Online Appendix

Imperfect Information and Slow Recoveries in the Labor Market

Anushka Mitra *

Federal Reserve Board

This draft: May 2026

A Empirical Appendix

This appendix provides empirical material that complements the main text. It (i) documents the data sources and construction of key time series, (ii) summarizes additional evidence on labor market recoveries and forecast behavior, (iii) reports further SVAR, local-projection, and historical-decomposition results, and (iv) presents robustness checks, including controls for uncertainty shocks. The theoretical appendix (Section B) develops the information structure and Bayesian updating problem, characterizes the search-and-matching model with imperfect information, and describes the solution method and analytical results that underpin the quantitative exercises.

A.1 Data Sources

The baseline sample period for the empirical analysis is 1968Q4–2019Q4, unless noted otherwise.

1. **Unemployment and worker flows.** The unemployment rate and the unemployment-to-employment (UE) transition rate are constructed from the Current Population Survey (CPS).

*anushka.mitra@frb.gov.

2. **Job-to-job transitions.** The CPS redesign in 1994 introduced a question that permits identifying whether a respondent is at the same job as in the previous month. Consequently, job-to-job (EE) transitions are observable only from 1994 onward; analyses involving EE transitions therefore use the restricted sample 1994–2019.

3. **Vacancies.** Vacancies are measured using the composite help-wanted index constructed by Barnichon (2010), which is available for the full baseline sample.

4. **Output and wages.** Real GDP is from the Bureau of Economic Analysis (BEA). Average hourly earnings are from the Bureau of Labor Statistics (BLS).

5. **Productivity.** Aggregate productivity is measured using Fernald's utilization-adjusted Solow residual series (Fernald, 2014), as maintained and updated by the Federal Reserve Bank of San Francisco.

6. **Nowcast errors.** Nowcast errors for GDP are constructed using the median forecast from the Survey of Professional Forecasters (SPF), available beginning in 1968Q4. Throughout, a "nowcast error" refers to the realized outcome minus the corresponding SPF forecast at a given horizon; the sign convention is held fixed across exercises.

A.2 Recovery Patterns in the Labor Market

Labor market recoveries typically lag output recoveries. Main-text Figure **??** and Table **??** document this pattern for unemployment. Figure A1 provides the analogous exercise for vacancies. The recoveries following the 1990–91 recession, the 2001 recession, and the Great Recession (2007–09) are notably slower than earlier postwar recoveries: unemployment remains elevated and vacancy creation is subdued for several quarters after output begins to recover.

To quantify these patterns, I compute the number of quarters required for unemployment to reverse a given fraction of its recessionary increase. Let u^{peak} denote unemployment at the cyclical peak during a recession and u^{pre} its level immediately prior to the recession. Let t_{trough} be the NBER trough date. Define the "recovery share" at horizon h as

$$\text{share}(h) \equiv \frac{u^{\text{peak}} - u_{t_{\text{trough}}+h}}{u^{\text{peak}} - u^{\text{pre}}}.$$

I report the horizons at which share(h) reaches 25%, 50%, 75%, and 100% in Table **??**.

Figure A1: Vacancy Recovery Across Recessions

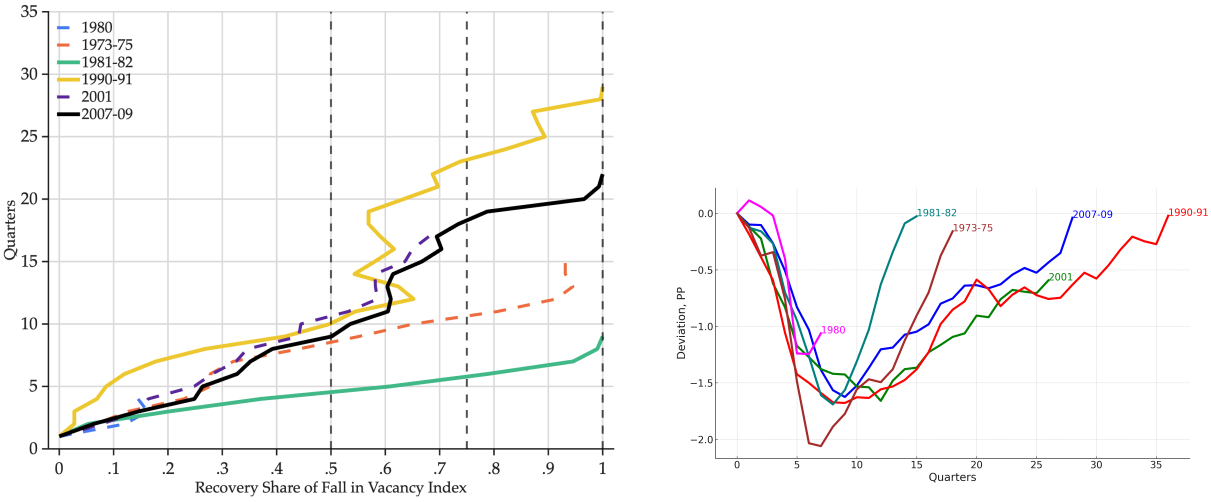

(a) Recovery-share representation

(b) Vacancy deviations from peak

Note: Panel (a) mirrors the main-text unemployment-recovery representation and plots the share of the recessionary vacancy decline that has been reversed over time. Vertical dashed lines mark 50%, 75%, and 100% recovery thresholds. Panel (b) reports the vacancy paths (deviations from recession peak) by episode. Vacancy data are based on the Barnichon (2010) composite help-wanted index.

Job-finding and separation recovery dynamics. I run the same recovery-horizon exercise for both the job-finding and separation rates using CPS series from BLS. I first construct a monthly job-finding probability using the standard Shimer (2012) approximation,

$$f_t = 1 - \frac{U_{t+1} - U_{t+1}^{<5w}}{U_t},$$

where U_t is total unemployment (BLS series LNS13000000) and $U_t^{<5w}$ is short-duration unemployment, less than 5 weeks (BLS series LNS13008396). I then recover the monthly separation rate from the unemployment law of motion using the unemployment rate u_t (BLS series LNS14000000):

$$u_{t+1} = u_t + s_t(1 - u_t) - f_t u_t \quad \Longrightarrow \quad s_t = \frac{u_{t+1} - u_t + f_t u_t}{1 - u_t}.$$

I then average to quarterly frequency and define recovery shares as

$$\text{share}^f(h) = \frac{f_{t_{\text{trough}}+h} - f^{\text{trough}}}{f^{\text{pre}} - f^{\text{trough}}},$$

Table A1: Job-Finding and Separation Recovery Across Recessions

Recessions	(1) 25%	(2) 50%	(3) 75%	(4) 100%	(5) NBER
Panel A: Job-Finding Rate					
2007-09	10	18	27	36	25
2001	5	10	NA	NA	40
1990-91	3	16	16	24	31
1981-82	2	3	5	8	4
1973-75	7	11	13	NA	12
1969-70	5	8	NA	NA	36
Average:					
Total	5	11	15	23	25
Pre 2000	4	10	11	16	21
Post 2000	8	14	27	36	33
Panel B: Separation Rate					
2007-09	1	2	7	18	25
2001	2	5	8	17	40
1990-91	4	4	9	11	31
1981-82	1	2	4	4	4
1973-75	2	11	NA	NA	12
1969-70	5	8	NA	NA	36
Average:					
Total	3	5	7	13	25
Pre 2000	3	6	7	8	21
Post 2000	2	4	8	18	33

Note: Panel A reports the number of quarters required for the quarterly CPS-based job-finding rate to recover 25%, 50%, 75%, and 100% of its recessionary decline. Panel B reports the same thresholds for the reversal of the recessionary increase in the separation rate. The job-finding rate is constructed from BLS CPS series LNS13000000 and LNS13008396 using the Shimer-style monthly approximation; the separation rate is recovered using BLS unemployment rate series LNS14000000 and the unemployment law of motion. Both series are averaged to quarters. NA means the threshold is not reached before the next recession starts.

and

$$\text{share}^{s}(h) = \frac{s^{\text{peak}} - s_{t_{\text{trough}}+h}}{s^{\text{peak}} - s^{\text{pre}}},$$

where f^{pre} and s^{pre} are the quarter before the recession starts, f^{trough} is the minimum job-finding rate between recession start and the next recession start, and s^{peak} is the corresponding maximum separation rate. Table A1 reports the quarters needed to recover 25%, 50%, 75%, and 100%.

Figure A2: Unemployment Projections from the Survey of Professional Forecasters

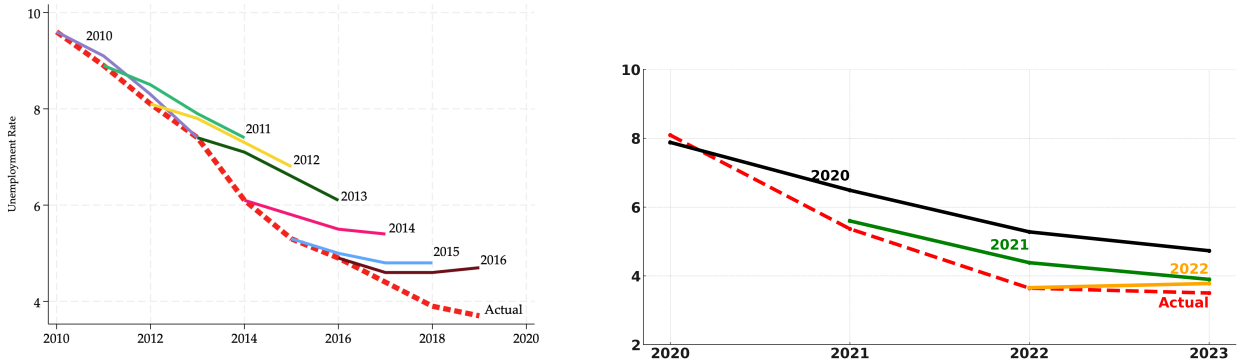

(a) Great Recession (b) COVID-19

Note: Colored lines show median SPF projections at 1–3 year horizons; the solid red line is realized unemployment.

A.3 Expectations of Forecasters and Policymakers

A standard assumption in many macroeconomic models is that agents immediately infer the nature of aggregate disturbances. In practice, forecasts and policy projections often exhibit sizable and persistent errors, consistent with gradual learning about the underlying shocks. This subsection documents forecast patterns for unemployment using professional forecasters and policymakers.

A.3.1 Professional forecasters

Survey of Professional Forecasters. Figure A2 reports SPF median unemployment projections at horizons of 1–3 years during the Great Recession and the COVID-19 episode. Two qualitative facts stand out. First, long-horizon projections adjust gradually over the recovery. Second, forecast errors remain nontrivial for multiple quarters, indicating that forecasters do not instantaneously infer the persistence of the driving shocks.

The SPF does not report a consistent panel of long-horizon unemployment projections prior to 2009. For earlier episodes, Figure A5a therefore reports 1-year-ahead forecasts, which also display persistent forecast errors around recessions. Because the panel includes both recovery and expansion phases, it also serves as a direct check against asymmetric-loss explanations: overprediction is concentrated around recoveries and does not persist systematically during expansions.

Livingston Survey. The Livingston Survey (fielded semiannually, beginning in 1946) provides an additional, longer-horizon perspective on professional forecasts. Figure A3 shows 1- and 2-year-ahead median unemployment forecasts for the 1973–75 and 2007–09

Figure A3: Unemployment Projections from the Livingston Survey

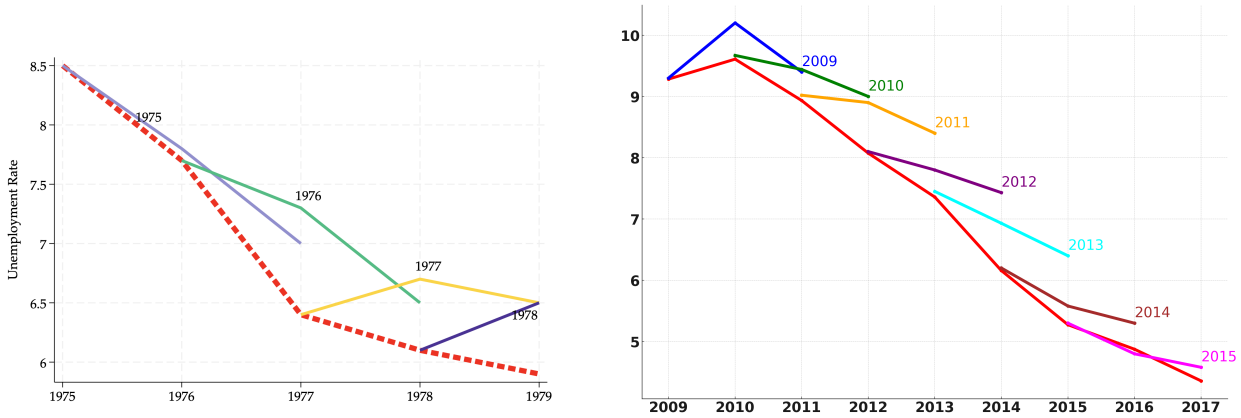

(a) 1973–75 recession

(b) Great Recession

Note: Colored lines show median Livingston forecasts at 1–2 year horizons; the solid red line is realized unemployment.

recessions. In both cases, forecasts adjust gradually over the recovery, consistent with imperfect real-time inference about shock persistence.

Natural rate misperception. Figure A4 evaluates whether the systematic overprediction of unemployment during recoveries can be attributed to misperception about the natural rate rather than the persistence of cyclical shocks. For each SPF forecast vintage, the dark bars report the standard forecast error $u_{t+h} - u^f_{t,h}$, while the orange bars report a gap-adjusted error $(u_{t+h} - u^*_{t+h}) - (u^f_{t,h} - \bar{u}^{LR}_t)$, where $\bar{u}^{LR}_t$ is the SPF median long-run unemployment rate (the "UBAR" variable) and u^*_{t+h} is the CBO natural rate (NROU) at horizon h. The gap-adjusted measure removes the component of the forecast error attributable to forecasters anchoring to a natural rate that differs from the CBO benchmark. The remaining residual can therefore be interpreted as the portion of the forecast error not accounted for by natural-rate misperception.

