Why Are Some Recoveries Weak and Others Strong?

Paula Donaldson
Brown University

Johannes F. Wieland*
S.F. Fed, U.C. San Diego, NBER

September 14, 2026

Abstract

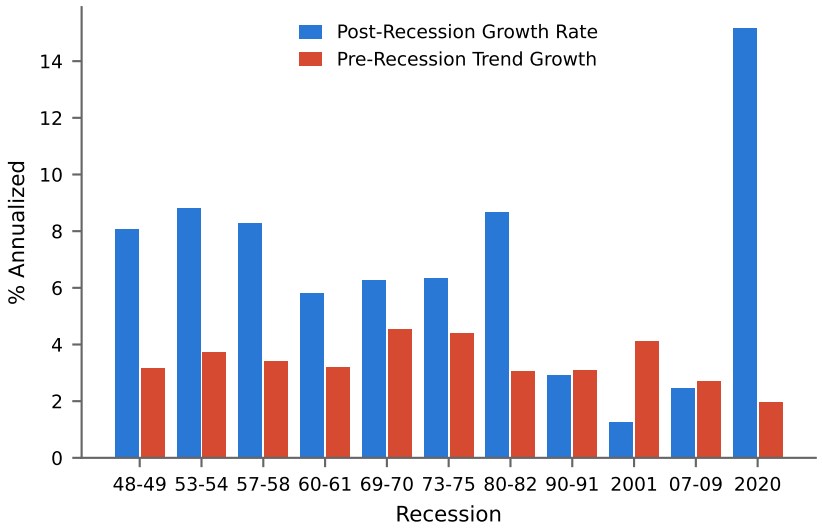
Why were the recoveries from the 1990-1, 2001, and 2007-9 recessions weak relative to other postwar recessions? Leveraging heterogeneous exposure to the national business cycle across U.S. states, we estimate the trajectory of more exposed U.S. states relative to less exposed U.S. states. For the 1990-1, 2001, and 2007-9 recessions we estimate that more exposed states experienced a stronger boom-bust cycle and for the other postwar recessions we estimate a deeper V-shaped recession in more exposed states. The cross-sectional estimates plausibly aggregate to scaled national recession effects and support the view that different types of shocks explain the difference in recovery strength across postwar U.S. recessions. Closing the output gap requires very different monetary policy responses across these recession types.

*Donaldson: paula_donaldson@brown.edu. Wieland: jfwieland@ucsd.edu. We are grateful to seminar participants and Andrew Atkeson, Adrien Auclert, Olivier Coibion, James Hamilton, Amy Handlan, Joshua Hausman, Juan Herreño, Nir Jaimovich, John Mondragon, and Valerie Ramey for helpful comments. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1 Introduction

The recoveries from the 1990-1, 2001, and 2007-9 recessions have been weak. The blue bars in Figure 1 show that private domestic demand[1] grows by only 2.2% on average in the four quarters following these recessions, much less than the 8.4% average growth rate in the four quarters following other recessions. The red bars show that growth after each of the 1990-1, 2001, and 2007-9 recessions also falls short of the pre-recession trend, computed as the annualized growth rate from the last business cycle peak to the current business cycle peak, so private domestic demand does not catch up to its pre-recession trend. We call such recoveries weak. This contrasts with the strong recoveries from the other recessions, in which recovery growth exceeds the pre-recession trend and demand catches back up. This paper asks why the recoveries from the 1990-1, 2001, and 2007-9 recessions are weak and why the recoveries from the other recessions are strong.

Figure 1: Growth of Private Domestic Demand after Recession Trough



Note: The post-recession growth rate is the growth rate of private domestic demand (consumption expenditures plus private investment) for the four quarters after the recession trough. The pre-recession trend is the annualized growth rate between the prior business cycle peak and the current business cycle peak, except for 1948-9 where the trend is linearly extrapolated from 1947Q1-1948Q4.

