in a Quantitative Business Cycle Framework," in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Vol. 1, Amsterdam: Elsevier, 1999, pp. 1341–1393.

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

Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and Stephen J Terry, "Real credit cycles," *American Economic Review*, 2026, *116* (4), 1274–1308.

Born, Benjamin and Johannes Pfeifer, "The New Keynesian Wage Phillips Curve: Calvo vs. Rotemberg," *Macroeconomic Dynamics*, 2020, *24* (5), 1017–1041.

Brianti, Marco and Vito Cormun, "Expectation-Driven Boom-Bust Cycles," *Journal of Monetary Economics*, 2024, *146*, 103575.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek, "The Macroeconomic Impact of Financial and Uncertainty Shocks," *European Economic Review*, 2016, *88*, 185–207.

Chahrour, Ryan and Kyle Jurado, "News or Noise? The Missing Link," *American Economic Review*, 2018, *108* (7), 1702–1736.

Chatterjee, Pratiti, David Gunawan, and Robert Kohn, "The Interaction between Credit Constraints and Uncertainty Shocks," *Journal of Money, Credit and Banking*, 2025, *57* (8).

Christiano, Lawrence J., Cosmin L. Ilut, Roberto Motto, and Massimo Rostagno, "Monetary Policy and Stock Market Boom-Bust Cycles," ECB Working Paper 955, European Central Bank 2008.

_ , _ , _ , and _ , "Monetary Policy and Stock Market Booms," NBER Working Paper 16402, National Bureau of Economic Research 2010.

_ , Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 2000, *46* (2), 281–313.

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

Forni, Mario, Luca Gambetti, Antonio Granese, Luca Sala, and Stefano Soccorsi, "An American Macroeconomic Picture: Supply and Demand Shocks in the Frequency Domain," *American Economic Journal: Macroeconomics*, 2025, *17* (3), 311–341.

_ , _ , Marco Lippi, and Luca Sala, "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 122–152.

Gilchrist, Simon and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102* (4), 1692–1720.

Granese, Antonio, "Two Main Business Cycle Shocks Are Better than One," *Journal of Economic Dynamics and Control*, 2026, *182*, 105231.

Greenwood, Robin and Samuel G. Hanson, "Issuer Quality and Corporate Bond Returns," *Review of Financial Studies*, 2013, *26* (6), 1483–1525.

Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, 2012, *102* (1), 238–271.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, "When Credit Bites Back," *Journal of Money, Credit and Banking*, 2013, *45* (s2), 3–28.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, "Measuring Uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Kase, Hanno, Leonardo Melosi, and Matthias Rottner, "Estimating Nonlinear Heterogeneous Agent Models with Neural Networks," BIS Working Paper 1241, Bank for International Settlements January 2025.

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

Leduc, Sylvain and Zheng Liu, "Uncertainty Shocks Are Aggregate Demand Shocks," *Journal of Monetary Economics*, 2016, *82*, 20–35.

Lintner, John, "Distribution of Incomes of Corporations among Dividends, Retained Earnings, and Taxes," *American Economic Review*, 1956, *46* (2), 97–113.

Liu, Zheng and Pengfei Wang, "Credit Constraints and Self-Fulfilling Business Cycles," *American Economic Journal: Macroeconomics*, 2014, *6* (1), 32–69.

López-Salido, David, Jeremy C. Stein, and Egon Zakrajšek, "Credit-Market Sentiment and the Business Cycle," *Quarterly Journal of Economics*, 2017, *132* (3), 1373–1426.

Lorenzoni, Guido, "A Theory of Demand Shocks," *American Economic Review*, 2009, *99* (5), 2050–2084.

Ludvigson, Sydney C., Sai Ma, and Serena Ng, "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 369–410.

Maliar, Lilia, Serguei Maliar, and Pablo Winant, "Deep Learning for Solving Dynamic Economic Models," *Journal of Monetary Economics*, 2021, *122*, 76–101.

Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycles Worldwide," *Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

Scheidegger, Simon, "Deep Learning for Solving and Estimating Dynamic Models in Economics and Finance," 2026. Lecture notes, arXiv:2605.14493.

Schularick, Moritz and Alan M. Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008," *American Economic Review*, 2012, *102* (2), 1029–1061.

Uhlig, Harald, "What Moves Real GNP?," 2003. Unpublished manuscript.

Wei, Yanhao and Zhenling Jiang, "Estimating Parameters of Structural Models Using Neural Networks," 2025. Working paper, arXiv:2502.04945.

Appendix A: Data Construction and Empirical Implementation

This appendix records the construction of the variables used in Section 2. The objective is to make the empirical inputs reproducible while keeping measurement details out of the main text.

A.1. Baseline VAR variables

Table A.1: Construction of the baseline variables

Variable	Construction	Source
Output	100 log of real GDP divided by the civilian noninstitutional population aged sixteen and older.	BEA National Income and Product Accounts; BLS population series.
Consumption	100 log of real consumption of nondurable goods and services per capita. Nominal expenditure shares are used to allocate real GDP to these two components.	BEA National Income and Product Accounts.
Investment	100 log of real gross private domestic investment plus consumer durable expenditure, divided by population. Nominal expenditure shares are used to form the aggregate.	BEA National Income and Product Accounts.
Hours	100 log of total hours in the nonfarm business sector divided by the civilian noninstitutional population aged sixteen and older.	BLS nonfarm business hours and population series.
Interest rate	Quarterly average of the effective federal funds rate, divided by four so that the rate is expressed at a quarterly frequency.	Federal Reserve Board.
TFP	Cumulated quarterly growth in utilization-adjusted total factor productivity. The reported annualized growth rate is divided by four before cumulation.	Fernald (2014).
Credit spread	Quarter-end Gilchrist–Zakrajšek corporate bond spread, divided by four.	Gilchrist and Zakrajšek (2012).
Macro uncertainty	Quarterly average of the monthly twelve-month-ahead macro uncertainty index, expressed as 100 log of the index.	Jurado et al. (2015).

The four real aggregates follow the construction used by Angeletos et al. (2020). The baseline

estimation begins in 1973:I and ends in 2024:IV. The current estimation input removes 2020:II and 2020:III, the two observations most affected by the abrupt pandemic shutdown and reopening. Apart from these exclusions, quarterly observations are aligned by calendar date before the sample is selected.

A.2. Comparison and auxiliary variables

The specification used for comparison with Angeletos et al. (2020) adds labor productivity and the labor share. Labor productivity is 100 log of nonfarm business output per hour, and the labor share is 100 log of the BLS nonfarm business labor-share index. Appendix specifications also consider inflation, measured as $100 \log(P_t/P_{t-1})$ using the GDP implicit price deflator. Inflation is not part of the eight-variable baseline VAR.