The adjustment largely eliminates the mean overprediction in the 2009 recovery, has little effect in 2001, and amplifies overprediction in 2020 when $\bar{u}^{LR}_t < u^*_{t+h}$. On average across recoveries, the natural-rate adjustment reduces the mean forecast error by about one-third, but does not eliminate it. This suggests that natural-rate misperceptions explain part, but not all, of forecasters' overprediction during recoveries, leaving a sizable residual consistent with misperceived cyclical persistence. Because the SVAR is identified from GDP nowcast errors rather than unemployment forecast errors, this alternative mechanism does not mechanically affect the identification in Section **??**.

Figure A4: Unemployment Forecast Errors: Natural Rate Adjustment

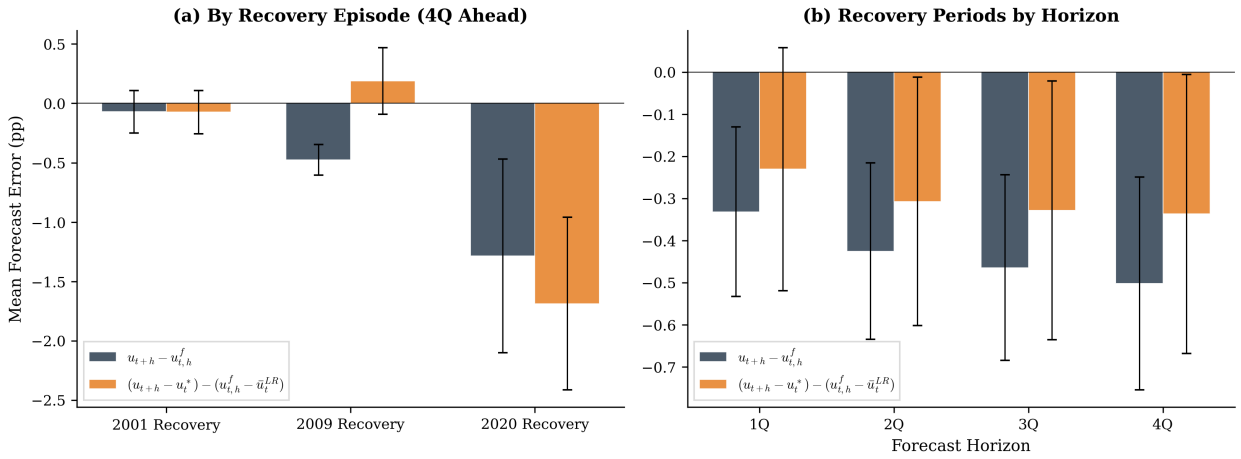

Note: Dark bars report the mean forecast error $u_{t+h} - u^f_{t,h}$ for SPF median unemployment forecasts; orange bars report the gap-adjusted error $(u_{t+h} - u^*_{t+h}) - (u^f_{t,h} - \bar{u}^{LR}_t)$, where $\bar{u}^{LR}_t$ is the SPF median long-run unemployment rate (UBAR) and u^*_{t+h} is the CBO natural rate (NROU) at horizon h. Recovery periods are defined as the 12 quarters following each NBER trough. Panel (a) reports 4-quarter-ahead errors by recovery episode. Panel (b) pools all post-1996 recovery quarters by forecast horizon. Whiskers show 95% confidence intervals. Negative values indicate overprediction. Sample: 1996Q3–2023Q4 (restricted to quarters with UBAR data available).

A.3.2 Policymakers

Federal Reserve Greenbook. To illustrate that real-time inference is challenging even for policymakers, Figure A5 reports unemployment projections from the Federal Reserve Board's December 2008 Greenbook under alternative scenarios. The projections imply a substantially faster recovery than realized outcomes, underscoring the scope for misperception about shock persistence during the Great Recession.

A.4 SVAR: Additional Results

This subsection reports additional SVAR results and robustness. Figure A6 plots the identified shock series.

A.4.1 COVID-19 recession

The COVID-19 recession features an unusually rapid labor market recovery: unemployment rose from 3.5% (February 2020) to 14.7% (April 2020) and declined to 3.9% by end-2021, with a large role for temporary layoffs. To assess real-time beliefs, I examine SPF forecasts at short horizons (1–4 quarters ahead; Figure A7a) and longer horizons (1–3 years ahead; Figure A2b). Forecasts adjust rapidly by 2021, consistent with a comparatively transitory interpretation of the shock.

Figure A5: Unemployment Projections: SPF and Greenbook

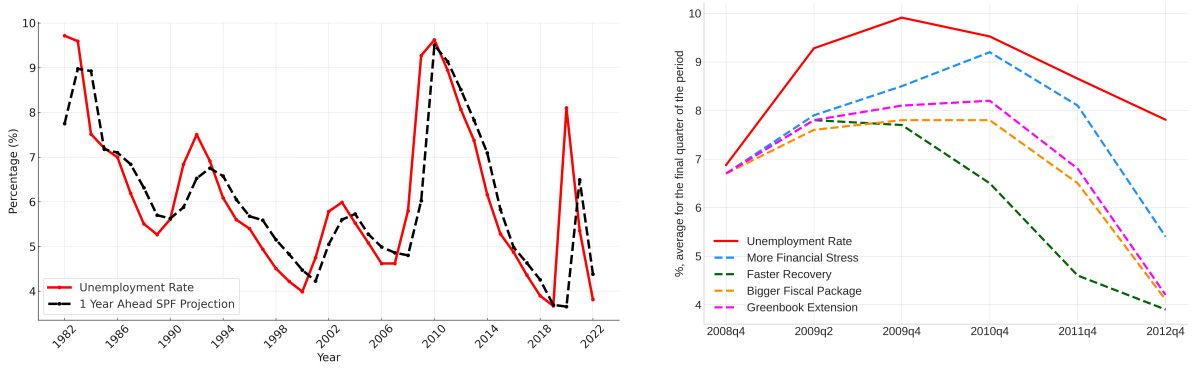

(a) SPF: 1-year-ahead forecast

(b) Greenbook: December 2008

Note: Panel (a) plots the SPF median 1-year-ahead unemployment forecast (dashed) and realized unemployment (solid), covering both recession/recovery and expansion periods; overprediction is concentrated in recoveries rather than expansions. Panel (b) plots Greenbook unemployment projections under alternative scenarios (dashed) and realized unemployment (solid).

To quantify shock contributions, I extend the SVAR sample through the pandemic, 1968Q4–2022Q4. Figure A7b reports the historical decomposition for unemployment. In this extension, fluctuations during COVID-19 are predominantly accounted for by identified TFP shocks, with a smaller role for noise shocks, consistent with rapid expectation adjustment.

VAR impulse responses. Figure A8 reports SVAR impulse responses. The imposed restrictions ensure that noise shocks do not move measured TFP on impact, while affecting forecast errors and real activity. The resulting dynamics are consistent with gradual learning: forecast errors remain persistent following identified shocks, indicating that agents do not instantaneously infer the underlying shock type.

Distinguishing noise from news shocks. A natural concern is whether the identified noise shocks instead capture news shocks—advance information about future productivity—rather than expectational errors. While Chahrour and Jurado (2018) emphasize that news and noise can be observationally similar in VARs that include expectations, the identification here is designed to isolate a non-fundamental information disturbance in the spirit of Blanchard, L'Huillier, and Lorenzoni (2013). I treat the long-run TFP shock as the fundamental innovation that accounts for low-frequency movements in measured TFP, and classify residual innovations to the signal and nowcast error as non-fundamental. The key restriction is that the noise shock has zero contemporaneous impact on TFP and,

Figure A6: Identified Shock Series

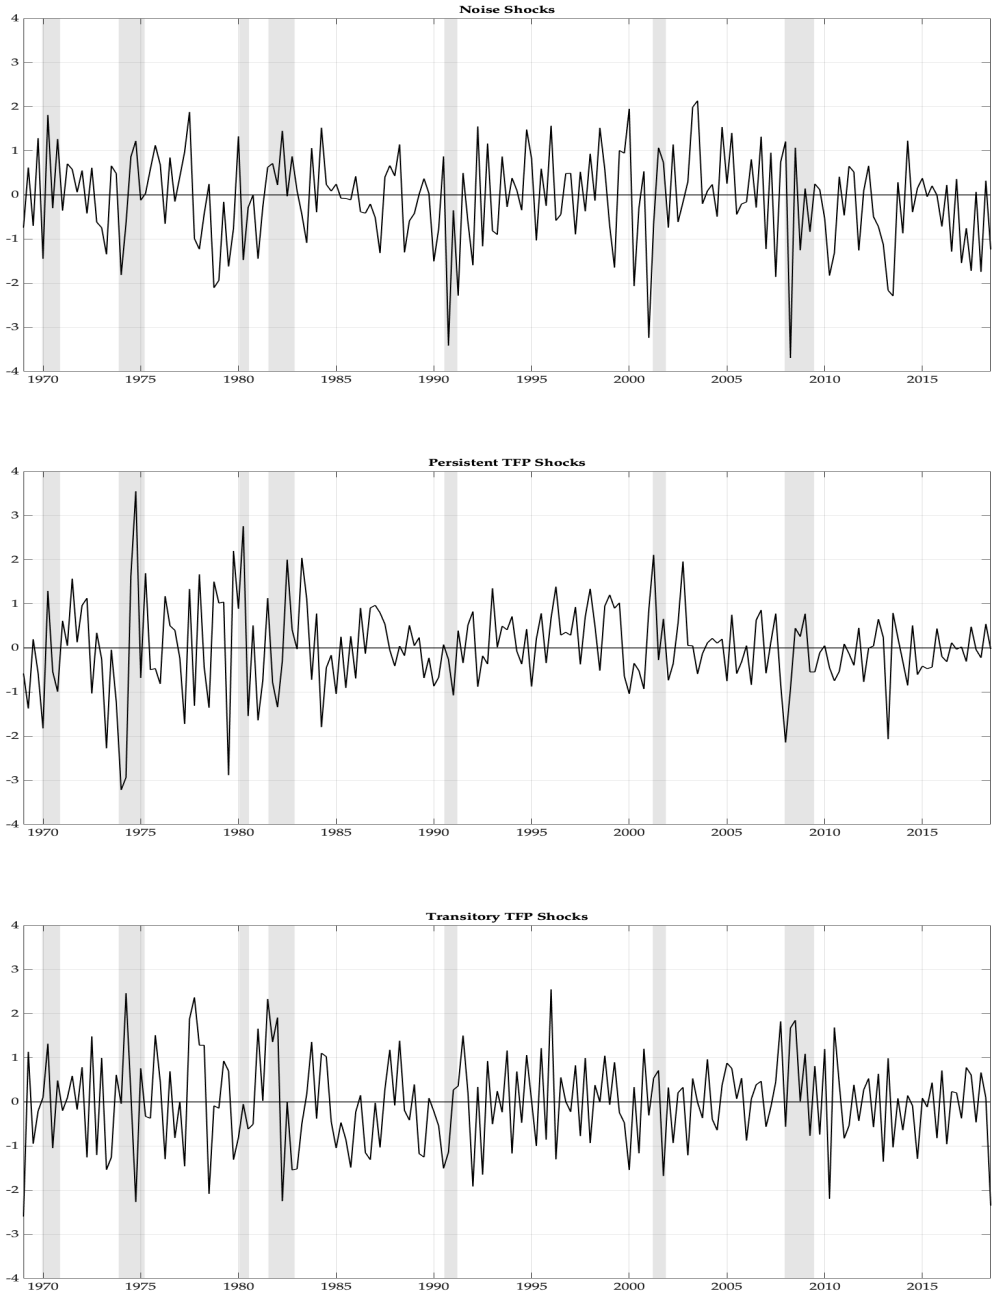

Note: Time series of identified noise shocks, persistent TFP shocks, and transitory TFP shocks.

Figure A7: COVID-19: Projections and Contributions of Shocks

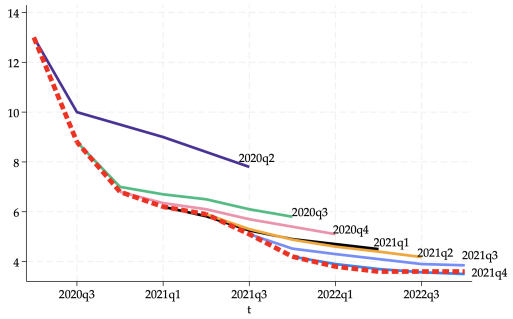

(a) SPF: 1–4-quarter-ahead projections

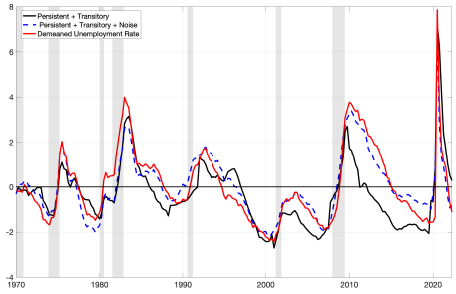

(b) Historical decomposition (SVAR extension)

Note: Panel (a) reports SPF median short-horizon unemployment projections during COVID-19. Panel (b) reports the unemployment historical decomposition in the SVAR extension through 2022Q4.