An extensive literature has sought to explain why only some recessions are followed by

[1]The sum of consumption expenditures and private investment. The patterns for GDP are similar except around the Korean War, Vietnam War, and Afghanistan War: Military expenditure for the Korean War makes the 1948-9 recovery and the 1953-4 pre-trend look stronger; the Vietnam War raises the 1969-70 pre-trend growth rate and reduces the post-recession growth rate; and the Afghanistan War raises the post-2001 growth rate. The recession classification is unchanged except that the 1969-70 is weak on the basis of GDP.

weak recoveries: Hysteresis (Reifschneider, Wascher, and Wilcox, 2015; Benigno and Fornaro, 2018; Eggertsson, Mehrotra, and Robbins, 2019; Garga and Singh, 2021; Cerra, Fatás, and Saxena, 2023), changes in beliefs (Kozlowski, Veldkamp, and Venkateswaran, 2019, 2020), bad luck (Galí, Smets, and Wouters, 2012), an exogenous slowdown in trend growth (Fernald, Hall, Stock, and Watson, 2017; Fernald, Inklaar, and Ruzic, 2025), investment boom-bust cycles (Beaudry, Galizia, and Portier, 2018; Rognlie, Shleifer, and Simsek, 2018; Schaal and Taschereau-Dumouchel, 2023), credit boom-bust cycles (Jordà, Schularick, and Taylor, 2013; Mian, Sufi, and Verner, 2017; Bordalo, Gennaioli, Shleifer, and Terry, 2026), different sectoral incidences of shocks (Beraja and Wolf, 2022), and the secular decline in manufacturing (Leamer, 2021). The debate is not only academic but matters for real policy decisions: the 2021 *Economic Report of the President* argued that "the pre-COVID U.S. economy possessed fewer macroeconomic vulnerabilities than it had in the lead-up to the Great Recession" (p. 52) and that this difference would support a more robust recovery. In contrast, Treasury Secretary Janet Yellen argued that without the American Rescue Plan "we would have a long, slow recovery, like we had after the financial crisis."[2]

In this paper, we propose a new identification strategy based on regional data to discriminate between these theories and uncover the main economic mechanisms that can explain the slow recoveries from the 1990-1, 2001, and 2007-9 recessions relative to the other postwar recessions. We start with the observation that the theories imply different counterfactual paths for aggregate output. Figure 2a illustrates these counterfactuals around a hypothetical GDP path similar to the Great Recession. Hysteresis or changes in beliefs due to the recession imply a persistent gap between realized GDP and its counterfactual path absent the recession. An exogenous slowdown in trend growth or bad luck instead implies that growth would have been weak even absent the recession. These theories predict that counterfactual and realized GDP path converge as the effect of the recession dies out. And boom-bust explanations imply that GDP was higher than the counterfactual prior to the recession, and following the recession converges to counterfactual GDP.

The challenge in testing these differential predictions is that the counterfactual path of aggregate output is unobserved. Common strategies to recover counterfactuals include extrapolating from existing trends (e.g., Blanchard, Cerutti, and Summers, 2015), extrapolating from other countries or past recessions (e.g., Jordà, Schularick, and Taylor, 2013), or using a structural model to infer shocks and counterfactuals (e.g., Galí, Smets, and Wouters, 2012). That disagreement about the correct counterfactual remains reflects that each of these methods embodies important trade-offs. Trend extrapolation holds the country and the recession fixed, but must assume that the trend continues on its existing path. In practice,

[2]Interview with Jake Tapper, February 7, 2021. https://www.cnn.com/TRANSCRIPTS/2102/07/sotu.01.html

trend extrapolations are sensitive to the chosen horizon (Ball, 2015), and would be invalid under some theories such as boom-bust cycles or if aggregate output is non-stationary. In comparing a recession to past recessions or across countries one needs to determine which differences need to be controlled for in order to generate a valid counterfactual (Howard, Martin, and Wilson, 2011). Finally, structural models may imbue the researchers priors, such as over what mechanisms are most important, as well as model misspecification in the analysis (Kehoe, Midrigan, and Pastorino, 2018).

Figure 2b illustrates how regional data can help recover a recession counterfactual. For simplicity, suppose that there are two regions, one is completely exposed and one is completely unexposed to the national recession. If the two regions are on parallel trends absent the recession, then the unexposed region provides the correct regional counterfactual for the exposed region. If these regional counterfactuals are also informative about the aggregate recession counterfactual, then regional data helps discriminate among theories for weak recoveries.

Section 2 formalizes these identification and aggregation problems. In the regional setting we can identify the *relative recession impulse response function* (RRIRF): the relative effect on output in a region that experiences a 1% deeper recession due to its greater exposure to the aggregate business cycle shock compared to another region (Figure 2d). To identify the RRIRF we need an instrument for national recession exposure, since variation based on realized relative recession depth across regions will conflate the national recession with local shocks. We measure exposure to the national recession using a Bartik instrument that interacts regional industry shares with the cyclical component of national industry growth. In the first stage we predict relative output growth over the national recession using this Bartik predicted recession depth. In the second stage we use instrumented relative output growth over the recession to predict relative output growth before, during, and after the recession. If more and less exposed regions are on parallel trends conditional on controls, then this IV strategy recovers the RRIRF.