The auxiliary financial series are constructed separately because their samples differ from that of the baseline VAR. The real per-capita S&P index is $100 \log[S_t/(0.01 P_t N_t)]$, where S_t is the quarter-end S&P price index, P_t is the GDP deflator, and N_t is population. The excess bond premium is the quarter-end series of Gilchrist and Zakrajšek (2012), divided by four. The high-yield ratio is the share of nonfinancial corporate bond issuance rated high yield, constructed by Greenwood and Hanson (2013); the available quarterly series runs from 1983:I through 2016:I. These variables are projected on the contemporaneous baseline VAR vector and two of its lags, as described in equations (8) and (9).

A.3. Prior and posterior computation

The VAR contains two lags and an intercept. The Minnesota prior centers the first own lag of every variable at one. Its four hyperparameters are $(0.2, 0.5, 2, 10^5)$: the first controls overall tightness, the second shrinks cross-variable lags relative to own lags, the third increases shrinkage with lag length, and the fourth makes the intercept effectively diffuse. Innovation variances from univariate AR(1) regressions scale the prior across variables. The covariance matrix has an inverse-Wishart prior. The Gibbs sampler uses 50,000 draws, discards the first 49,000, and retains the remaining 1,000 draws for posterior summaries. Each draw is checked for stability before its frequency-domain rotation and impulse responses are computed.

Appendix B: Derivations of the Model

This appendix derives every optimality condition reported in Section 3, characterises the deterministic steady state in closed form, proves the results stated in Section 3.7, lists the complete equilibrium system, and describes how the noise-shock experiment and the impulse responses are computed. Throughout, φ_t denotes the multiplier on the production function and should not be confused with the investment adjustment cost parameter ϕ or the wage adjustment cost parameter ϕ_w. Table B.1 at the end of the appendix collects all notation.

B.1 The Household's Problem

Setup. Substituting the labor demand curve (12) directly into both the objective (10) and the budget constraint (11), and attaching the multiplier $\beta^t \Lambda_t(j)$ to the date-t budget constraint, household j chooses $\{c_t(j), B_{t+1}(j), W_t(j)\}_{t\geq 0}$ to maximise

$$\begin{aligned}\mathcal{L}_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ & \frac{(c_t(j) - \iota c_{t-1}(j))^{1-\sigma} - 1}{1-\sigma} \\ & - \frac{\alpha}{1+\nu} \left[\left(\frac{W_t(j)}{W_t} \right)^{-\eta_w} n_t \right]^{1+\nu} \\ & + \Lambda_t(j) \Bigg[\big(1 - \Theta_{w,t}(j)\big) \frac{W_t(j)}{P_t} \left(\frac{W_t(j)}{W_t} \right)^{-\eta_w} n_t \\ & \qquad + \frac{B_t(j)}{P_t} + d_t - T_t - c_t(j) - \frac{B_{t+1}(j)}{P_t(1+i_t)} \Bigg] \Bigg\}. \end{aligned} \tag{B.1}$$

The household takes W_t, n_t, P_t, i_t, d_t and T_t as given.

Consumption. Because habit is internal, $c_t(j)$ enters $\mathcal{L}_0$ at two dates: directly through date-t felicity, and indirectly through the habit stock $\iota c_t(j)$ that appears in date-$(t+1)$ felicity. Writing $X_t(j) \equiv (c_t(j) - \iota c_{t-1}(j))^{-\sigma}$ for the flow marginal felicity and noting that

$$\frac{\partial}{\partial c_t(j)} \left[\frac{(c_{t+1}(j) - \iota c_t(j))^{1-\sigma}}{1-\sigma} \right] = -\iota \left(c_{t+1}(j) - \iota c_t(j) \right)^{-\sigma} = -\iota X_{t+1}(j), \tag{B.2}$$

the derivative of (B.1) with respect to $c_t(j)$ collects $\beta^t X_t(j)$ from date t, $-\beta^{t+1} \iota \, \mathbb{E}_t X_{t+1}(j)$ from date $t+1$, and $-\beta^t \Lambda_t(j)$ from the budget constraint. Dividing by β^t and setting the result to zero gives the household-level condition

$$\Lambda_t(j) = X_t(j) - \iota \beta \, \mathbb{E}_t[X_{t+1}(j)] \, . \tag{B.3}$$

The correction term is strictly negative for $\iota > 0$, so $\Lambda_t(j) < X_t(j)$: the household values an extra unit of wealth less than the flow term alone would suggest, because consuming it raises tomorrow's reference stock. Under external habit the second term is absent and $\Lambda_t(j) = X_t(j)$. After symmetry is imposed, this condition becomes (14).

Bonds. Differentiating (B.1) with respect to $B_{t+1}(j)$ gives

$$-\beta^t \frac{\Lambda_t(j)}{P_t(1+i_t)} + \beta^{t+1}\,\mathbb{E}_t\left[\frac{\Lambda_{t+1}(j)}{P_{t+1}}\right] = 0 \quad \Longrightarrow \quad \frac{\Lambda_t(j)}{P_t(1+i_t)} = \beta\,\mathbb{E}_t\left[\frac{\Lambda_{t+1}(j)}{P_{t+1}}\right]. \tag{B.4}$$

Multiplying through by $P_t/\Lambda_t(j)$ and using $m_{t,t+1}(j) \equiv \beta\Lambda_{t+1}(j)/\Lambda_t(j)$ and $\Pi_{t+1} = P_{t+1}/P_t$ delivers the household-level bond-pricing condition

$$1 = (1+i_t)\,\mathbb{E}_t\left[\frac{m_{t,t+1}(j)}{\Pi_{t+1}}\right]. \tag{B.5}$$

After symmetry is imposed, this condition becomes (15). The discount factor $m_{t,t+1}(j)$ prices real payoffs; a nominal payoff must additionally be deflated, which is the source of the $1/\Pi_{t+1}$ term. The condition holds exactly under uncertainty, so no separate Fisher equation is imposed anywhere in the model.

The posted wage. The wage $W_t(j)$ enters (B.1) at date t through three channels and at date $t+1$ through one. We take them in turn, using $\partial n_t(j)/\partial W_t(j) = -\eta_w n_t(j)/W_t(j)$ throughout.

(i) labor disutility at t. Since a higher posted wage reduces the hours demanded of variety j,

$$\frac{\partial}{\partial W_t(j)}\left[-\frac{\alpha}{1+\nu}n_t(j)^{1+\nu}\right] = -\alpha n_t(j)^{\nu}\frac{\partial n_t(j)}{\partial W_t(j)} = +\alpha\eta_w\, n_t(j)^{\nu}\frac{n_t(j)}{W_t(j)}, \tag{B.6}$$

which is positive: raising the wage is attractive in part because it economises on work effort.