Figure A8: Impulse Responses from the SVAR

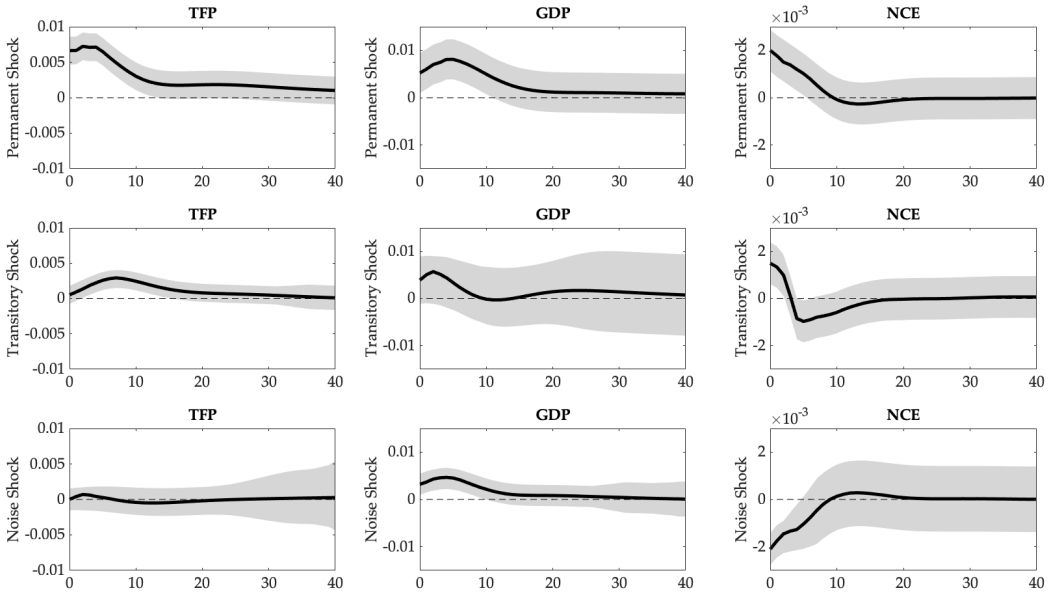

Note: Impulse responses to persistent, transitory, and noise shocks identified by the SVAR described in equation (**??**) in the main text. Baseline sample: 1968Q4–2019Q4.

Figure A9: FEVD after Max-Share Identification

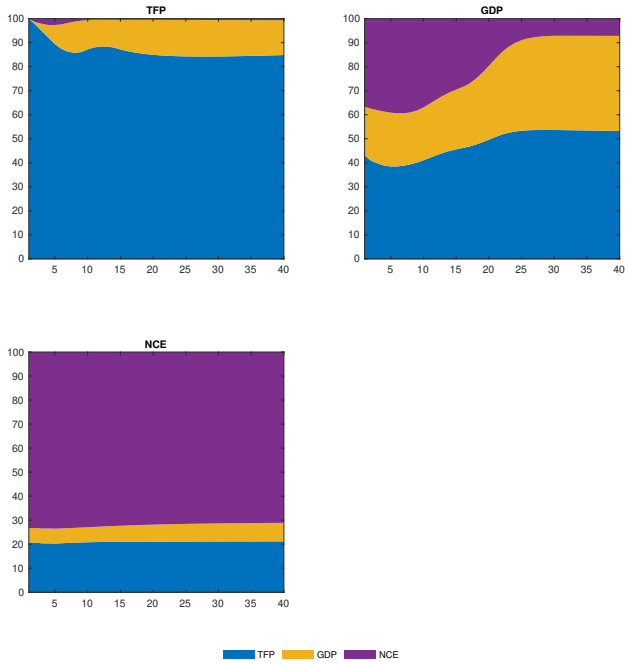

Note: FEVD for TFP, GDP, and nowcast errors. Colors denote contributions of persistent, transitory, and noise shocks.

crucially, does not generate systematic medium-run movements in TFP in the estimated impulse responses. Under the canonical news interpretation, an expectations shock should eventually be reflected in realized productivity. In contrast, the estimated noise shock moves forecast errors, GDP, and labor market variables but does not generate systematic subsequent movements in TFP, reinforcing the view that it reflects misperceptions about fundamentals rather than genuine news about future productivity. Taken together, the long-run TFP identification and the TFP impulse responses support interpreting the identified disturbance as noise rather than news, even though observational equivalence remains a relevant concern in principle.

Figure A9 reports the forecast error variance decomposition (FEVD) for TFP, GDP, and nowcast errors. The persistent shock accounts for the largest share of long-horizon TFP forecast error variance by construction; noise shocks account for a substantial share of forecast error variance for nowcast errors and a nontrivial share for GDP at short horizons.

A.5 Smooth Local Projections

This subsection summarizes the smooth local projection (SLP) estimator following Barnichon and Brownlees (2019). Let y_t denote an outcome and u_t a shock series. Standard

local projections estimate, for each horizon h,

$$y_{t+h} = \alpha_h + \beta_h u_t + \sum_{p=1}^{P} \delta_{p,h} w_{t-p} + \mu_{t,h}, \tag{1}$$

where w_{t-p} collects lagged controls (e.g. lags of y_t and other observables).

SLP imposes smoothness across horizons by approximating the horizon-dependent coefficients using a B-spline basis $\{B_k(h)\}_{k=1}^{K}$:

$$\alpha_h \approx \sum_{k=1}^{K} a_k B_k(h), \qquad \beta_h \approx \sum_{k=1}^{K} b_k B_k(h), \qquad \delta_{p,h} \approx \sum_{k=1}^{K} c_{p,k} B_k(h).$$

Substituting into (1) yields

$$y_{t+h} \approx \sum_{k=1}^{K} a_k B_k(h) + \left(\sum_{k=1}^{K} b_k B_k(h) \right) u_t + \sum_{p=1}^{P} \left(\sum_{k=1}^{K} c_{p,k} B_k(h) \right) w_{t-p} + \mu_{t,h}. \tag{2}$$

Estimation uses ridge regularization, with the shrinkage parameter selected by K-fold cross-validation. Figures A11 and A10 compare LP and SLP impulse responses; the SLP responses are close to LP responses while reducing sampling noise across horizons.

A.6 Additional Empirical Results

A.6.1 Impulse response to transitory TFP shocks

Figure A12 reports impulse responses to transitory TFP shocks. Relative to noise shocks, responses are smaller in magnitude and less persistent.

A.6.2 Historical decomposition

Let the reduced-form VAR be

$$y_t = A_1 y_{t-1} + \cdots + A_p y_{t-p} + e_t,$$

with structural shocks u_t satisfying $e_t = Q u_t$ and $\mathbb{E}[u_t u_t'] = I$. The corresponding vector moving-average (VMA) representation is

$$y_t - \mathbb{E}[y_t] = \sum_{j=0}^{\infty} \Psi_j e_{t-j} = \sum_{j=0}^{\infty} \Theta_j u_{t-j}, \qquad \Theta_j \equiv \Psi_j Q. \tag{3}$$

Figure A10: Smoothed Impulse Response to Persistent TFP Shocks

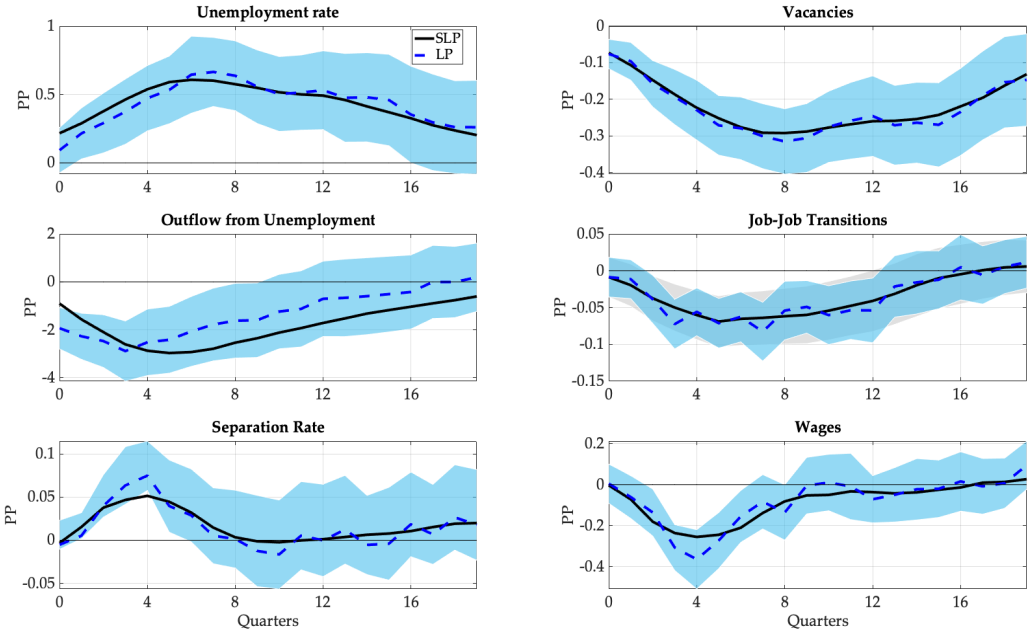

Note: LP and SLP impulse responses for the persistent TFP shock. Shaded area: 95% LP confidence interval.

Figure A11: Smoothed Impulse Response to Noise Shocks

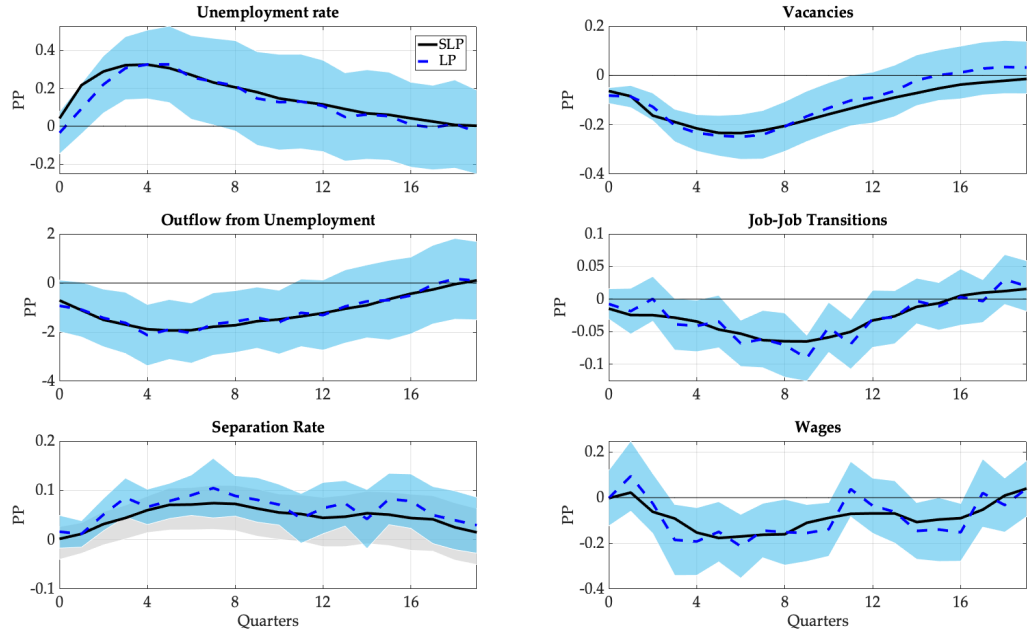

Note: LP and SLP impulse responses for the noise shock. Shaded area: 95% confidence interval.

Figure A12: Impulse Response to Transitory TFP Shocks

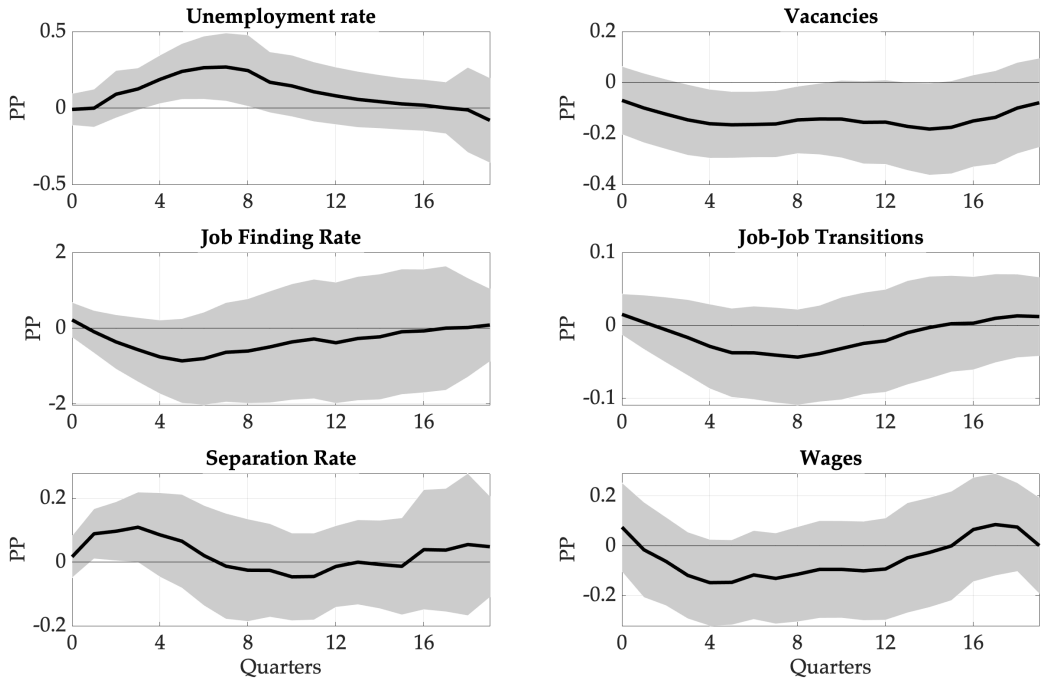

Note: Impulse responses of labor market variables to a transitory TFP shock. Shaded area: 95% confidence interval.

For a finite sample, the historical decomposition writes each component of $y_t - \mathbb{E}[y_t]$ as the cumulative contribution of each structural shock:

$$y_t - \mathbb{E}[y_t] = \sum_{\ell=1}^{n} \sum_{j=0}^{t-1} \Theta_j^{(\ell)} u_{t-j}^{(\ell)}, \tag{4}$$

where $\Theta_j^{(\ell)}$ denotes the column of Θ_j associated with shock ℓ.

Figure A13 shows the decomposition for unemployment. Noise shocks contribute materially to unemployment fluctuations in the 1990–91, 2001, and 2007–09 recessions. Figures A14a and A14b report analogous decompositions for the job-finding rate and vacancy posting.