Parallel trends requires that regions that are differentially exposed to the aggregate recessions are not differentially exposed to local shocks. This is the standard Bartik assumption, which we assess by isolating variation from particular industries and overidentification tests following Goldsmith-Pinkham, Sorkin, and Swift (2020). Parallel trends also requires that regions are not on differential pre-trends conditional on controls. In our specifications we control for the Bartik-implied trend and a local trend far in advance of the recession, and we assess the sensitivity of our results to different ways of constructing these trends.

An advantage of the regional analysis is that the RRIRF remains identified even if national recession exposure is measured with error. This is likely because the Bartik instrument

is constructed using detrended national industry growth, and separating cycle from trend is notoriously challenging (Stock and Watson, 1999; Ball, 2015). However, measurement error in the cross-section only matters to the extent that it violates parallel trends. This means our setting can accommodate time-varying measurement error that is common across industries, classical measurement error, and measurement error correlated with the business cycle. We can even handle idiosyncratic measurement error correlated with industry trends, as this measurement error largely averages out through industrial diversification.[3] In contrast, any measurement error biases the aggregate trend-cycle decomposition and thus the implied non-recession counterfactual in time series analysis.

A second advantage of exploiting regional variation is that all units are subject to the same aggregate shock, so there is a single well-defined *aggregate recession impulse response function* (ARIRF): the aggregate effect on output given a 1% deeper aggregate recession (Figure 2c). In a cross-country setting, we may compare countries subject to different aggregate shocks with heterogeneous ARIRF. Estimates of the RRIRF across countries may then load on differences in the ARIRFs by assigning negative weights to some countries, and therefore not correspond to a well-defined aggregate recession counterfactual.

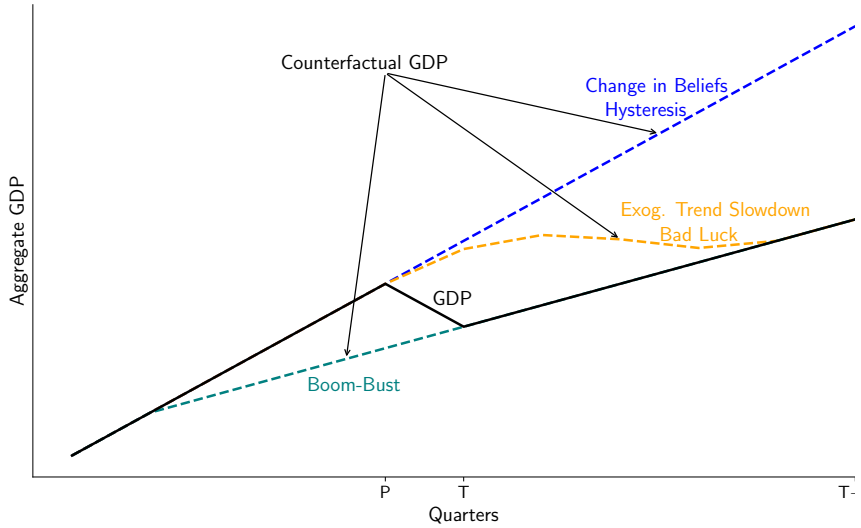
The price of this robustness is that the regional analysis does not directly identify the aggregate recession counterfactual. However, we show that if regional cyclical output has a one-factor structure over aggregate cyclical output in a recession, then the RRIRF exactly coincides with the ARIRF. In this case, the RRIRF inherits the different predicted counterfactuals across theories from the ARIRF and directly discriminates among them. The one-factor structure allows for arbitrary interactions between regions and industries within periods, such as multiplier effects, input-output structures, price stickiness, and regional and industry spillovers, so long as these interactions are stable over the cycle. It also allows these relationship to vary across recessions. It does rule out regional recession dynamics that differ from the aggregate shock itself, such as regional convergence forces. In line with the literature (Blanchard and Katz, 1992; Yagan, 2019; House, Proebsting, and Tesar, 2026), we find that these forces are sufficiently weak such that the RRIRF closely corresponds to the ARIRF.