(ii) labor income at t. Using (12) we may write $W_t(j)n_t(j) = W_t(j)^{1-\eta_w}W_t^{\eta_w}n_t$, so that

$$\frac{\partial\left[W_t(j)n_t(j)\right]}{\partial W_t(j)} = (1-\eta_w)\,n_t(j), \tag{B.7}$$

which is negative because demand for the variety is elastic.

(iii) The adjustment cost at t. From (13) and $\Pi_{w,t}(j) = W_t(j)/W_{t-1}(j)$ with $W_{t-1}(j)$ predetermined,

$$\frac{\partial\Theta_{w,t}(j)}{\partial W_t(j)} = \frac{\phi_w\big(\Pi_{w,t}(j)-1\big)}{W_{t-1}(j)} = \frac{\phi_w\big(\Pi_{w,t}(j)-1\big)\Pi_{w,t}(j)}{W_t(j)}. \tag{B.8}$$

Combining (ii) and (iii), the date-t income term contributes

$$\Lambda_t(j)\frac{n_t(j)}{P_t}\Big[(1-\eta_w)\big(1-\Theta_{w,t}(j)\big) - \phi_w\big(\Pi_{w,t}(j)-1\big)\Pi_{w,t}(j)\Big]. \tag{B.9}$$

(iv) The adjustment cost at $t+1$. Since $\Pi_{w,t+1}(j) = W_{t+1}(j)/W_t(j)$ depends on the wage posted today,

$$\frac{\partial \Pi_{w,t+1}(j)}{\partial W_t(j)} = -\frac{\Pi_{w,t+1}(j)}{W_t(j)}, \qquad \frac{\partial \Theta_{w,t+1}(j)}{\partial W_t(j)} = -\frac{\phi_w\big(\Pi_{w,t+1}(j)-1\big)\Pi_{w,t+1}(j)}{W_t(j)}. \tag{B.10}$$

Because the cost enters the date-$(t+1)$ budget constraint multiplied by $W_{t+1}(j)n_{t+1}(j)/P_{t+1}$, and $W_{t+1}(j)/W_t(j) = \Pi_{w,t+1}(j)$, this term contributes

$$\beta\,\mathbb{E}_t\left[\Lambda_{t+1}(j)\phi_w\big(\Pi_{w,t+1}(j)-1\big)\Pi_{w,t+1}(j)^2\frac{n_{t+1}(j)}{P_{t+1}}\right]. \tag{B.11}$$

Note that a higher wage today lowers next period's wage inflation and therefore *reduces* next period's adjustment cost; this is the forward-looking term that makes wage setting dynamic.

Summing (i)–(iv), setting the total to zero, multiplying through by $P_t/\left(\Lambda_t(j)n_t(j)\right)$, and imposing symmetry ($W_t(j) = W_t$, $n_t(j) = n_t$, $\Theta_{w,t}(j) = \Theta_{w,t}$, $\Pi_{w,t}(j) = \Pi_{w,t}$, $c_t(j) = c_t$, $B_t(j) = B_t$, $\Lambda_t(j) = \Lambda_t$, and $m_{t,t+1}(j) = m_{t,t+1}$) gives, with $\mathrm{mrs}_t \equiv \alpha n_t^{\nu}/\Lambda_t$ and $w_t = W_t/P_t$,

$$\phi_w\big(\Pi_{w,t}-1\big)\Pi_{w,t} = (1-\eta_w)\big(1-\Theta_{w,t}\big) + \eta_w\frac{\mathrm{mrs}_t}{w_t} + \mathbb{E}_t\left[m_{t,t+1}\phi_w\big(\Pi_{w,t+1}-1\big)\Pi_{w,t+1}^2\frac{n_{t+1}}{n_t}\frac{1}{\Pi_{t+1}}\right], \tag{B.12}$$

which is (16). The first term on the right is the marginal revenue loss from a higher wage, the second the marginal saving in work effort, and the third the discounted marginal reduction in future adjustment costs. Symmetry is legitimate here because all households face the same aggregate state and the adjustment cost is convex in the household's own wage growth, so the wage distribution is degenerate at every date and no dispersion state is required.

Remark 2. Setting $\phi_w = 0$ implies $\Theta_{w,t} = 0$ and collapses (B.12) to $0 = (1-\eta_w) + \eta_w\,\mathrm{mrs}_t/w_t$, that is $w_t = \mu_w\,\mathrm{mrs}_t$ with $\mu_w = \eta_w/(\eta_w - 1)$, the constant-markup result of Erceg et al. (2000). This provides a direct numerical check on the coded model: imposing $\phi_w = 0$ should reproduce this static relationship period by period.

Remark 3. Born and Pfeifer (2020) write the household-level cost in the more general form $\frac{\phi_w}{2}\big(\Pi_{w,t}(j)/\Gamma_{t-1,t} - 1\big)^2\Xi_t$, where $\Gamma_{t-1,t}$ is an indexation term and Ξ_t a nominal scaling base. Specification (13) is the special case with no indexation and the cost measured as a share of the wage bill, which is why (B.12) contains no backward-looking term. Their slope-equivalence formula maps a given ϕ_w into the average contract duration of an observationally equivalent staggered-contract economy, which we use as an external check on the estimate of ϕ_w.

B.2 The Firm's Problem

Setup. Attaching multipliers λ_t to the flow of funds (20), q_t to the capital accumulation equation (18), φ_t to the production function (17), and μ_t to the enforcement constraint (21), the

firm's date-t Lagrangian is

$$\begin{aligned}\mathcal{L}_t = d_t + \mathbb{E}_t[m_{t,t+1}V_{t+1}(k_{t+1}, B_{t+1})] \\ + \lambda_t \left\{ y_t + \frac{B_{t+1}}{P_t R_t} - d_t - \kappa \left(d_t - \bar{d}\right)^2 - w_t n_t - I_t - \frac{B_t}{P_t} \right\} \\ + q_t \left\{ (1-\delta)k_t + I_t - \frac{\phi}{2}\left(\frac{I_t}{k_t} - \delta\right)^2 k_t - k_{t+1} \right\} \\ + \varphi_t \left\{ z_t k_t^{\theta} n_t^{1-\theta} - y_t \right\} + \mu_t \left\{ \zeta_t \left[Q_t k_t - \frac{B_{t+1}}{P_t(1+i_t)} \right] - w_t n_t \right\}. \end{aligned} \tag{B.13}$$

The firm takes w_t, i_t, P_t, $m_{t,t+1}$ and the aggregate collateral price Q_t as given. The valuation rule (23) is imposed on the symmetric equilibrium only *after* the conditions below are derived; if it were substituted into (B.13) beforehand, an individual firm would incorrectly perceive that its own investment and payout choices move the contractual price at which its collateral is assessed.