A.6.3 Persistence of unemployment

To quantify persistence across recessions, I compute the number of quarters required for unemployment to reverse 50% of its recessionary increase (as in Section A.2). Using the historical decomposition, I construct counterfactual unemployment paths attributed to each structural shock and apply the same recovery metric to each path. Table A2 reports

Figure A13: Historical Contribution of Shocks to Unemployment

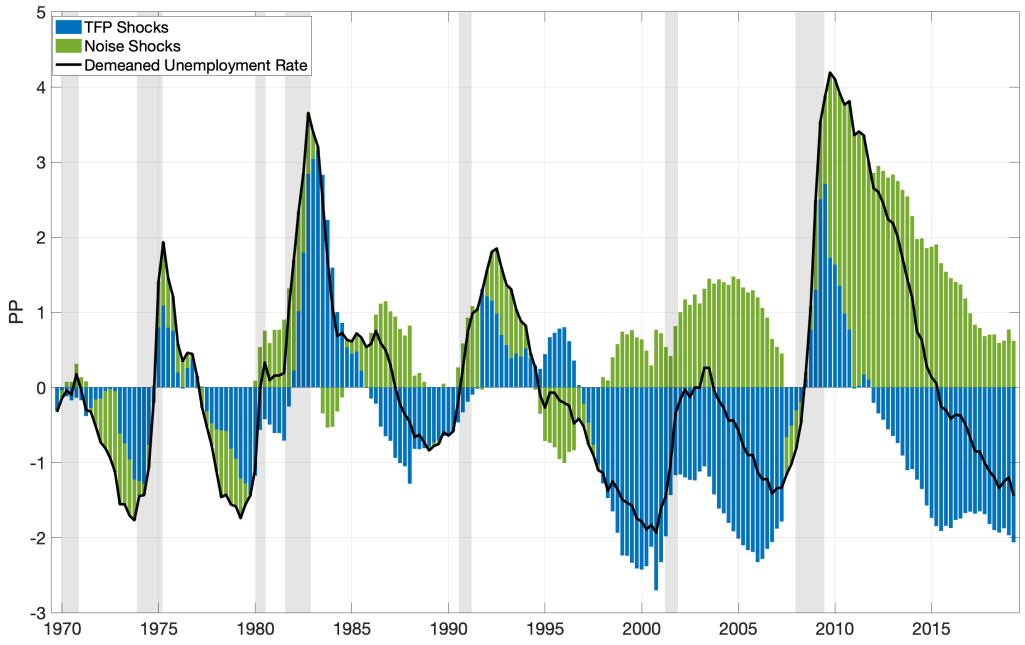

Note: Historical decomposition based on (4). Bars report cumulative contributions of identified shocks; the line plots demeaned unemployment.

the implied contributions. In the Great Recession, noise shocks account for a sizable share of the time required to reverse half of the rise in unemployment; on average, noise shocks account for a nontrivial share across recessions.

A.6.4 Responses of forecast errors

Figures A16 and A17 compare the responses of realized unemployment and SPF forecasts (and forecast errors) to identified shocks.

A.6.5 Subsample analysis

To assess whether shock properties change over time, I compare summary statistics for identified shocks and nowcast errors before and after 1990. Table A3 reports the results. Post-1990, identified noise shocks become more volatile, while persistent TFP shocks are less volatile. Nowcast errors also become more volatile, consistent with changes in the information environment and/or the composition of shocks.

Table A2: Contribution of Noise Shocks to the 75% Recovery Horizon

	No of quarters for 75% recovery	
Recession	Data	Share explained by Noise shocks
2007-09	22	35%
2001	18	33%
1990-91	14	28%
1981-82	9	33%
1973-75	20	29%
Average	16.6	32%

Note: The table reports the number of quarters required to reverse 75% of the recessionary increase in unemployment. Contributions are computed from historical-decomposition-implied counterfactual paths.

Table A3: Summary Statistics Pre- and Post-1990

	1968-1989		1990-2019	
	Mean	SD	Mean	SD
Unemployment Rate	5.68	1.65	5.80	1.83
GDP Nowcast Error	0.06	1.72	0.25	2.59
Unemployment Nowcast Error	0.07	0.695	-0.03	1.20
Noise Shock	0.04	0.745	-0.05	1.27
Persistent Shock	-0.21	1.34	0.03	0.68
Transitory Shock	0.18	0.83	0.27	0.89

Note: Summary statistics for 1968–1989 and 1990–2019.

Figure A14: Historical Contribution to Job-Finding and Vacancies

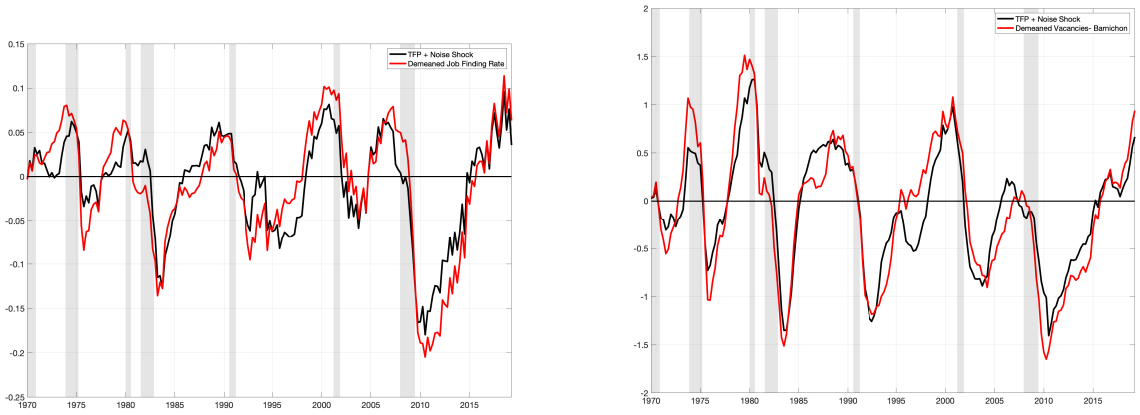

(a) Job-finding rate (b) Vacancy posting

Note: Historical decompositions based on (4).

Figure A15: Historical Contribution to Wages and Investment

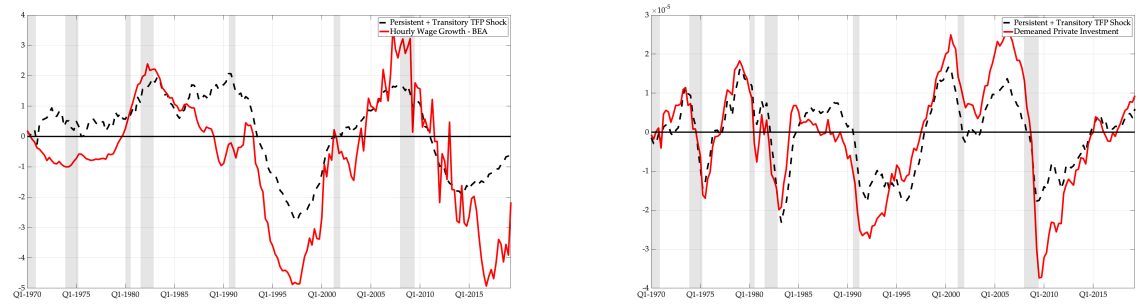

(a) Wage growth (b) Private investment

Note: Historical decompositions based on (4).

A.7 Robustness: Controlling for Uncertainty Shocks

A first concern is that the baseline three-variable SVAR may load monetary, financial, or other omitted shocks onto the persistent, transitory, or noise shocks. To address this, I estimate a four-variable VAR that adds unemployment and leaves the row and column associated with a fourth shock unrestricted:

$$(5) \qquad \begin{bmatrix} z_t \\ y_t \\ nce_t \\ u_t \end{bmatrix} = \Sigma_j^p B_j \begin{bmatrix} z_{t-p} \\ y_{t-p} \\ nce_{t-p} \\ u_{t-p} \end{bmatrix} + \begin{bmatrix} + & + & 0 & * \\ + & + & + & * \\ * & * & - & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} \epsilon_t \\ \eta_t \\ \nu_t \\ \mu_t \end{bmatrix}.$$

Figure A16: Unemployment: Realized and Expected

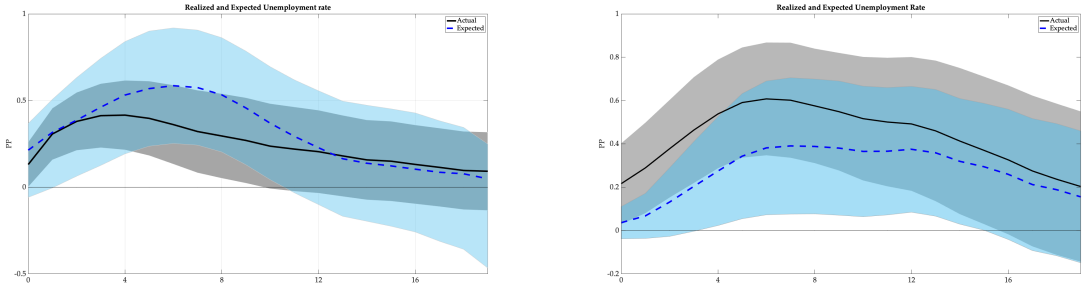

(a) Noise shocks (b) Persistent TFP shocks

Note: Responses of realized unemployment and SPF expected unemployment to one-standard-deviation shocks (baseline sample 1968Q4–2019Q4).

Figure A17: Response of Unemployment Forecast Errors

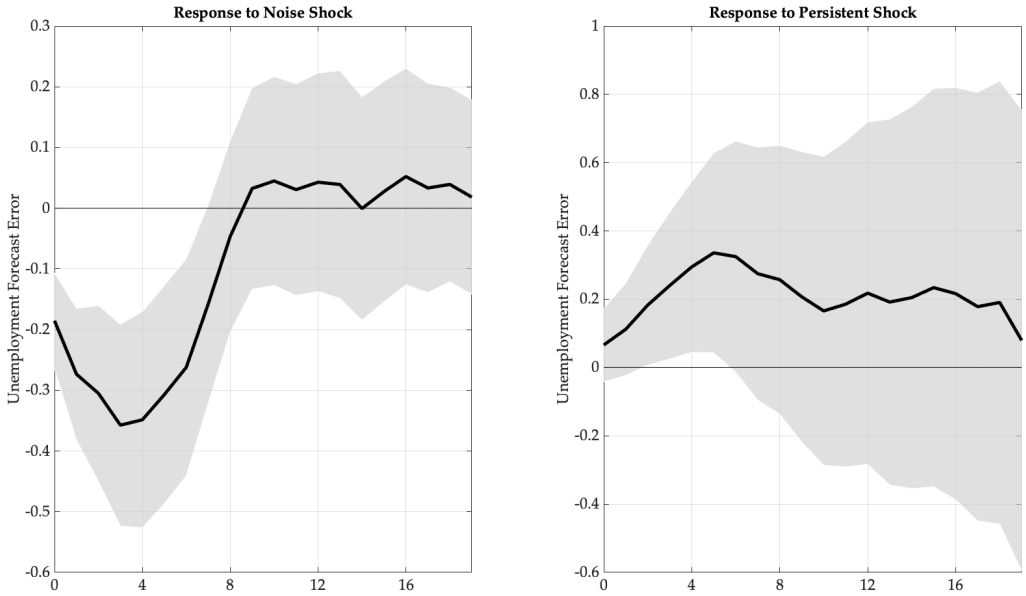

Note: Impulse responses of SPF unemployment forecast errors to identified shocks (baseline sample 1968Q4–2019Q4).

As in the baseline, I assume that ϵ_t maximizes the forecast error variance of TFP at a long-run horizon. Figure A18 shows the historical decomposition of unemployment. Noise shocks still explain substantial variation in the 1990–91, 2001, and 2007–09 recessions. In contrast, the unrestricted shock contributes more in the 1973–75, 1980, and 1981–82 recessions, consistent with the importance of non-TFP shocks such as monetary or oil shocks in those episodes.

Figure A18: 4 Variable VAR: Historical Contribution of Shocks to Unemployment Rate

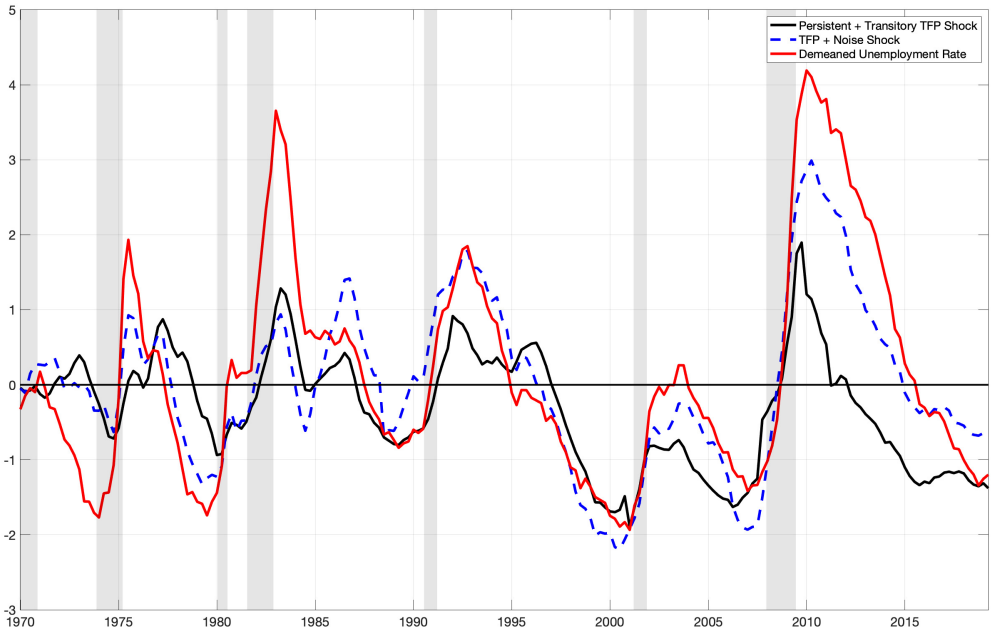

Note: This figure shows the historical decomposition of the unemployment rate following equation 4. The black line is the cumulative contribution of the identified persistent and transitory TFP shocks to demeaned unemployment (red line). The dashed blue line is the contribution of the TFP shocks and noise shocks. The remaining movement is explained by the fourth shock μ_t.