Most related to our cross-sectional approach is Beraja, Hurst, and Ospina (2019) who estimate the regional wage Phillips Curve estimates together and a regional DSGE model to determine the relative role of demand and supply shocks during the Great Recession and its recovery. Relative to their work, we directly estimate the RRIRF and impose less structure for aggregation, but our estimates do not speak to the relative importance of demand and supply shocks.

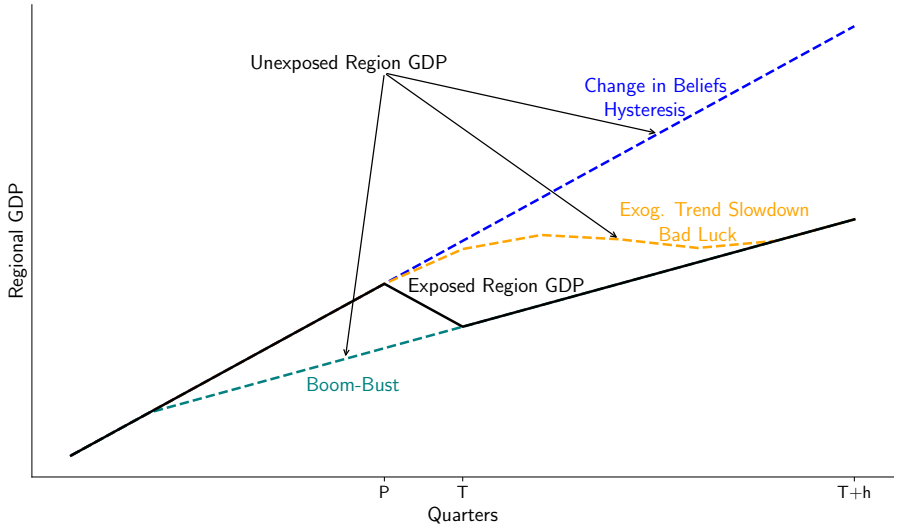
[3] See Appendix E for details.