Payout. Differentiating (B.13) with respect to d_t,

$$1 - \lambda_t \left[1 + 2\kappa \left(d_t - \bar{d}\right)\right] = 0 \quad \Longrightarrow \quad \lambda_t = \frac{1}{1 + 2\kappa \left(d_t - \bar{d}\right)}. \tag{B.14}$$

The marginal value of internal funds falls below unity when the firm pays out more than its target and rises above unity when it pays out less. When $\kappa = 0$ we have $\lambda_t \equiv 1$: internal and external funds are perfect substitutes and the firm is indifferent about the timing of payouts.

Output. Since y_t appears in (20) and (17) but not in (21)—the constraint restricts the wage bill, not revenue—the credit wedge does not route through the production multiplier:

$$\lambda_t = \varphi_t. \tag{B.15}$$

This is the step that makes the labor condition below take the form of a wedge on the wage rather than on the marginal product, and it distinguishes the present specification from one in which the constraint bounds output.

labor. Differentiating with respect to n_t and using (B.15),

$$-\lambda_t w_t + \varphi_t (1-\theta)\frac{y_t}{n_t} - \mu_t w_t = 0 \quad \Longrightarrow \quad w_t \left(1 + \frac{\mu_t}{\lambda_t}\right) = (1-\theta)\frac{y_t}{n_t}, \tag{B.16}$$

which is (25). The wedge is μ_t/λ_t rather than μ_t because the shadow value of the constraint must be expressed per unit of internal funds, whose marginal value is λ_t. In the special case $\kappa = 0$ the two coincide.

Investment. Differentiating with respect to I_t,

$$-\lambda_t + q_t\left[1-\phi\left(\frac{I_t}{k_t}-\delta\right)\right]=0 \quad\Longrightarrow\quad \lambda_t = q_t\left[1-\phi\left(\frac{I_t}{k_t}-\delta\right)\right]. \tag{B.17}$$

Capital. Because (21) constrains the predetermined stock k_t rather than the choice k_{t+1}, the latter appears only in the continuation value and in (18):

$$\mathbb{E}_t[m_{t,t+1}V_{k,t+1}] - q_t = 0 \quad\Longrightarrow\quad q_t = \mathbb{E}_t[m_{t,t+1}V_{k,t+1}]\,. \tag{B.18}$$

Had the constraint been written in k_{t+1}, this condition would carry an additional contemporaneous term $\mu_t\zeta_t Q_t$, and investing today would directly relax today's borrowing limit. Under the timing adopted here the collateral benefit of capital appears instead in the envelope condition, and hence one period later in the Euler equation.

Debt. Differentiating with respect to B_{t+1}, which enters both the flow of funds and the enforcement constraint,

$$\mathbb{E}_t[m_{t,t+1}V_{B,t+1}] + \frac{\lambda_t}{P_t R_t} - \frac{\mu_t\zeta_t}{P_t(1+i_t)} = 0. \tag{B.19}$$

The second term is the marginal benefit of borrowing, discounted at the tax-advantaged rate R_t; the third is the marginal cost of consuming pledgeable collateral, since one unit of new debt reduces the right-hand side of (21) by $\zeta_t/\big(P_t(1+i_t)\big)$.

Envelope conditions. The individual state k_t enters (B.13) in three places. From the capital accumulation term, writing $u_t \equiv I_t/k_t - \delta$ and using $\partial u_t/\partial k_t = -I_t/k_t^2$,

$$\frac{\partial}{\partial k_t}\left[-\frac{\phi}{2}u_t^2 k_t\right] = -\frac{\phi}{2}\left[2u_t\left(-\frac{I_t}{k_t^2}\right)k_t + u_t^2\right] = \phi\, u_t\frac{I_t}{k_t} - \frac{\phi}{2}u_t^2. \tag{B.20}$$

From the production function, $\partial y_t/\partial k_t = \theta y_t/k_t$. From the enforcement constraint, holding the aggregate price Q_t fixed, $\partial/\partial k_t = \mu_t\zeta_t Q_t$. Collecting terms and using (B.15),

$$V_{k,t} = q_t\,\Xi_t + \lambda_t\,\theta\frac{y_t}{k_t} + \mu_t\zeta_t Q_t, \qquad \Xi_t \equiv (1-\delta) + \phi\left(\frac{I_t}{k_t}-\delta\right)\frac{I_t}{k_t} - \frac{\phi}{2}\left(\frac{I_t}{k_t}-\delta\right)^2. \tag{B.21}$$

The first two terms are the familiar user-cost and marginal-product components; the third is the collateral dividend, present only because the constraint is written in k_t. Since B_t enters only through the flow of funds,

$$V_{B,t} = -\frac{\lambda_t}{P_t}. \tag{B.22}$$

The factor $1/P_t$ is required because B_t is a nominal face value while the firm's budget is measured in goods; it is the point at which unanticipated inflation revalues the firm's liabilities.

Assembled Euler equations. Leading (B.21) one period and substituting into (B.18) gives (26). Leading (B.22) one period, substituting into (B.19) and multiplying through by P_t,

$$-\mathbb{E}_t\left[m_{t,t+1}\lambda_{t+1}\frac{P_t}{P_{t+1}}\right]+\frac{\lambda_t}{R_t}-\frac{\mu_t\zeta_t}{1+i_t}=0 \quad\Longrightarrow\quad \mathbb{E}_t\left[\frac{m_{t,t+1}\lambda_{t+1}}{\Pi_{t+1}}\right]=\frac{\lambda_t}{R_t}-\frac{\mu_t\zeta_t}{1+i_t}, \tag{B.23}$$

which is (27). Comparing (B.23) with the household's condition (B.5) is instructive and is the basis of both propositions below: households discount the same nominal payoff at $1/(1+i_t)$, while firms value it at $\lambda_t/R_t-\mu_t\zeta_t/(1+i_t)$. The tax shield drives a wedge in one direction and the collateral cost in the other, and the multiplier μ_t is precisely what reconciles them.