A potential concern is that the identified noise shock may proxy for uncertainty shocks. To address this, I re-estimate the local projections controlling for the uncertainty shock series from Bloom (2009):

$$y_{t+h} = \alpha_h + \tilde{\beta}_h u_t^{\text{noise}} + \theta_h u_t^{\text{uncert}} + \sum_{p=1}^{P} \tilde{\delta}_{p,h} \tilde{w}_{t-p} + \tilde{\mu}_{t,h}, \tag{6}$$

where $\tilde{w}_{t-p}$ includes lags of y_t, u_t^{noise}, and u_t^{uncert}. Figure A19 compares baseline and uncertainty-controlled responses. The results are similar: controlling for uncertainty shocks does not materially change the estimated responses to noise shocks.

Figure A19: Impulse Responses to Noise Shocks Controlling for Uncertainty

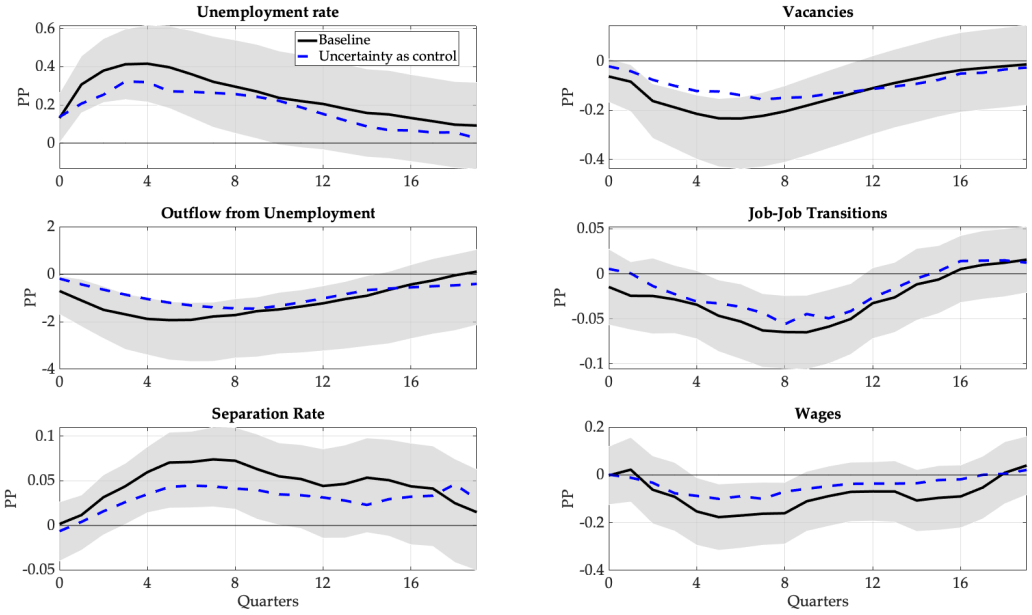

Note: Baseline and uncertainty-controlled impulse responses. The bands report 90% confidence intervals.

B Theoretical Appendix

This appendix collects the theoretical building blocks underlying the quantitative analysis. It has three aims. First, it formalizes the information structure and the associated Bayesian learning problem in a linear–Gaussian state-space representation, including the Kalman filter used in solving and estimating the model (Section B.1). Second, it reports auxiliary estimation and validation exercises, such as the re-estimation of the full-information benchmark and the internal validity check of the SVAR identification. Third, it provides a stylized analytical illustration of the mechanism under explicit restrictions, which isolates the role of imperfect information for unemployment dynamics. Throughout, timing and information sets are kept consistent with the main text: period-t decisions are chosen at the *start* of period t on the basis of prior beliefs and predetermined states, while the realizations of aggregate productivity and the public signal at the *end* of period t are used only to update beliefs for period t+1.

B.1 Information structure and Bayesian updating

This subsection restates the aggregate information structure in a form that is convenient for both the structural model and the empirical SVAR. I decompose aggregate productivity

into persistent and transitory components, introduce the noisy public signal observed by agents, and embed these objects in a linear state-space system. Under the maintained Gaussian assumptions, beliefs about the persistent component evolve according to a Kalman filter with a well-defined steady-state gain. The resulting filtering problem is the one used in the solution, simulation, and estimation of the model.

Signals and state decomposition. Aggregate productivity is decomposed into a persistent component x_t and a transitory component η_t:

$$z_t = x_t + \eta_t, \qquad \eta_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\eta^2), \tag{7}$$

$$x_t = \rho x_{t-1} + \epsilon_t, \qquad \epsilon_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\epsilon^2), \quad 0 < \rho < 1. \tag{8}$$

Agents also observe a public signal $\hat{s}_t$ about x_t that is contaminated by a persistent noise component a_t:

$$\hat{s}_t = x_t + a_t, \tag{9}$$

$$a_t = \rho_a a_{t-1} + \nu_t, \qquad \nu_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\nu^2), \quad |\rho_a| < 1. \tag{10}$$

The shocks $(\eta_t, \epsilon_t, \nu_t)$ are mutually independent. The persistent component x_t is the only source of long-run fluctuations in productivity; the temporary shock η_t and the noise component a_t affect beliefs but do not alter the long-run level of z_t.

Timing and information sets. Let $\mathcal{I}_t^{\text{beg}}$ denote information available at the start of period t and $\mathcal{I}_t^{\text{end}} = \mathcal{I}_t^{\text{beg}} \cup \{z_t, \hat{s}_t\}$ the end-of-period information set. In the notation of the main text, $\mathcal{I}_t^{\text{beg}}$ corresponds to $\mathcal{I}^{t-1}$. Define the prior and posterior for the persistent component:

$$x_{t|t-1} \equiv \mathbb{E}[x_t \mid \mathcal{I}_t^{\text{beg}}], \qquad x_{t|t} \equiv \mathbb{E}[x_t \mid \mathcal{I}_t^{\text{end}}].$$

Period-t controls are functions of $(x_{t|t-1}, \text{other predetermined states})$. The posterior $x_{t|t}$ pins down next period's prior via

$$x_{t+1|t} = \rho\, x_{t|t}. \tag{11}$$

State-space form. Let the latent state be $\xi_t \equiv (x_t, a_t)'$ and the observation vector $y_t \equiv (z_t, \hat{s}_t)'$. Equations (8)–(10) and (7)–(9) can be written as

$$\xi_t = F\xi_{t-1} + \varepsilon_t^{\xi}, \qquad F \equiv \begin{pmatrix} \rho & 0 \\ 0 & \rho_a \end{pmatrix}, \qquad \varepsilon_t^{\xi} \equiv \begin{pmatrix} \epsilon_t \\ \nu_t \end{pmatrix}, \qquad \mathbb{E}[\varepsilon_t^{\xi}(\varepsilon_t^{\xi})'] = Q \equiv \begin{pmatrix} \sigma_\epsilon^2 & 0 \\ 0 & \sigma_\nu^2 \end{pmatrix}, \tag{12}$$

$$y_t = H\xi_t + v_t, \qquad H \equiv \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \qquad v_t \equiv \begin{pmatrix} \eta_t \\ 0 \end{pmatrix}, \qquad \mathbb{E}[v_t v_t'] = R \equiv \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & 0 \end{pmatrix}. \tag{13}$$

The second measurement equation is observed without measurement error in the model (the "noise" in $\hat{s}_t$ is exactly a_t, which is part of the latent state). In the numerical implementation, I introduce a numerically negligible measurement error in the second equation (equivalently, replace $R_{22} = 0$ by $R_{22} = \epsilon_R$ with ϵ_R close to machine precision) to ensure stability of matrix inversions.

Kalman filter recursions. Let $\xi_{t|t-1} \equiv \mathbb{E}[\xi_t \mid \mathcal{I}_t^{\text{beg}}]$ and $\xi_{t|t} \equiv \mathbb{E}[\xi_t \mid \mathcal{I}_t^{\text{end}}]$. Let $P_{t|t-1}$ and $P_{t|t}$ be the forecast and posterior error covariance matrices. The Kalman filter is

$$\xi_{t|t-1} = F\xi_{t-1|t-1}, \qquad P_{t|t-1} = FP_{t-1|t-1}F' + Q, \tag{14}$$

$$\iota_t \equiv y_t - H\xi_{t|t-1}, \qquad S_t \equiv HP_{t|t-1}H' + R, \tag{15}$$

$$G_t \equiv P_{t|t-1}H'S_t^{-1}, \tag{16}$$

$$\xi_{t|t} = \xi_{t|t-1} + G_t\iota_t, \qquad P_{t|t} = (I - G_tH)P_{t|t-1}. \tag{17}$$

The posterior for x_t is the first element of $\xi_{t|t}$. Writing (17) componentwise makes the role of the two signals explicit:

$$x_{t|t} = x_{t|t-1} + g_{x,z}\big(z_t - x_{t|t-1}\big) + g_{x,s}\big(\hat{s}_t - x_{t|t-1} - a_{t|t-1}\big), \tag{18}$$

$$a_{t|t} = a_{t|t-1} + g_{a,z}\big(z_t - x_{t|t-1}\big) + g_{a,s}\big(\hat{s}_t - x_{t|t-1} - a_{t|t-1}\big), \tag{19}$$

where $(g_{x,z}, g_{x,s})$ and $(g_{a,z}, g_{a,s})$ are the corresponding rows of G_t.

Steady state. Under standard conditions, $P_{t|t-1}$ converges to a fixed point $\bar{P}$ that solves the discrete-time algebraic Riccati equation

$$\bar{P} = F\bar{P}F' + Q - F\bar{P}H'(H\bar{P}H' + R)^{-1}H\bar{P}F'. \tag{20}$$

The steady-state Kalman gain is

$$\bar{G} = \bar{P}H'(H\bar{P}H' + R)^{-1}. \tag{21}$$

In the quantitative implementation, the solution of (20) and the associated gain $\bar{G}$ are computed numerically. The linearized equilibrium conditions in the main text are solved using this steady-state filter, so that decision rules are functions of the perceived state $(x_{t|t-1}, a_{t|t-1}, \dots)$.

Variance-share (signal-to-noise) parameterization. To summarize the informativeness of the public signal $\hat{s}_t$, it is useful to report the variance share

$$\mathcal{K} \equiv \frac{\mathrm{Var}(x_t)}{\mathrm{Var}(\hat{s}_t)} = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_a^2}, \qquad \sigma_x^2 = \frac{\sigma_\epsilon^2}{1-\rho^2}, \quad \sigma_a^2 = \frac{\sigma_\nu^2}{1-\rho_a^2}. \tag{22}$$

This is the "signal-to-noise ratio" reported in the calibration tables. It measures the fraction of variation in the public signal that is attributable to true innovations in the persistent component x_t. It is distinct from the Kalman gain in (21): $\mathcal{K}$ is a summary statistic implied by $(\rho, \rho_a, \sigma_\epsilon, \sigma_\nu)$ and the Riccati equation (20). In the estimation, I estimate $(\sigma_\epsilon, \sigma_\nu, \rho_a)$ and report the implied $\mathcal{K}$ (with standard errors computed by the delta method).

B.2 Estimation: full-information benchmark

This subsection reports the re-estimation of the full-information model used as a benchmark in the quantitative analysis. The full-information economy corresponds to the case in which agents observe x_t directly and there is no signal-extraction problem. In the state-space representation above, this can be represented by setting $a_t \equiv 0$ so that $\hat{s}_t = x_t$ and no filtering step is needed. I re-estimate the parameters governing the shock processes and matching block by matching the model-implied impulse responses to the empirical responses to identified persistent TFP shocks, using the same GMM criterion as in the main text (cf. equation **??**).

Table B1 reports the estimated parameters and standard errors (computed by the delta method based on the outer-product-of-gradients approximation to the GMM covariance matrix).

Table B1: Estimated Parameters from IRF Matching: Full-Information Model

Parameters	Interpretation	Value	Target
Ψ	Match efficiency	0.48	Unemployment Rate = 0.055
κ	Cost of hiring	8.23	$U - E = 0.28$
μ	Scale parameter of search cost	0.082	$E - E = 0.025$
$1 - \sigma$	Separation rate	0.010	$E - U = 0.010$
ϕ	SS productivity from bad job	0.68	Δ Wage of E-E = 0.045
Parameters	**Interpretation**	**Estimate**	**Std. Error**
λ	Renegotiation frequency	0.88	0.13
ξ	Probability of finding a good job	0.18	0.05
η_h	Hiring cost convexity	0.34	0.09

Note: Parameter estimates from the impulse-response matching exercise in equation **??**. Standard errors are computed by the delta method.

B.3 Internal validity of the SVAR identification

To assess whether the SVAR identification recovers the intended structural shocks in the estimated model, I conduct an internal validity exercise. Starting from the imperfect-information benchmark calibration, I simulate the model for $T = 10{,}000$ periods, construct artificial data on unemployment, vacancies, and the other observables used in the empirical VAR, and re-estimate the SVAR using the same lag length, sample frequency, and identifying restrictions as in the data.

For each draw, I obtain (i) SVAR-implied impulse responses to the identified "noise" and TFP shocks, and (ii) model-implied impulse responses generated from the linearized decision rules. Figure B1 compares these objects for the benchmark calibration. The SVAR IRFs estimated on model-simulated data track the model-implied IRFs very closely across all horizons and variables, indicating that—within the model environment—the SVAR identification successfully recovers the targeted structural shocks.

B.4 Analytical properties under explicit restrictions

This subsection provides a stylized analytical illustration of the learning mechanism in a stripped-down environment. The system below abstracts from on-the-job search, capital accumulation, and wage staggering; its purpose is to make the role of imperfect information transparent rather than to replicate all margins of the full model.