Figure 2: Explanations for Weak Recoveries

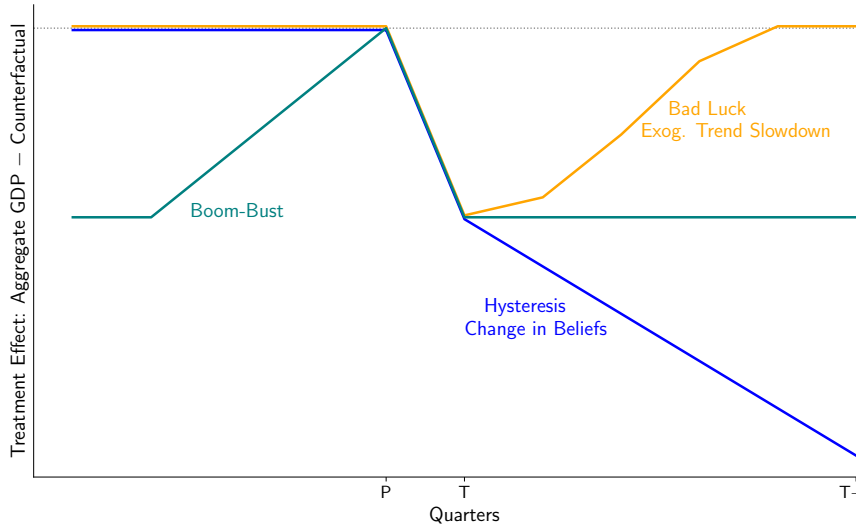
(a) Aggregate Counterfactuals



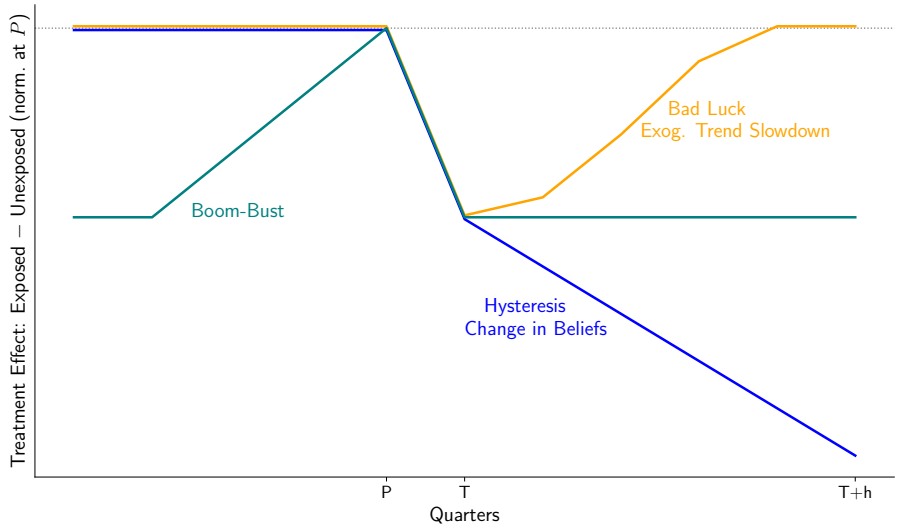
(b) Regional Identification



(c) Aggregate Recession Impulse Response Function



(d) Relative Recession Impulse Response Function



In Section 3 we construct the cyclical Bartik instrument based on state industry shares from the quarterly state personal income accounts at the beginning of all post-war recessions and combine them with detrended national industry growth rates from the NBER recession peak to the NBER recession trough. We use leave-one-out growth rate to remove any mechanical correlation of the Bartik with local shocks.

We illustrate how the data inform our estimates in Section 4 by exploiting the equivalence between the Bartik approach and weighted GMM using exposure shares (Goldsmith-Pinkham, Sorkin, and Swift, 2020). For pre-1990 recessions the Bartik instrument loads most strongly on the manufacturing sector. We show that states with an above-average manufacturing share experience a deeper recession but also a stronger recovery after the national recession trough. We find a similar pattern in the Covid recession, in which an above-average hospitality and food service share predicts both a deeper recession and stronger recovery. In contrast, a higher construction share during the Great Recession or higher professional services and IT share in 2001 implies a weaker recovery despite a deeper recession. More exposed states also experience relatively strong growth prior to the latter two recession. These recession-recovery patterns indicate that recessions with weak and strong national recoveries have different RRIRF.

We estimate the RRIRFs in Section 5. For recessions with a strong national recovery (pre-1990 and 2020), we estimate a deeper V-shape recession in more exposed U.S. states relative to less exposed states. The more exposed states see a larger fall in income, GDP, and employment and a larger increase in unemployment during the recession, but then grow more quickly in the recovery compared to less exposed states. We cannot statistically reject a complete relative recovery in these recessions, though some point estimates suggest modest hysteresis effects.

For recessions with a weak national recovery, 1990-1, 2001, and the Great Recession, we find that more exposed states experience a stronger boom-bust cycle.[4] In the recovery phase, the more exposed states do not grow more despite their deeper local recession. However, they do grow relatively more quickly prior to the recession. Thus, the RRIRF displays an inverse V-shape consistent with boom-bust explanations for weak recoveries (Jordà, Schularick, and Taylor, 2013; Beaudry, Galizia, and Portier, 2018; Rognlie, Shleifer, and Simsek, 2018). Again, the point estimates suggest some modest hysteresis effects in the more exposed states relative to less exposed states, but we cannot statistically reject a recovery to the pre-boom trend.

Our estimates for strong national recoveries support plucking models of the business cycle

[4]In practice, this effect is identified entirely by the 2001 and Great Recession due to limited cross-sectional variation in the 1990-1 recession. See Section 5.4 and Appendix Figure 20c.

(Friedman, 1993; Dupraz, Nakamura, and Steinsson, 2025), in which recessions are a drop in output below the balanced growth path. For the weak national recoveries of 1990-1, 2001, and the Great Recession, our evidence is consistent with an "upward" pluck, in which the recession is a drop from an unsustainable boom. To the extent that hysteresis effects are present, these are quantitatively similar across the two types of recessions, which suggests that they are not the main explanation for the different recovery patterns.