B.3 Collateral Valuation

Proof of Lemma 1. Dividing the valuation rule (23) by the investment condition (B.17),

$$Q_t=\frac{q_t}{\lambda_t}=\frac{q_t}{q_t\left[1-\phi\left(I_t/k_t-\delta\right)\right]}=\left[1-\phi\left(\frac{I_t}{k_t}-\delta\right)\right]^{-1}, \tag{B.24}$$

which requires only $q_t\neq 0$. Differentiating, $\partial Q_t/\partial(I_t/k_t)=\phi\left[1-\phi(I_t/k_t-\delta)\right]^{-2}=\phi\,Q_t^2$, which is strictly positive for $\phi>0$. At $I_t/k_t=\delta$ the bracket equals one, so $Q_{ss}=1$. If $\phi=0$ the bracket is identically one and $Q_t\equiv 1$. □

Two admissibility restrictions are implied and are checked numerically at every candidate parameter vector and along every simulated path:

$$\frac{I_t}{k_t}<\delta+\frac{1}{\phi} \qquad\text{and}\qquad d_t>\bar{d}-\frac{1}{2\kappa}, \tag{B.25}$$

the first ensuring a positive denominator in (37) and the second ensuring $\lambda_t>0$ in (B.14). Note also that (37) is not an additional independent equilibrium condition: it is implied by (23) together with (B.17), so imposing all three would over-determine the system.

B.4 Market Clearing and Walras' Law

Let $b_{t+1}\equiv B_{t+1}/P_t$ denote the real value at date t of the face value issued at t, so that the real repayment obligation at $t+1$ is $B_{t+1}/P_{t+1}=b_{t+1}/\Pi_{t+1}$. Written this way, the flow of funds (20) becomes

$$y_t+\frac{b_{t+1}}{R_t}=d_t+\kappa\left(d_t-\bar{d}\right)^2+w_tn_t+I_t+\frac{b_t}{\Pi_t}, \tag{B.26}$$

which makes the debt-deflation channel explicit: an unanticipated increase in Π_t reduces the real burden b_t/Π_t of inherited liabilities and thereby loosens the firm's budget.

Substituting the tax rule $T_t = b_{t+1}/R_t - b_{t+1}/(1+i_t)$ into the symmetric household budget constraint (11) cancels the two terms in $1/(1+i_t)$ and leaves

$$c_t = \big(1 - \Theta_{w,t}\big) w_t n_t + \frac{b_t}{\Pi_t} + d_t - \frac{b_{t+1}}{R_t}. \tag{B.27}$$

Substituting d_t from (B.26), the remaining debt terms and the financing subsidy cancel and we obtain (34),

$$y_t = c_t + I_t + \kappa \left(d_t - \bar{d}\right)^2 + \Theta_{w,t} w_t n_t. \tag{B.28}$$

Resources are therefore absorbed by consumption, investment, and the two adjustment costs, both of which are genuine deadweight losses. The cancellation of nominal terms is an accounting result and does not imply that the nominal block is irrelevant: inflation still enters (B.26), and the policy rate still enters both the pledgeability of collateral in (21) and the effective cost of debt in (B.23).

B.5 The Deterministic Steady State

Set all innovations to zero, $z = 1$, $\zeta = \zeta_{ss}$, and impose the inflation target $\Pi = \Pi^* = 1$. A constant real wage then implies $\Pi_w = 1$ from (28), so $\Theta_w = 0$ and the wage adjustment cost vanishes in steady state. We adopt the standard convention that the payout target $\bar{d}$ equals the steady-state payout, so $d = \bar{d}$ and, from (B.14), $\lambda = 1$. The steady state can then be solved recursively in the following order.

Step 1: rates. From (B.5) with $m = \beta$ and $\Pi = 1$,

$$1 + i_{ss} = \frac{1}{\beta}, \qquad R_{ss} = 1 + i_{ss}(1-\tau) = \frac{1 - \tau + \tau\beta}{\beta}, \qquad \beta R_{ss} = 1 - \tau + \tau\beta. \tag{B.29}$$

Step 2: capital prices. With $I = \delta k$ the investment rate equals δ, so (B.17) gives $q_{ss} = \lambda_{ss} = 1$, and (23) gives $Q_{ss} = 1$. Also $\Xi_{ss} = 1 - \delta$.

Step 3: the multiplier. Evaluating (B.23) at the steady state,

$$\beta = \frac{1}{R_{ss}} - \mu_{ss}\zeta_{ss}\beta \quad \Longrightarrow \quad \mu_{ss} = \frac{1 - \beta R_{ss}}{\zeta_{ss}\,\beta R_{ss}} = \frac{\tau(1-\beta)}{\zeta_{ss}\,(1 - \tau + \tau\beta)}, \tag{B.30}$$

using (B.29). It is convenient to define the *pledgeability-weighted wedge*

$$\Omega \equiv \mu_{ss}\zeta_{ss} = \frac{\tau(1-\beta)}{1 - \tau + \tau\beta}, \tag{B.31}$$

which depends on τ and β alone and is independent of ζ_{ss}.

Step 4: the capital-output ratio. Evaluating (26) at the steady state with $q = \lambda = Q = 1$ and

$\Xi = 1 - \delta$,

$$1 = \beta\left[(1-\delta) + \theta\frac{y}{k} + \Omega\right] \quad \Longrightarrow \quad \frac{y}{k} = \frac{1}{\theta}\left[\frac{1}{\beta} - (1-\delta) - \Omega\right]. \tag{B.32}$$

The collateral dividend Ω acts exactly like a reduction in the required marginal product of capital: because capital is useful as collateral as well as in production, the firm is willing to hold more of it than the frictionless user cost would justify, and the capital-output ratio is correspondingly higher.

Step 5: quantities and the wage. From (17), $y/k = (n/k)^{1-\theta}$, so

$$\frac{n}{k} = \left(\frac{y}{k}\right)^{\frac{1}{1-\theta}}, \qquad \frac{y}{n} = \left(\frac{y}{k}\right)^{-\frac{\theta}{1-\theta}}, \qquad w = \frac{(1-\theta)}{1+\mu_{ss}} \cdot \frac{y}{n}, \tag{B.33}$$

the last from (B.16) with $\lambda_{ss} = 1$. Normalising hours to a target $\bar{n}$ then fixes $k = \bar{n}/(n/k)$, $y = (y/k)k$, $I = \delta k$ and, from (B.28) with $\Theta_w = 0$ and $d = \bar{d}$, $c = y - \delta k$.

Step 6: debt. The enforcement constraint (21) holds with equality, so with $Q_{ss} = 1$ and $\Pi = 1$,

$$b_{ss} = (1 + i_{ss})\left[k - \frac{wn}{\zeta_{ss}}\right] = \frac{1}{\beta}\left[k - \frac{(1-\theta)y}{\zeta_{ss}\,(1+\mu_{ss})}\right]. \tag{B.34}$$

Step 7: the labor-disutility scale. With $\Pi_w = 1$ and $\Theta_w = 0$, (B.12) reduces to $w = \mu_w\,\mathrm{mrs}$, and $\Lambda_{ss} = (1 - \iota\beta)\,[(1-\iota)c]^{-\sigma}$ from (B.3). The scale parameter α is then recovered residually as $\alpha = w\,\Lambda_{ss}/\left(\mu_w \bar{n}^{\nu}\right)$, so that hours equal their target by construction.