Figure B1: Internal Validity of the SVAR in Model-Simulated Data

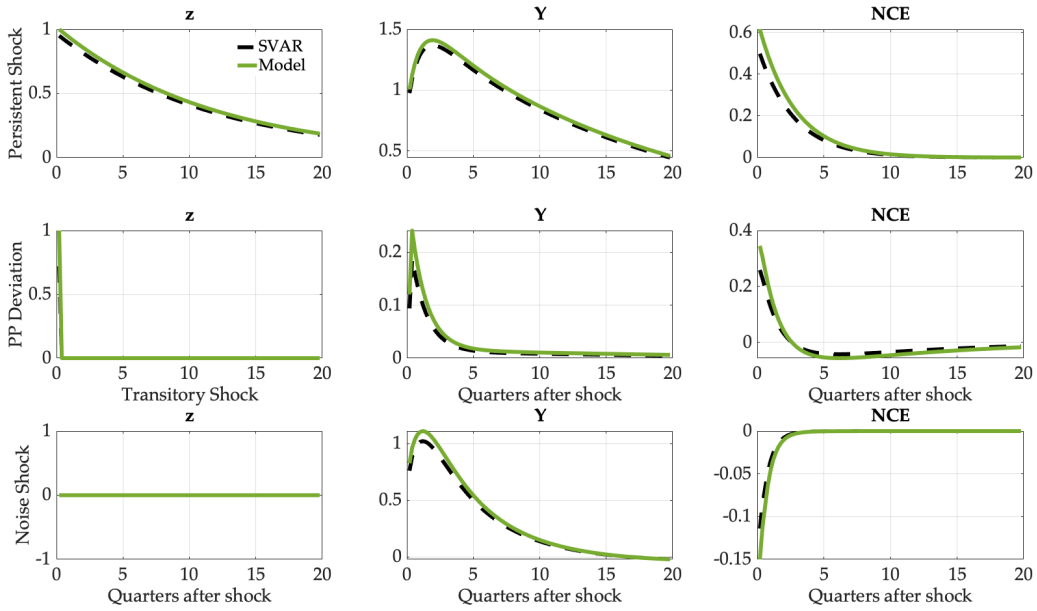

Note: Dashed lines: SVAR IRFs estimated on model-simulated data ($p = 8$, $T = 10{,}000$). Solid lines: model-implied IRFs from the linearized solution.

Setup. Consider

$$x_t = \rho x_{t-1} + \epsilon_t, \qquad 0 < \rho < 1, \tag{23}$$

$$s_t = x_t + \nu_t, \tag{24}$$

$$\hat{x}_t = m_t + K_F\,(s_t - m_t), \tag{25}$$

$$m_{t+1} = \rho \hat{x}_t, \tag{26}$$

$$u_t = \varphi_u u_{t-1} - \psi m_t, \qquad 0 < \varphi_u < 1,\ \psi > 0, \tag{27}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$ are i.i.d. and independent, and $K_F \in (0,1)$ denotes the scalar steady-state Kalman gain in this simple signal-extraction problem. The prior m_t is defined as

$$m_t = \mathbb{E}[x_t \mid s_0, \ldots, s_{t-1}],$$

so period-t decisions use the prior m_t, consistent with the pre-commitment timing in the main model. For comparison, I also consider a contemporaneous full-information (FI) benchmark in which agents observe x_t when choosing u_t. This auxiliary benchmark abstracts from the one-period decision lag and isolates the role of signal extraction itself.

Assume $x_{-1} = m_{-1} = u_{-1} = 0$.

Full-information benchmark. Under this contemporaneous full-information benchmark, decisions respond immediately to innovations in x_t. For a one-time persistent shock $\epsilon_0 < 0$ and $\epsilon_t = 0$ for $t \geq 1$,

$$u_0^{FI} = -\psi x_0 = -\psi \epsilon_0 > 0,$$

and the subsequent path is given by

$$u_t^{FI} = \varphi_u u_{t-1}^{FI} - \psi x_t = \varphi_u u_{t-1}^{FI} - \psi \rho^t \epsilon_0, \qquad t \geq 1.$$

Three illustrative properties. The following statements highlight how imperfect information modifies the dynamics of u_t.

1. **Impact delay under imperfect information.** Consider a one-time innovation $\epsilon_0 < 0$ with $\epsilon_t = 0$ for $t \geq 1$ and $\nu_t \equiv 0$. Under imperfect information, $m_0 = 0$ because the shock has not yet been incorporated into the prior, so

   $$u_0^{II} = \varphi_u u_{-1} - \psi m_0 = 0.$$

   After observing $s_0 = x_0 = \epsilon_0$, the posterior is $\hat{x}_0 = K_F \epsilon_0$ and hence $m_1 = \rho \hat{x}_0 = \rho K_F \epsilon_0$. Thus

   $$u_1^{II} = \varphi_u u_0^{II} - \psi m_1 = -\psi \rho K_F \epsilon_0 > 0.$$

   By contrast, under the contemporaneous FI benchmark $u_0^{FI} > 0$. In the prior-based imperfect-information economy, the unemployment response cannot peak at $t = 0$, whereas it can in the FI comparison benchmark.

2. **Belief persistence and the noise-to-signal ratio.** Define the belief error $e_t \equiv x_t - m_t$. Along the one-time fundamental-innovation path with $\nu_t \equiv 0$, we have $e_0 = \epsilon_0$ and, for $t \geq 1$,

   $$e_t = \rho(1 - K_F) e_{t-1}.$$

   Hence the response of e_t to a one-time innovation $\epsilon_0 \neq 0$ (and $\epsilon_t = 0$ for $t \geq 1$, $\nu_t \equiv 0$) is

   $$e_t = \Big(\rho(1 - K_F)\Big)^t \epsilon_0, \qquad t \geq 0.$$

   With nonzero signal noise away from this path, the belief-error recursion picks up an additional innovation term proportional to ν_{t-1}, but the decay rate of the deterministic component remains $\rho(1 - K_F)$. In the scalar Kalman filter, K_F is strictly decreasing in the noise-to-signal ratio $\sigma_\nu^2/\sigma_\epsilon^2$. Hence the autoregressive root $\rho(1 - K_F)$ and the half-life of the belief error e_t are increasing in $\sigma_\nu^2/\sigma_\epsilon^2$.

A convenient way to isolate the effect of imperfect information on unemployment is to consider the incremental response

$$\Delta u_t \equiv u_t^{II} - u_t^{FI}.$$

Since the decision-relevant state under imperfect information is $m_t = x_t - e_t$ while under the contemporaneous FI benchmark it is x_t, we have $\Delta m_t = -e_t$ and

$$\Delta u_t = \varphi_u \Delta u_{t-1} - \psi \Delta m_t = \varphi_u \Delta u_{t-1} + \psi e_t.$$

Because the deterministic component of e_t decays at rate $\rho(1 - K_F)$ while φ_u is fixed, the persistence of Δu_t (and the associated half-life) is increasing in the noise-to-signal ratio.

3. **Hump-shaped responses to pure noise shocks.** Consider a pure noise shock sequence with $\nu_0 \neq 0$, $\nu_t = 0$ for $t \geq 1$, and $\epsilon_t \equiv 0$. Then $x_t \equiv 0$ for all t, but beliefs move. Starting from $m_0 = 0$ and using the steady-state Kalman filter,

$$\hat{x}_0 = K_F \nu_0, \qquad m_1 = \rho \hat{x}_0 = \rho K_F \nu_0.$$

For $t \geq 1$, the prior evolves according to $m_{t+1} = \rho(1 - K_F) m_t$, so

$$m_t = \lambda_m^{t-1} m_1, \qquad t \geq 1, \qquad \lambda_m \equiv \rho(1 - K_F) \in (0, \rho).$$

With $u_0 = 0$, the unemployment response is

$$u_t = -\psi \sum_{k=1}^{t} \varphi_u^{t-k} m_k = -\psi m_1 \sum_{k=1}^{t} \left(\varphi_u^{t-k} \lambda_m^{k-1} \right) = -\psi m_1 \frac{\varphi_u^t - \lambda_m^t}{\varphi_u - \lambda_m}, \qquad t \geq 1,\ \varphi_u \neq \lambda_m.$$

Under the sufficient conditions $0 < \lambda_m < \varphi_u < 1$ and $\varphi_u + \lambda_m > 1$, the sequence $\{|u_t|\}_{t \geq 1}$ is single-peaked (for example, $|u_2| > |u_1|$ if $\varphi_u + \lambda_m > 1$), while $u_t \to 0$ as $t \to \infty$ since both φ_u^t and λ_m^t tend to zero. Thus, even when fundamentals do not move at all, persistent noise in the signal can generate hump-shaped unemployment dynamics via the interaction of sluggish belief updating and the intrinsic persistence of u_t.

The first result follows from comparing prior-based decisions under imperfect information with a contemporaneous full-information benchmark. The latter two results show how the degree of informational noise and the relative location of the roots $(\varphi_u, \rho(1 - K_F))$ shape the persistence and hump-shapedness of unemployment in a transparent benchmark. In

the full model, additional margins—capital, wage rigidity, and on-the-job search—interact with this core learning channel but do not overturn its basic logic.

B.5 Quantitative Results Referenced in the Main Text

This subsection collects additional quantitative results from the estimated benchmark model that are referenced, but not shown in full, in the main text. The focus is on the imperfect-information model with staggered wage contracts and on-the-job search (OJS), estimated as described in Section **??**. I report impulse responses of aggregate quantities and labor market variables to the identified shocks, which complement the empirical and model-based responses shown in the main text.

B.5.1 Impulse Responses from the Model

This subsection documents model-implied impulse responses for a broader set of variables than those plotted in the main text. In all cases, the responses are computed from the log-linearized imperfect-information, sticky-wage, OJS model using the solution method described in Section **??** and the estimated information parameters from Table **??**. The figures below show how output, investment, and key labor market variables respond over time to (i) a noise shock to beliefs about the persistent component of productivity and (ii) a persistent productivity shock.

A negative noise shock induces a downturn driven purely by misperceptions: beliefs about future productivity deteriorate when the signal is observed even though fundamentals move little on impact. Because period-t controls are already fixed, the belief revision affects vacancy posting, job-finding rates, and on-the-job search from the next decision period onward, which in turn lowers employment, output, and investment. The dynamics are hump-shaped for unemployment and protracted for vacancies, in line with the empirical impulse responses.

A positive persistent productivity shock, by contrast, raises realized productivity on impact and expected future productivity after agents update their beliefs. The improvement in expected profitability then strengthens vacancy posting and hiring, stimulates investment, and lowers unemployment. Relative to a full-information benchmark, the imperfect-information model delivers more gradual adjustment, as agents learn about the shock over time rather than instantaneously.

Figure B2: Model-Implied Impulse Response Functions to a Negative Noise Shock

Note: Model-implied impulse responses from the imperfect-information, sticky-wage, on-the-job-search specification to a noise shock in the public signal about the persistent component of productivity. The panels report the dynamic responses of output, investment, and selected labor market variables.

Figure B3: Model-Implied Impulse Response Functions to a Positive Persistent Productivity Shock

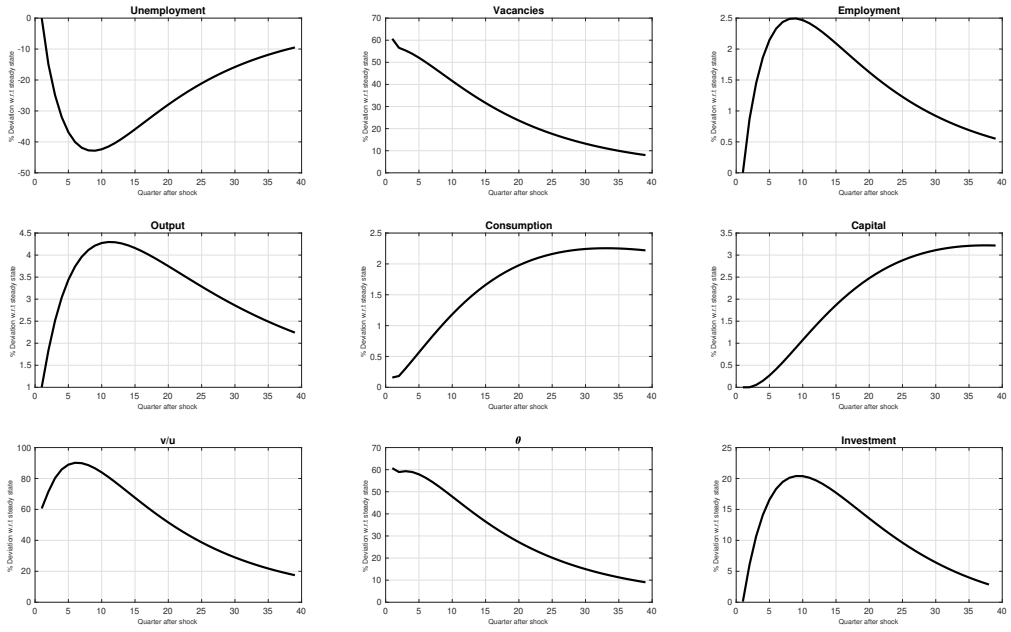

Note: Model-implied impulse responses from the imperfect-information, sticky-wage, on-the-job-search specification to a persistent productivity shock. The panels report the dynamic responses of output, investment, and selected labor market variables.

Figure B4: 4–8-Quarter-Ahead Projections in the Model

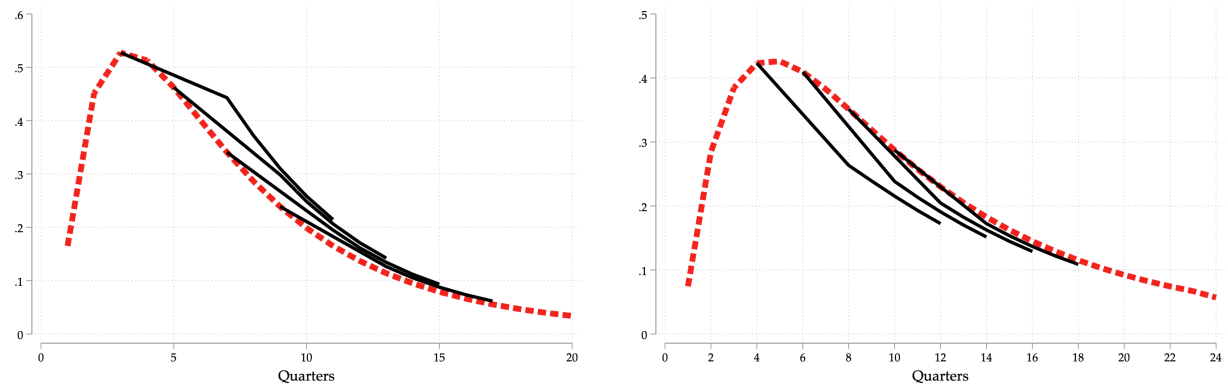

(a) Noise Shock

(b) Persistent Productivity Shock

Note: This figure shows the 4–8-quarter-ahead projections by agents in the model in response to a noise shock (a) and a persistent TFP shock (b). The solid thick black line is the actual response of unemployment due to these shocks.