We can also rule out that differences in countercyclical transfer policy account for the differential recovery patterns. First, our estimates for the RRIRFs change little when transfers are included or excluded from state income. Second, the transfer share in income follows the opposite pattern of the RRIRFs: in the 1990-1, 2001, and the Great Recession, the transfer share falls during the boom, rises in the recession, and remains persistently elevated in the recovery phase. In the other recessions the transfer share is only briefly elevated during the recession and normalizes quickly during the recovery. This shows that countercyclical transfer policy works, if anything, to offset the differences we find across recessions. We do not have direct evidence on countercyclical monetary policy, but the long-term real interest rate falls more on average in the 1990-1, 2001, and the Great Recession compared to the other recessions, which again suggests that countercyclical policy is shrinking the differences across recessions.[5]

Following Goldsmith-Pinkham, Sorkin, and Swift (2020) we decompose our RRIRF into the RRIRFs implied by exposure to a particular sector and the associated sectoral weights. This allows us to determine whether changes in the RRIRF across recessions are due to changes in recession trajectories from particular sectoral exposures or changes in sectoral exposure weights. Across broad sectoral and recession exposures, we find evidence of V-shaped RRIRF in more exposed states for the pre-1990 recessions and COVID-19, and boom-bust RRIRF for the 1990-1, 2001, and Great Recession. We also do not reject overidentification restrictions at almost all horizons. This suggests that changing recovery patterns extend beyond the manufacturing sector emphasized by Leamer (2021). Furthermore, changes in sectoral weights play a minor role in explaining weak versus strong recoveries.

In short, our evidence supports the view that the difference in recovery strength in Figure 1 is best explained by the recessions being caused by different types of shocks. We then show that distinguishing between these two types of recessions matters for policy. We map the RRIRFs for the two recession types into paths for the output gap by exploiting our aggregation results, and by assuming that the output gap is zero prior to the V-shaped

[5]This is despite the 1990-1, 2001, and the Great Recession being on average slightly shallower, measured in terms of GDP or private domestic demand relative to the pre-recession trend, compared to the other recessions and despite the presence of the ZLB.

recession and prior to the boom-bust cycle. Following Barnichon and Mesters (2023), McKay and Wolf (2023), and Caravello, McKay, and Wolf (2024), we combine these paths with the causal effects matrix of monetary policy on output to determine the counterfactual long-run real interest rate R^* that closes current and future output gaps during the cycle.

Our counterfactual calls for long-term real interest rates to be one percentage point lower than realized around the trough of V-shaped recessions. In boom-bust cycles the long-run R^* is about two percentage points higher in the boom and one percentage point lower in the bust than realized. Comparing across recession types, boom-bust cycles display both much higher R^* in the boom and lower R^* in the recovery then V-shaped recessions, despite the same recession depth. This suggests that appropriately calibrating countercyclical monetary policy depends importantly on a correct diagnosis of the recession.

2 Identification and Aggregation

We next describe the objects of interest, their data generating process, and identification. We then discuss the aggregation problem from regional to aggregate objects. We work with a general DGP and show how a structural model maps into the DGP in Appendix A.

2.1 The Object of Interest We assume that log aggregate output y_t can be decomposed into a cyclical component $\hat{y}_t$ and a trend component $\bar{y}_t$,

$$y_t = \hat{y}_t + \bar{y}_t. \tag{1}$$

To the extent that business cycles have permanent hysteresis effects, we include them in the cyclical component $\hat{y}_t$.

We take as given the path of cyclical output $\{\hat{y}_t\}$ around a particular recession. Let P and T be the recession peak and trough respectively, so $\hat{y}_T - \hat{y}_P$ is the cyclical recession depth. The dynamics of $\{\hat{y}_t\}$ are encoded in the *aggregate recession impulse response function* (ARIRF) which is defined as,

$$ARIRF_h \equiv \frac{\hat{y}_{T+h} - \hat{y}_P}{\hat{y}_P - \hat{y}_T} \qquad \Rightarrow \qquad \hat{y}_{T+h} - \hat{y}_P = ARIRF_h(\hat{y}_P - \hat{y}_T) \tag{2}$$

The ARIRF measures what fraction of the cyclical recession depth persists h period after the recession trough (or before if $h < 0$). It is equal to -1 at $h = 0$ and to 0 at $h = P - T$, which are normalizations to level and scale. Given a path $\{\hat{y}_t\}$ around a recession, we can always define the $ARIRF_h$ in this way. But from an ARIRF we cannot derive the path $\{\hat{y}_t\}$, because the level and scale of the recession is not known.

The ARIRF may be recession-specific. As shown in Figure 2c, explanations for slow recoveries differ in their predictions for the ARIRF. For example, the exogenous trend slowdown explanation implies that the ARIRF converges to zero as the cyclical component vanishes, whereas hysteresis explanations predict a permanent cyclical component. Boom-bust cycles predict that the ARIRF is increasing in the boom period prior to the recession peak. Since the ARIRF can discriminate among these theories it is our object of interest.