Existence and two useful restrictions. Two conditions must hold for this steady state to be well defined, and both are imposed as constraints during estimation rather than checked only at the initial calibration. First, (B.32) requires a positive marginal product of capital,

$$\Omega = \frac{\tau(1-\beta)}{1 - \tau + \tau\beta} < \frac{1}{\beta} - (1-\delta), \tag{B.35}$$

which restricts τ relative to the steady-state real rate and depreciation. Note that this restriction falls on Ω, and hence on τ and β, and not on ζ_{ss}: by (B.31) the collateral dividend is independent of the enforcement parameter, because a lower ζ_{ss} raises the multiplier μ_{ss} in exactly offsetting proportion.

Second, (B.34) requires steady-state debt to be positive. Since $\zeta_{ss}\,(1+\mu_{ss}) = \zeta_{ss} + \Omega$, the requirement $\zeta_{ss}\,(1+\mu_{ss}) > (1-\theta)\,y/k$ is linear in ζ_{ss} and can be stated explicitly as a *lower* bound,

$$\zeta_{ss} \;>\; \underline{\zeta} \;\equiv\; (1-\theta)\frac{y}{k} - \Omega \;=\; \frac{1-\theta}{\theta}\left[\frac{1}{\beta} - (1-\delta) - \Omega\right] - \Omega, \tag{B.36}$$

where the second equality substitutes (B.32). If enforcement is too weak, the working-capital loan alone exhausts pledgeable collateral and the firm cannot support positive intertemporal debt; as $\zeta_{ss} \downarrow \underline{\zeta}$, $b_{ss} \downarrow 0$ and the model collapses to one in which firms are financed entirely by

equity and the intraperiod loan.

Remark 4. It is worth being precise about how (B.36) relates to the restriction on the enforcement parameter in Chatterjee et al. (2025), since the two are easily conflated. Setting $\Omega = 0$ in (B.36) recovers $\frac{1-\theta}{\theta}\left[\frac{1}{\beta} - (1-\delta)\right]$, which is exactly the threshold they derive. The algebraic object is therefore the same, but the inequality runs the other way and for a different reason. In their environment there is no intertemporal debt and the constraint binds only if enforcement is weak enough, so the threshold is an *upper* bound required for $\mu_{ss} > 0$. Here Proposition 1 shows that the tax advantage of debt, not the tightness of enforcement, is what makes the constraint bind: $\mu_{ss} > 0$ for *any* $\zeta_{ss} > 0$ whenever $\tau > 0$. The same threshold instead marks the point below which senior debt can no longer be supported, and so becomes a lower bound. At the calibration of Table 1 it takes the value $\underline{\zeta} \approx 0.061$.

Is there an upper bound? Neither restriction bounds ζ_{ss} from above, and no further condition does either. The natural candidate is that the steady-state payout be positive. Using (B.26) in steady state with $d = \bar{d}$ and $\Pi = 1$,

$$\bar{d} = y - wn - \delta k - b_{ss}\frac{R_{ss}-1}{R_{ss}}, \qquad \frac{R_{ss}-1}{R_{ss}} = \frac{(1-\beta)(1-\tau)}{1-\tau+\tau\beta}, \tag{B.37}$$

so the payout is what remains of output after the wage bill, depreciation and interest. As ζ_{ss} rises, μ_{ss} falls and the wage bill rises toward $(1-\theta)y$, while b_{ss} rises toward k/β; both reduce $\bar{d}$, so (B.37) is strictly decreasing in ζ_{ss}. Substituting the limiting values and using (B.32) gives

$$\lim_{\zeta_{ss}\to\infty} \bar{d} = k\left[\frac{1-\beta}{\beta} - \Omega - \frac{(1-\beta)(1-\tau)}{\beta\,(1-\tau+\tau\beta)}\right] = \frac{k(1-\beta)}{\beta\,(1-\tau+\tau\beta)}\Big[(1-\tau+\tau\beta) - \tau\beta - (1-\tau)\Big] = 0, \tag{B.38}$$

the bracket vanishing identically. The steady-state payout is therefore strictly positive for every $\zeta_{ss} > \underline{\zeta}$ and approaches zero only in the limit of perfect enforcement, in which the firm levers up until the entire operating surplus is absorbed by interest and depreciation. The admissible region is thus the half-line $\zeta_{ss} \in (\underline{\zeta}, \infty)$, and the estimation imposes only (B.35) and (B.36).

A steady-state calibration target. Equations (B.32) and (B.34) imply that the ratio of debt to quarterly output is

$$\frac{b_{ss}}{y} = \frac{1}{\beta}\left[\frac{k}{y} - \frac{1-\theta}{\zeta_{ss}\,(1+\mu_{ss})}\right], \tag{B.39}$$

which is strictly increasing in ζ_{ss} and depends on it alone once the remaining calibrated parameters are fixed. The steady state therefore maps ζ_{ss} one-to-one into observed corporate leverage, which provides a natural over-identifying check on the estimate reported in Table 1: at $\beta = 0.99$, $\tau = 0.35$, $\theta = 0.33$ and $\delta = 0.025$, equation (B.39) implies a debt-to-quarterly-output ratio of roughly six at $\zeta_{ss} = 0.15$, against a value near 0.09 that would reproduce the ratio of 3.36 targeted by Jermann and Quadrini (2012). We report this diagnostic alongside the estimates.

B.6 Proofs of the Results in Section 3.7

Proposition 1. The proof in the text is complete; (B.30) above derives the same expression as part of the recursive characterisation of the steady state. Two comparative statics are worth recording. The wedge μ_{ss} is strictly decreasing in ζ_{ss}, so a single parameter controls the average severity of the friction; and it is strictly increasing in τ, since $\partial\Omega/\partial\tau = (1-\beta)/\left(1-\tau+\tau\beta\right)^2 > 0$. The economics is that the tax shield induces firms to lever up until enforcement stops them, so a larger shield produces a tighter constraint in equilibrium, not a looser one.