B.5.2 Projections from the Model

In this section, I present the 4–8-quarter-ahead projections by agents in the model in response to a persistent TFP shock and a noise shock. When faced with a persistent TFP shock, due to imperfect information, agents attribute part of the shock to noise as well as transitory productivity, and hence their projections underreact to the actual unemployment rate. However, the reverse happens when they face a noise shock. They similarly attribute some part of the shock to persistent or transitory productivity and hence initially expect unemployment to be higher than it actually is (since true productivity has not changed). As they learn, they place more and more weight on the shock being noise and their projections move closer to the actual.

Figure B5a shows the model-generated one-, two- and three-year ahead unemployment projections in the model after the Great Recession. Here, all three shocks, identified from the VAR, act together each period while simulating the imperfect-information model with noise shocks. Since all three shocks act, the projections underreact if the contribution of the persistent shock dominates the contribution of the noise shocks as well as transitory shocks. Similarly, as the contribution of the noise shocks dominates, the projections overestimate the unemployment rate. As seen in the historical decomposition of the unemployment rate in the data in Figure A13, the contribution of the noise shocks to the movement in unemployment dominates after 2012. Thus, in the model, initially, as the productivity shocks have higher weight, the unemployment rate is underestimated by the agents in the model. However, from 2012, the contribution of the noise shocks

Figure B5: Unemployment Rate: Projections and Actual– Model vs Data

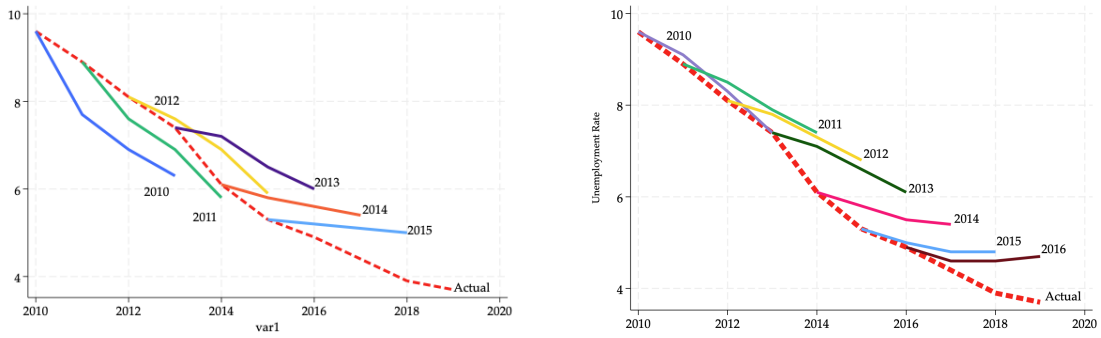

(a) Unemployment Projections: Model

(b) Unemployment Projections: Data

Note: Panel (a) shows the model-implied forecasts for the unemployment rate 1-, 2- and 3-year ahead. The dashed black line is the model-simulated unemployment rate for the Great Recession. While simulating the model, each period all three shocks act. In Panel (b) the various colored lines represent the median long-run (1-, 2- and 3-year ahead) projections of the unemployment rate from the Survey of Professional Forecasters during the Great Recession. The dashed red line is the actual unemployment rate.

increases but the agents are unable to discern the shock from a true persistent productivity shock and hence keep expecting higher unemployment rates in the future. However, as the shock is truly noise, the actual unemployment rate is lower than expected. This is similar to the pattern seen in the data in Figure B5b. It is important to note that the noise shocks are unique in generating overestimation of long-run unemployment projections. For all structural shocks, the long-run expectations underestimate the unemployment rate. Thus, noise shocks can be a potential solution to the consistent pattern observed in the data where the long-run unemployment forecasts are overestimated by professional forecasters.

B.5.3 Unemployment Dynamics across Recessions: Data vs Model

The calibrated model is simulated to generate counterfactual unemployment rate series for five recessions between 1968 and 2019. This exercise shows that imperfect information explains the slow recovery of the unemployment rate in the last three recessions. For this exercise, the model is normalized to match the starting unemployment rate for each of the recessions. While simulating the imperfect-information model, each period all three identified shocks from the VAR are incorporated. For the full-information model, I only introduce the persistent and the transitory shocks each period. Furthermore, the full-information model is re-estimated as described in the previous section to match the empirical IRFs to the persistent TFP shocks. The estimated parameters for the full-information model are presented in the Appendix.

Figure B6: Model-Implied Recovery of Unemployment for Recessions

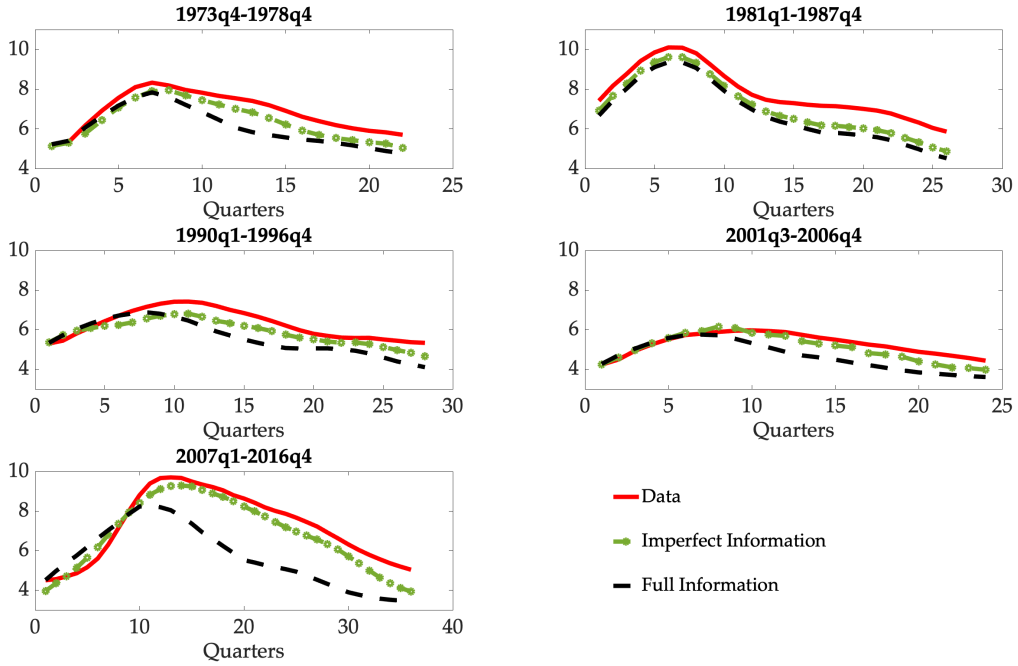

Note: This figure plots the model-implied simulated unemployment rate for the recalibrated full-information model (dashed blue line) and the imperfect-information model (solid green line) for major recessions between 1968–2019.

B.5.4 Comparing Mechanisms in the Model

This subsection compares the persistence and volatility of unemployment implied by the model under different combinations of wage-setting and on-the-job search (OJS), with and without imperfect information. I consider four specifications: (i) flexible wages without OJS, (ii) flexible wages with OJS, (iii) sticky wages without OJS, and (iv) sticky wages with OJS. For each specification, I construct a full-information benchmark and an imperfect-information counterpart (with the same structural parameters but shutting down noise shocks where indicated), feed in the identified structural shocks from the SVAR, and evaluate the model's implications for recovery dynamics and business cycle moments. The full-information model is re-estimated following Appendix Section B.2 to match the empirical impulse responses to persistent TFP shocks.

Persistence This exercise isolates the role of imperfect information about the *persistent vs. transitory* decomposition of TFP by shutting down noise shocks entirely. Both the full-information and imperfect-information specifications are driven only by the identified persistent and transitory TFP shocks from the SVAR; the noise innovation ν_t is set to

zero. The only difference between the two is whether agents observe the persistent and transitory components separately (full information) or must infer them from the aggregate signal under Bayesian learning (imperfect information). Any added persistence under imperfect information therefore reflects slow learning about shock persistence alone, not belief distortions from noise.

To quantify unemployment persistence, I compare the average duration required for unemployment to return to its pre-recession trough across the recessions between 1968 and 2019. For each recession, let u_t denote the unemployment rate, u^{pre} the pre-recession trough, and u^{peak} the cyclical peak during the recession. The full-recovery horizon in the data, $h^{100,\text{data}}$, is defined as the smallest $h \geq 0$ such that

$$\frac{u^{\text{peak}} - u_{t_{\text{trough}}+h}}{u^{\text{peak}} - u^{\text{pre}}} \geq 1.$$

The average $h^{100,\text{data}}$ across recessions is approximately 17 quarters (excluding the short-lived 1980 recession). For each model specification, I simulate the economy using the estimated persistent and transitory TFP shock series (with noise shocks shut down), compute the model-implied full-recovery horizon recession by recession, and then average across recessions in the same way.

Figure B7 summarizes these results. The x-axis reports the average full-recovery horizon in quarters; the percentages on top of each bar are the share of the empirical benchmark (17 quarters) explained by the corresponding model specification. Within each of the four wage/OJS configurations, the blue bar reports the contribution of the re-estimated full-information model, and the green segment reports the incremental contribution of imperfect information—*with noise shocks still shut down*—when agents must learn the persistent/transitory split from the aggregate signal. The total height of the bar therefore measures the contribution of the imperfect-information model without noise in that configuration.

Two observations emerge. First, across all four wage/OJS configurations, the imperfect-information model (blue plus green) moves the implied full-recovery horizon closer to the empirical benchmark than the full-information model alone. Second, the incremental contribution of learning—captured by the green segment—is quantitatively non-negligible in each configuration, indicating that imperfect information is a robust source of additional persistence over and above wage rigidity and on-the-job search. This echoes the logic in Wright (1986), who shows that imperfect information can generate persistent dynamics in search models through endogenous learning, albeit in a different informational environment.

Figure B7: Average Duration to Return Unemployment to its Pre-Recession Trough Across Models: Without Noise Shocks

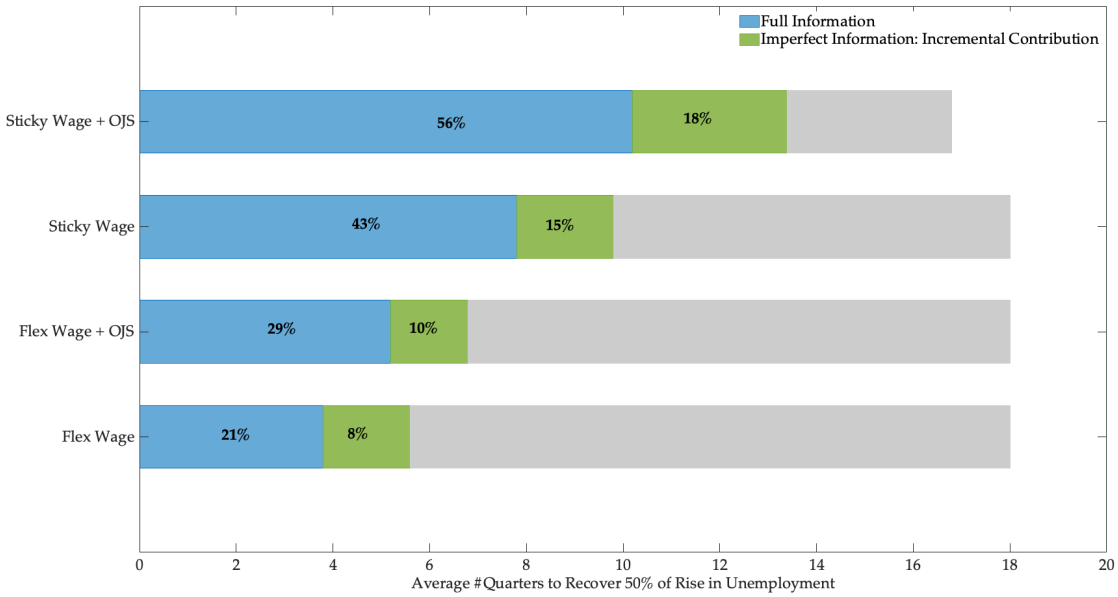

Note: The figure plots the model-implied duration, in quarters, from the beginning of each recession to the point where unemployment returns to its pre-recession trough. The measure is averaged across recessions between 1968 and 2019. The percentages above the bars express this duration as a share of the empirical benchmark (17 quarters). For each wage/OJS configuration, the blue bar corresponds to the re-estimated full-information model; the green segment shows the incremental contribution of imperfect information in the same specification when noise shocks are shut down. The total height of each bar is thus the contribution of the imperfect-information model (without noise) in that configuration.

Table B2: Duration of Recovery of Unemployment Rate Across Recessions

Recession	Data	Full Information	Imperfect Information
1973-75	NA	14	17
1980-82	22	17	21
1990-91	21	16	24
2001	NA	14	21
2007-09	33	22	32

Note: The table reports the number of quarters required for the unemployment rate to return to its pre-recession trough across five recessions between 1973 and 2019. For each episode, the model is initialized to match the observed unemployment rate at the onset of the recession. The imperfect-information model is simulated with all three identified shocks (persistent, transitory, and noise). The full-information model is driven only by the persistent and transitory TFP shocks and is re-estimated to match the empirical impulse responses to the persistent TFP shock, as described in Appendix Section B.2. Note that here I treat the 1980 and 1981-82 recessions as a single episode, given the short duration of the 1980 recession and the fact that the unemployment rate did not return to its pre-1980 trough before the onset of the 1981-82 recession.

Using the same full-recovery metric, Table B2 reports the recession-by-recession recovery duration for the re-estimated full-information and imperfect-information benchmark models: the number of quarters it takes for unemployment to return to its pre-recession trough level in each of the five major recessions between 1973 and 2019. For each episode, the model is initialized to match the observed unemployment rate at the start of the recession. The imperfect-information model is simulated using all three identified shocks from the SVAR (persistent, transitory, and noise), while the full-information model is driven only by the persistent and transitory TFP shocks and re-estimated as in Appendix Section B.2. The imperfect-information model tracks the observed recovery durations more closely, particularly in recessions with notably slow labor market recoveries.