But the ARIRF cannot be measured directly for the same reason that one cannot directly measure the cyclical path $\{\hat{y}_t\}$: the cyclical component $\hat{y}_t$ is unobserved. And filtering $\hat{y}_t$ from y_t based on structural or time-series models is notoriously challenging, particularly if $\hat{y}_t$ contains permanent components or is non-zero prior to a recession. This challenge motivates our cross-sectional approach.

2.2 Cross-Sectional Identification We next show that we can more reliably identify the *relative recession impulse response function* (RRIRF) in the cross-section. We then show that the RRIRF is closely related to the ARIRF.

There are I regions indexed by i. Log regional output can be decomposed into a cyclical, a trend component, and an idiosyncratic component,

$$y_{it} = \hat{y}_{it} + \bar{y}_{it} + \mathring{y}_{it}. \tag{3}$$

where the idiosyncratic component $\mathring{y}_{it}$ is weakly correlated across regions and aggregates to zero.

We define the relative recession impulse response function for location i as,

$$RRIRF_{ih} \equiv \frac{\hat{y}_{i,T+h} - \hat{y}_{iP} - (\hat{y}_{T+h} - \hat{y}_P)}{\hat{y}_{iP} - \hat{y}_{iT} - (\hat{y}_P - \hat{y}_T)} \tag{4}$$

The RRIRF measures what fraction of the *relative* recession of location i persists h periods after the national recession trough. Analogous to the ARIRF it is normalized to -1 at $h = 0$ and to 0 at $h = P - T$.

The main restriction we impose is that the RRIRF is not location specific,

$$RRIRF_{ih} = RRIRF_h. \tag{5}$$

Homogeneity of the RRIRF across regions is consistent with our overidentification tests, and the evidence in Angeletos, Collard, and Dellas (2020) that aggregate business cycle dynamics are best explained by one shock.[6] Violation of homogeneity also does not invalidate

[6]Homogeneity will fail if aggregate dynamics are driven by two or more shocks with different dynamics and regions load differentially on those shocks. Donaldson (2026) shows how to recover causal effects in these

our identification, but implies that our estimator will converge to a weighted average of location-specific RRIRF. [7] Implicit in (5) is that local cyclical output scales linearly with local recession depth. This linearity is less restrictive than homogeneity, as $RRIRF_h$ can always be interpreted as a linear approximation to the true nonlinear RRIRF.

The RRIRF cannot be identified by a regression of output growth on the output decline over the recession,

$$y_{i,T+h} - y_{i,P} = \delta_h + \beta_h(y_{i,P} - y_{i,T}) + \Omega X_{iP} + u_{i,P,T+h}. \tag{6}$$

This is because local recession depth and output growth are also determined by the trend and idiosyncratic components, which are unlikely to follow the same dynamics as the cyclical component $\hat{y}_{it}$.

To identify $RRIRF_h$ we take a Bartik approach. There are G industries indexed by g. Log output in industry g is again composed of a cyclical, a trend, and an idiosyncratic component,

$$y_{gt} = \hat{y}_{gt} + \bar{y}_{gt} + \mathring{y}_{gt}. \tag{7}$$

where the idiosyncratic component $\mathring{y}_{gt}$ also aggregates to zero and is weakly correlated across industries.

We assume that we can observe the trend component with error,

$$\tilde{\bar{y}}_{gt} \equiv \bar{y}_{gt} + \varepsilon_{gt}. \tag{8}$$

Importantly, we allow for ε_{gt} to have arbitrary mean, so that the industry trend component may be systematically over- or underestimated. For example, if aggregate trend growth slows because of weaker productivity or population growth, then extrapolating from prior trends overstates the cyclical component in all industries. Such systematic bias implies that one cannot uncover $\hat{y}_t = y_t - \bar{y}_t$ by aggregating $\tilde{\bar{y}}_{gt}$ over many industries and appealing to a law of large numbers. In fact, a result of our approach is that $\{\hat{y}_t\}$ remains unknown even when the RRIRF is identified and equal to the ARIRF, because $\{\hat{y}_t\}$ cannot be uniquely inferred from the ARIRF.

settings.