Proposition 2 and Remark 1. It is convenient to prove both at once. Suppose $\kappa = 0$, so $\lambda_t \equiv 1$ by (B.14). Then (B.23) reads

$$\mathbb{E}_t\left[\frac{m_{t,t+1}}{\Pi_{t+1}}\right] = \frac{1}{R_t} - \frac{\mu_t\zeta_t}{1+i_t}. \tag{B.40}$$

The household's condition (B.5) gives $\mathbb{E}_t\left[m_{t,t+1}/\Pi_{t+1}\right] = 1/(1+i_t)$ at every date. Equating the two and multiplying by $(1+i_t)$,

$$\mu_t\zeta_t = \frac{1+i_t}{R_t} - 1 = \frac{1+i_t-\left[1+i_t(1-\tau)\right]}{R_t} = \frac{\tau i_t}{R_t} \quad\Longrightarrow\quad \mu_t = \frac{\tau i_t}{\zeta_t R_t}. \tag{B.41}$$

If in addition $\tau = 0$, then $R_t = 1+i_t$ and (B.41) gives $\mu_t\zeta_t = 0$, hence $\mu_t = 0$ whenever $\zeta_t > 0$. With $\mu_t = 0$ the enforcement constraint drops out of (B.16) and (26), and ζ_t disappears from the subsystem that determines real allocations and prices. It remains in the enforcement constraint and can therefore affect the selected debt–payout pair, but those financing choices are payoff-irrelevant when $\kappa = 0$. This proves Proposition 2.

Remark 1 is the case $\kappa = 0$ with $\tau > 0$. Expression (B.41) shows that the multiplier is then determined *statically* by the policy rate, the tax parameter and the enforcement parameter, with no reference to the collateral price Q_t, the capital stock k_t, the wage bill, or the level of debt. When $i_t > 0$, this expression implies $\mu_t > 0$ and the constraint holds with equality, determining debt compatible with the desired wage bill. When $i_t = 0$, $\mu_t = 0$ and the constraint may be slack. A negative i_t with $\tau > 0$ would imply a negative multiplier, so the unrestricted-debt, frictionless-payout conditions do not support that case with a nonnegative ICC multiplier. The result is therefore a statement about admissible equilibria, not a binding assumption. In those equilibria, collateral revaluation does not directly move the wedge. When $\kappa > 0$, λ_t is endogenous and state dependent, and the wedge depends on its expected path through (38). Collateral revaluation can then affect the wedge and the regime through the joint determination of investment, debt and payouts.

B.7 The Complete Equilibrium System

Collecting the conditions above, a symmetric equilibrium is characterised by the following nineteen conditions in the nineteen endogenous variables

$$\{c_t, X_t, \Lambda_t, n_t, w_t, \Pi_{w,t}, \Pi_t, i_t, R_t, y_t,$$
$$k_{t+1}, I_t, d_t, \lambda_t, q_t, Q_t, \mu_t, b_{t+1}, T_t\},$$

given the exogenous processes $\{z_t, \zeta_t, h_t, \chi_t, a_t\}$ and the predetermined variables $\{c_{t-1}, k_t, b_t, w_{t-1}\}$.

1. $X_t = (c_t - \iota c_{t-1})^{-\sigma}$
2. $\Lambda_t = X_t - \iota\beta\,\mathbb{E}_t X_{t+1}$
3. $1 = (1 + i_t)\,\mathbb{E}_t\left[m_{t,t+1}/\Pi_{t+1}\right]$, with $m_{t,t+1} = \beta\Lambda_{t+1}/\Lambda_t$
4. the wage Phillips curve (B.12)
5. $y_t = z_t k_t^{\theta} n_t^{1-\theta}$
6. $k_{t+1} = (1-\delta)k_t + I_t - \frac{\phi}{2}\left(I_t/k_t - \delta\right)^2 k_t$
7. $\lambda_t = \left[1 + 2\kappa\left(d_t - \bar{d}\right)\right]^{-1}$
8. $w_t\left(1 + \mu_t/\lambda_t\right) = (1-\theta)y_t/n_t$
9. $\lambda_t = q_t\left[1 - \phi\left(I_t/k_t - \delta\right)\right]$
10. $Q_t = q_t/\lambda_t$
11. the capital Euler equation (26)
12. the debt Euler equation (B.23)
13. $S_t \equiv \zeta_t[Q_t k_t - b_{t+1}/(1+i_t)] - w_t n_t \geq 0$, $\mu_t \geq 0$, $\mu_t S_t = 0$, with S_t an algebraic slack
14. the flow of funds (B.26)
15. $R_t = 1 + i_t(1-\tau)$
16. $\Pi_t = \Pi_{w,t}\, w_{t-1}/w_t$
17. the policy rule (29)
18. $T_t = b_{t+1}/R_t - b_{t+1}/(1+i_t)$
19. goods market clearing (B.28)

Four accounting relations are available—the household budget constraint, the flow of funds, the government budget, and goods market clearing—of which any three imply the fourth. The list above includes the last three, so the household budget constraint is the redundant one and is

omitted, exactly as Appendix B.4 shows. Equation (37) is likewise omitted, being implied by items 9 and 10.

Item 13 is a complementarity condition, not an equality imposed at every date. The equality is recovered in binding states; slack states have $\mu_t = 0$. The DEQN enforces this condition with (35) while allowing future regime changes to enter expectations. Proposition 1 establishes binding only at the deterministic steady state. Technology is stationary in (30), so no balanced-growth transformation is needed.

B.8 Implementation of the Productivity Signal and Impulse Responses

Signal state and timing. The network conditions on today's fully resolved state: any maturing event has already been confirmed or not confirmed and its delivery incorporated into z_t. Given (χ_t, a_t), draw confirmation F_{t+1}, arrival $A_{t+1} \sim \text{Bernoulli}(\lambda_{\text{news}})$, $H_{t+1} \sim 8 + \text{Binomial}(8, p_H)$, and a new magnitude $\tilde{a}_{t+1} \sim N(0, \sigma_0^2)$ independently. Then

$$(\chi_{t+1}, a_{t+1}) = \begin{cases} (\chi_t - 1, a_t), & \chi_t > 1, \\ (H_{t+1}, \tilde{a}_{t+1}), & \chi_t \leq 1,\ A_{t+1} = 1, \\ (0, 0), & \chi_t \leq 1,\ A_{t+1} = 0. \end{cases} \tag{B.42}$$

Delivery at $t+1$ is $D_{t+1} = \mathbf{1}\{\chi_t = 1\} F_{t+1} a_t$, independent of any new signal at that date. Thus a countdown of one resolves in the next state, not before today's decision. The same transition is used for simulation and expectations. Discrete arrival, horizon and confirmation outcomes are enumerated with their probabilities; new magnitudes are integrated by Gauss–Hermite quadrature. Two states replace the fixed-horizon relay chain used in the preliminary perturbation calculations.