Business Cycle Statistics Across Specifications Table B3 compares business cycle statistics obtained by simulating the imperfect-information model and the re-estimated full-information model under the four wage/OJS configurations to their empirical counterparts for the US economy over 1968–2019. The table reports the standard deviation of output (Y), the unemployment rate (U), vacancies (V), job-to-job transitions ($E-E$), unemployment-to-employment transitions ($U-E$), and the hiring rate, computed from model-simulated series that are HP-filtered with smoothing parameter $\lambda_{\text{HP}} = 1{,}600$ and from US data treated analogously.

Across all four specifications, introducing imperfect information brings the model-implied volatilities of unemployment, vacancies, and transition rates closer to their empirical counterparts, relative to the full-information benchmarks. The gains are particularly pronounced in environments with wage rigidity and OJS, where the interaction between

Table B3: Business Cycle Statistics Across Model Specifications

	Data (SD)	Flex Wage , No OJS		Flex Wage , OJS		Sticky Wage , No OJS		Sticky Wage , OJS	
		Full Info	Imperfect Info	Full Info	Imperfect Info	Full Info	Imperfect Info	Full Info	Imperfect Info
Y	0.019	0.009	0.014	0.011	0.017	0.013	0.021	0.018	0.027
U	0.162	0.029	0.068	0.052	0.098	0.087	0.128	0.121	0.153
V	0.182	0.032	0.091	0.072	0.136	0.101	0.176	0.131	0.193
U-E	0.069	0.019	0.031	0.027	0.042	0.032	0.061	0.048	0.077
E-E	0.102	0.017	0.039	0.042	0.063	0.036	0.055	0.069	0.086

Note: The table reports the standard deviation of key labor market variables in the data and in the model under alternative wage-setting and on-the-job search (OJS) specifications. Model series are generated by simulating the full-information and imperfect-information versions of the model under each specification and HP-filtering the simulated data with smoothing parameter $\lambda_{\text{HP}} = 1{,}600$.

staggered wages, search on the job, and gradual belief updating generates sizable amplification of labor market fluctuations.

B.6 List of Model Equations

This subsection collects the core equilibrium conditions of the model used in the quantitative analysis. Unless otherwise noted, period-t choices are made at the beginning of t based on $\mathcal{I}^{t-1}$ and on the prior belief about the persistent productivity component, while beliefs are updated at the end of t after observing $(z_t, \hat{s}_t)$.

Information and beliefs. Aggregate productivity satisfies

$$z_t = x_t + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{28}$$

$$x_t = \rho x_{t-1} + \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \quad |\rho| < 1, \tag{29}$$

and agents observe a noisy public signal about the persistent component:

$$\hat{s}_t = x_t + a_t, \tag{30}$$

$$a_t = \rho_a a_{t-1} + \nu_t, \qquad \nu_t \sim \mathcal{N}(0, \sigma_\nu^2), \quad |\rho_a| < 1. \tag{31}$$

Let the state and measurement vectors be

$$\xi_t \equiv \begin{pmatrix} x_t \\ a_t \end{pmatrix}, \qquad y_t \equiv \begin{pmatrix} z_t \\ \hat{s}_t \end{pmatrix},$$

with linear state-space representation

$$\xi_t = F\xi_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0, Q), \tag{32}$$

$$y_t = H\xi_t + v_t, \qquad v_t \sim \mathcal{N}(0, R). \tag{33}$$

Let $\xi_{t|t-1}$ and $P_{t|t-1}$ denote the prior mean and covariance, and $\xi_{t|t}$ and $P_{t|t}$ the posterior. The Kalman filter recursions are

$$\xi_{t|t-1} = F\xi_{t-1|t-1}, \qquad P_{t|t-1} = FP_{t-1|t-1}F' + Q, \tag{34}$$

$$\iota_t \equiv y_t - H\xi_{t|t-1}, \qquad S_t \equiv HP_{t|t-1}H' + R, \tag{35}$$

$$G_t \equiv P_{t|t-1}H'S_t^{-1}, \qquad \xi_{t|t} = \xi_{t|t-1} + G_t\iota_t, \tag{36}$$

$$P_{t|t} = (I - G_tH)P_{t|t-1}. \tag{37}$$

The posterior mean of the persistent component is $x_{t|t}$ (the first element of $\xi_{t|t}$). The prior used for decisions at the beginning of $t+1$ is

$$x_{t+1|t} = \rho\, x_{t|t}. \tag{38}$$

Search, matching, and composition of matches. Effective labor for a firm is

$$l_t = g_t + \phi b_t, \tag{39}$$

and in the aggregate

$$\bar{u}_t = 1 - \bar{g}_t - \bar{b}_t, \tag{40}$$

where $\bar{g}_t$ and $\bar{b}_t$ denote the aggregate stocks of good and bad matches.

Unemployed workers search with intensity s_{ut} and employed workers in bad matches search on the job with intensity s_{bt}. The efficiency units of search are

$$\bar{s}_t = s_{ut}\,\bar{u}_t + \sigma s_{bt}\bar{b}_t, \tag{41}$$

and the aggregate matching function is

$$\bar{m}_t = \Psi\, \bar{s}_t^{\alpha}\bar{v}_t^{1-\alpha}, \tag{42}$$

where $\bar{v}_t$ denotes aggregate vacancies. The probability that a unit of search yields a match

is

$$p_t = \frac{\bar{m}_t}{\bar{s}_t}, \tag{43}$$

with good and bad match probabilities

$$p_t^g = \xi\, p_t, \tag{44}$$

$$p_t^b = (1-\xi)\, p_t. \tag{45}$$

For a vacancy, the probability of a match is

$$q_t^m = \frac{\bar{m}_t}{\bar{v}_t}, \tag{46}$$

and the probabilities of hiring a good or bad worker are

$$q_t^g = \xi\, q_t^m, \tag{47}$$

$$q_t^b = (1-\xi)\left(1 - \frac{\sigma s_{bt} \bar{b}_t}{\bar{s}_t}\right) q_t^m. \tag{48}$$

The expected efficiency units of labor per vacancy are

$$q_t = q_t^g + \phi q_t^b, \tag{49}$$

and the hiring rate (in efficiency units) is

$$\chi_t = \frac{q_t v_t}{l_t}. \tag{50}$$

The laws of motion for good and bad matches are

$$\bar{g}_{t+1} = \sigma \bar{g}_t + \xi p_t \bar{s}_t, \tag{51}$$

$$\bar{b}_{t+1} = \sigma\Big(1 - s_{bt}\, \xi\, p_t\Big)\bar{b}_t + (1-\xi) p_t s_{ut} \bar{u}_t. \tag{52}$$

Firm problem and hiring condition. Output at the firm level is produced with Cobb–Douglas technology:

$$y_t = Z_t k_t^{\zeta} l_t^{1-\zeta}, \tag{53}$$

where $Z_t = \exp(z_t)$ is aggregate productivity. The representative firm's value is

$$F_t = \max_{k_t, \chi_t} \mathbb{E}_t \left\{ Z_t k_t^{\zeta} l_t^{1-\zeta} - \frac{\kappa}{1+\eta_h} \chi_t^{1+\eta_h} l_t - w_t l_t - r_t k_t + \Lambda_{t,t+1} F_{t+1} \right\}, \tag{54}$$

subject to (39), (50), and the laws of motion for g_t, b_t.

The first-order conditions for capital and hiring are

$$\text{capital:} \quad \mathbb{E}_t \left[Z_t \zeta \left(\frac{l_t}{k_t} \right)^{1-\zeta} - r_t \right] = 0, \tag{55}$$

$$\text{hiring:} \quad \kappa \chi_t^{\eta_h} = \mathbb{E}_t \left[\Lambda_{t,t+1} J_{t+1} \right]. \tag{56}$$

Define the value of a marginal efficiency unit of labor as

$$J_t \equiv \frac{F_t}{l_t}. \tag{57}$$

Household problem and Euler equation. The representative household solves

$$\Omega_t = \max_{\bar{c}_t, \bar{k}_{t+1}} \mathbb{E}_t \left\{ \log(\bar{c}_t) + \beta \Omega_{t+1} \right\} \tag{58}$$

subject to the budget constraint

$$\begin{aligned} \bar{c}_t + \bar{k}_{t+1} + c(s_{bt}) \sigma \bar{b}_t + c(s_{ut}) \bar{u}_t = \bar{w}_t \bar{g}_t + \phi \bar{w}_t \bar{b}_t + \bar{u}_t b \\ + (1 - \delta + r_t) \bar{k}_t + T_t + \Pi_t, \end{aligned} \tag{59}$$

and the laws of motion (51)–(52). The implied stochastic discount factor is

$$\Lambda_{t,t+1} = \beta \frac{\bar{c}_t}{\bar{c}_{t+1}}, \tag{60}$$

and the Euler equation for capital is

$$\frac{1}{\bar{c}_t} = \beta \, \mathbb{E}_t \left[\frac{1 + r_{t+1} - \delta}{\bar{c}_{t+1}} \right]. \tag{61}$$

Worker values and search intensities. Let U_t denote the value of unemployment, and V_t^g and V_t^b the values of good and bad matches. The flow benefit from unemployment is b, and the search cost function is $c(s_{jt}) = \mu s_{jt}^{1+1/\omega}$ for $j \in \{u, b\}$.

The value of unemployment is

$$U_t = \max_{s_{ut}} \mathbb{E}_t \Big\{ b - c(s_{ut}) + \Lambda_{t,t+1} \Big[\big(1 - s_{ut} p_t\big) U_{t+1} + s_{ut}(1-\xi) p_t V^b_{t+1} + s_{ut} \xi p_t V^g_{t+1} \Big] \Big\}. \tag{62}$$

The value of a good match is

$$V^g_t = \mathbb{E}_t \Big\{ w_t + \Lambda_{t,t+1} \Big[\sigma V^g_{t+1} + (1-\sigma) U_{t+1} \Big] \Big\}, \tag{63}$$

and the value of a bad match is

$$\begin{aligned} V^b_t = \max_{s_{bt}} \mathbb{E}_t \Big\{ & \phi w_t - \sigma c(s_{bt}) \\ & + \Lambda_{t,t+1} \Big[(1-\sigma) U_{t+1} + \sigma \big(1 - s_{bt}\, \xi p_t\big) V^b_{t+1} + \sigma s_{bt}\, \xi p_t V^g_{t+1} \Big] \Big\}. \end{aligned} \tag{64}$$

Wage contracts. Define the worker surplus in a good match as

$$H_t \equiv V^g_t - U_t, \tag{65}$$

and recall J_t from (57). Wages for good matches are determined by staggered Nash bargaining. In periods when a contract is renegotiated, the contract wage w^N_t solves

$$w^N_t = \arg\max_w \; H_t(w)^\gamma J_t(w)^{1-\gamma}, \tag{66}$$

where $\gamma \in (0,1)$ is the worker's bargaining weight. With probability λ the previous contract is kept, and with probability $1-\lambda$ the wage is renegotiated. The average wage evolves approximately as

$$\bar{w}_t = (1-\lambda)\bar{w}^N_t + \lambda \bar{w}_{t-1}, \tag{67}$$

where

$$\bar{w}_t = \int_{w,\vartheta} w \, dG_t(w,\vartheta), \tag{68}$$

$$\bar{w}^N_t = \int_{w,\vartheta} w^N_t(\vartheta) \, dG_t(w,\vartheta), \tag{69}$$

and G_t is the cross-sectional distribution of wages and composition across firms.

Resource constraint and government budget. Aggregate output is

$$\bar{y}_t = Z_t \bar{k}_t^{\zeta} \bar{l}_t^{1-\zeta}, \qquad \bar{l}_t \equiv \bar{g}_t + \phi \bar{b}_t, \tag{70}$$

and the resource constraint is

$$\bar{y}_t = \bar{c}_t + \bar{k}_{t+1} - (1-\delta)\bar{k}_t + \frac{\kappa}{1+\eta_h} \int_i \chi_t^{1+\eta_h} l_t \, di + c(s_{bt})\, \sigma \bar{b}_t + c(s_{ut}) \bar{u}_t. \tag{71}$$

The government budget constraint is

$$T_t + \left(1 - \bar{g}_t - \bar{b}_t\right) b = 0. \tag{72}$$

Recursive equilibrium. Given the laws of motion for (x_t, a_t) and the information structure described in (28)–(38), a recursive equilibrium is a collection of value functions $\{F_t, \Omega_t, U_t, V_t^g, V_t^b\}$, decision rules $\{\bar{c}_t, \bar{k}_{t+1}, k_t, \chi_t, s_{ut}, s_{bt}, v_t\}$, prices $\{r_t, w_t^N, \bar{w}_t\}$, and laws of motion for $\{\bar{g}_t, \bar{b}_t, \bar{u}_t\}$ such that: (i) households solve (58)–(61); (ii) firms solve (54)–(56); (iii) search and matching satisfy (41)–(52); (iv) wages satisfy (66)–(67); (v) the resource constraint (71) and government budget (72) hold; and (vi) beliefs evolve according to the Kalman filter (34)–(37) with pre-commitment timing.

References

Barnichon, Regis (2010). "Building a composite help-wanted index". In: *Economics Letters* 109.3, pp. 175–178.

Barnichon, Regis and Christian Brownlees (2019). "Impulse response estimation by smooth local projections". In: *Review of Economics and Statistics* 101.3, pp. 522–530.

Blanchard, Olivier J, Jean-Paul L'Huillier, and Guido Lorenzoni (2013). "News, noise, and fluctuations: An empirical exploration". In: *American Economic Review* 103.7, pp. 3045–70.

Bloom, Nicholas (2009). "The impact of uncertainty shocks". In: *Econometrica* 77.3, pp. 623–685.

Chahrour, Ryan and Kyle Jurado (2018). "News or noise? The missing link". In: *American Economic Review* 108.7, pp. 1702–36.

Fernald, John (2014). *A quarterly, utilization-adjusted series on total factor productivity*. Working Paper 2012-19. Federal Reserve Bank of San Francisco.

Wright, Randall (1986). "Job search and cyclical unemployment". In: *Journal of Political Economy* 94.1, pp. 38–55.