The two experiments. The signal intervention sets $(\chi_1, a_1) = (H, \bar{a})$ with $\bar{a} > 0$. The noise and genuine-news experiments condition respectively on $F_{1+H} = 0$ and $F_{1+H} = 1$ for this signal. They use the same decision rule and confirmation probability during anticipation; the conditioned outcome is not revealed to agents early. Under paired background draws their paths coincide before maturity. A noise realization delivers no productivity increment; confirmation delivers $\bar{a}$. Without background technology innovations or additional signals, the noise experiment leaves productivity unchanged at every date.

The macro-uncertainty experiment sets $\varepsilon_{h,1} > 0$ in (33) while pairing the underlying draws for enforcement and all other processes with the baseline. Enforcement draws are transformed using each path's volatility and the positive-support truncation. This preserves the implemented transition rather than artificially fixing realised enforcement.

Computing impulse responses. The nonlinear responses are generalised impulse responses. Starting states are sampled from long simulations with the DEQN. For each state, shocked and baseline paths use identical underlying random draws. Responses average their differences over both initial states and future innovations. Signal experiments use eligible starting states with no pending signal and condition on the specified maturity outcome. We report simulation counts, Monte Carlo uncertainty and domain exits. Absolute multiplier and slack paths are checked separately from their responses relative to baseline, so a zero multiplier response is not mistaken for a slack constraint. No pruning is applied to the DEQN simulations.

B.9 DEQN Solution and Validation

The fixed-parameter network has nine standardised inputs, two hidden layers of 128 units each, and six outputs. Each hidden layer uses 64 tanh and 64 softplus units. The output coordinates recover the full allocation through production, the policy and tax rules, payout optimality, capital accumulation, resource clearing, and the joint financing and investment conditions. Collateral slackness is imposed through (35). These algebraic conditions hold by construction; they do not establish that the remaining equilibrium conditions have been solved.

Following Azinovic et al. (2022) and Scheidegger (2026, Chapter 2), the loss sums scaled squared residuals of the six conditions described in Section 3.6. Future allocations use the same network, evaluated at states obtained from current choices and the exogenous transitions. Expectations are weighted averages of full nonlinear expressions, not products of separate expectations. Signal-resolution branches are enumerated, and continuous shocks use Gauss–Hermite quadrature, including the transformation for positive enforcement. Training uses 3 volatility, 5 enforcement, 3 productivity and 3 signal-magnitude nodes; validation raises each count to 7.

The network is initialised from the second-order perturbation solution. Residual scales are fixed using perturbation errors on a mixture of near-steady-state, simulated and signal states. Soft level and slope anchors use the deterministic steady state and first-order solution. Adam updates the weights, with the enforcement innovation scale and stochastic volatility introduced gradually. Training states combine near-steady-state draws, persistent simulated paths, targeted signal states, states near the collateral switching boundary, and a wider state-space sample. These choices follow the solution procedure in Scheidegger (2026, Chapters 2–3); they impose no desired impulse-response signs or binding dates.

Acceptance requires equation-by-equation residual checks on unseen ordinary, simulated and signal states with finer quadrature, comparisons with perturbation on the same states, and checks of local slopes, feasibility and the redundant household budget. Long-horizon simulations track individual paths and report domain exits and capital and debt distributions. Binding shares are reported by state type. For the parameterised network, these checks cover interior, edge and corner parameter vectors before estimation is allowed over that region. Solver errors are checked

again at the estimate using a separately trained fixed-parameter network. Small training loss or exact accounting alone is not an acceptance criterion.

B.10 Notation

Table B.1: Notation

Symbol	Definition
Quantities	
c_t	Real consumption
n_t	Aggregate hours
y_t	Real output
I_t	Real investment
k_t	Capital stock, predetermined at t
d_t	Real equity payout; $\bar{d}$ its steady-state target
b_{t+1}	Real value of newly issued debt, B_{t+1}/P_t
B_t	Nominal face value of debt due at t
T_t	Real lump-sum tax
Prices and nominal variables	
w_t	Real wage, W_t/P_t
W_t	Aggregate nominal wage index
P_t	Price level
Π_t	Gross price inflation, P_t/P_{t-1}
$\Pi_{w,t}$	Gross nominal wage inflation, W_t/W_{t-1}
i_t	Net nominal policy rate
R_t	Firm's effective gross nominal borrowing rate, $1 + i_t(1-\tau)$
Q_t	Contractual collateral price in goods, q_t/λ_t
Multipliers and auxiliary objects	
Λ_t	Household marginal utility of wealth
X_t	Flow marginal felicity, $(c_t - \iota c_{t-1})^{-\sigma}$
$m_{t,t+1}$	Stochastic discount factor, $\beta\Lambda_{t+1}/\Lambda_t$
mrs_t	Marginal rate of substitution, $\alpha n_t^{\nu}/\Lambda_t$
λ_t	Multiplier on the flow of funds
q_t	Multiplier on capital accumulation
φ_t	Multiplier on the production function; $\varphi_t = \lambda_t$
μ_t	Multiplier on the enforcement constraint
S_t, ψ_t	ICC slack and signed DEQN collateral coordinate

Table B.1, continued

Symbol	Definition
Ξ_t	User-cost term defined in (B.21)
$\Theta_{w,t}$	Wage adjustment cost as a share of the wage bill
Ω	Steady-state pledgeability-weighted wedge, $\mu_{ss}\zeta_{ss}$
Exogenous processes	
z_t	Total factor productivity
ε_t^0, D_t	Productivity surprise and signal delivery
χ_t, a_t	Remaining signal horizon and pending magnitude
H, F_t	Signal horizon and confirmation indicator
ζ_t	Enforcement parameter
h_t	Log volatility of innovations to ζ_t
$\varepsilon_{\zeta,t},\ \varepsilon_{h,t}$	Enforcement and log-volatility innovations
Parameters	
$\beta, \sigma, \iota, \nu, \alpha$	Discount factor, curvature, habit, inverse Frisch, disutility scale
η_w, μ_w	Elasticity of substitution across labor varieties; implied markup
ϕ_w	Rotemberg wage adjustment cost
θ, δ	Capital share, depreciation rate
ϕ	Investment adjustment cost
κ	Payout adjustment cost
τ	Tax advantage of debt
$\zeta_{ss}, \rho_\zeta, \bar{h}, \rho_h, \sigma_h$	Enforcement and volatility process parameters
ρ_z, σ_0	Productivity persistence and common surprise/news scale
$\lambda_{\text{news}}, p_H, \pi_{\text{confirm}}$	News arrival, horizon and confirmation parameters
ϕ_π, ϕ_y, Π^*	Policy rule coefficients and inflation target