[7] Goldsmith-Pinkham, Sorkin, and Swift (2020) (Assumption 3, Proposition 4) show that a location-industry-specific first stage that is weakly monotone in sign across locations is sufficient for the just-identified estimate from any single industry's share to be a nonnegatively-weighted average of location-specific effects. Appendix Figure 15 shows first stage monotonicity for the industry shares that explain most of the Bartik variation. Nonnegative weights on the overall Bartik estimate additionally require that the Rotemberg weights α_k combining across industries not be predominantly negative (Goldsmith-Pinkham, Sorkin, and Swift (2020)). Section 5.4 shows this second condition.

Based on the mismeasured trend, we extract the mismeasured cyclical component,

$$\tilde{\hat{y}}_{gt} = y_{gt} - \tilde{\bar{y}}_{gt} = \hat{y}_{gt} + \overset{\circ}{y}_{gt} + \varepsilon_{gt}, \tag{9}$$

which includes both trend mismeasurement as well as any idiosyncratic industry shocks.

We use this cyclical component to construct Bartik predicted recession depth based on industry shares in each location, s_{ig},

$$\tilde{\hat{y}}^b_{i,P,T} = \sum_{g=1}^{G} s_{ig}(\tilde{\hat{y}}_{g,P} - \tilde{\hat{y}}_{g,T}) \tag{10}$$

$$= \sum_{g=1}^{G} s_{ig}[(\hat{y}_{g,P} - \hat{y}_{g,T}) + (\overset{\circ}{y}_{g,P} - \overset{\circ}{y}_{g,T}) + (\varepsilon_{g,P} - \varepsilon_{g,T})]. \tag{11}$$

where $\tilde{\hat{y}}_{g,P} - \tilde{\hat{y}}_{g,T}$ is the imperfectly measured recession depth in industry g. We assume that idiosyncratic industry shocks are sufficiently weakly correlated across industries and regional industry shares are sufficiently diversified that $\sum_{g=1}^{G} s_{ig}(\overset{\circ}{y}_{g,P} - \overset{\circ}{y}_{g,T}) \to 0$ as G grows. Hence, the idiosyncratic component of the Bartik is asymptotically uncorrelated with the non-cyclical component of local output. In our empirical implementation we additionally use leave-one-out industry growth rates to eliminate mechanical correlation with location-specific shocks.

We then use Bartik predicted recession depth as an instrument for output growth over the recession in the first stage,

$$\text{First Stage:} \qquad y_{i,P} - y_{i,T} = \gamma_0 + \gamma_1 \tilde{\hat{y}}^b_{i,P,T} + \Gamma X_{iP} + \nu_{i,T}. \tag{12}$$

In the second stage we project output growth relative to the recession peak on the instrumented local output growth during the recession,

$$\text{Second Stage:} \qquad y_{i,T+h} - y_{i,P} = \delta_h + \beta_h(\widehat{y_{i,P} - y_{i,T}}) + \Omega X_{iP} + u_{i,P,T+h}, \tag{13}$$

The fact that $\tilde{\hat{y}}^b_{i,P,T}$ is constant for each i implies that the regression exploits only cross-sectional variation. Specifically, we compare the output trajectory of locations that have a deeper predicted recession with the output trajectory of locations that have a shallower predicted recession based on their industry structure.

This IV specification identifies the RRIRF if predicted recession depth is uncorrelated with the non-cyclical component of local output conditional on the controls X_{iP}.

Proposition 1 *If Bartik predicted recession depth is correlated with recession depth condi-*

tional on controls (relevance),

$$Cov_i[\tilde{\hat{y}}^b_{i,P,T}, \hat{y}_{i,P} - \hat{y}_{i,T} | X_{iP}] \neq 0 \tag{14}$$

and uncorrelated with differential local trends or idiosyncratic local shocks conditional on controls (exclusion),

$$Cov_i[\tilde{\hat{y}}^b_{i,P,T}, \bar{y}_{i,T+h} - \bar{y}_{i,P} + \mathring{y}_{i,T+h} - \mathring{y}_{i,P} | X_{iP}] = 0 \qquad \forall \ h \tag{15}$$

then as $I \to \infty$ the IV estimator of (12)-(13) converges to the RRIRF,

$$\beta^{IV}_h \to RRIRF_h. \tag{16}$$

Proof See Appendix C. ■

The relevance condition is standard. The exclusion restriction requires that locations that we measure to be more cyclically sensitive are not on otherwise different non-cyclical output trajectories than locations that we measure to be less cyclically sensitive conditional on controls. It is the parallel trends assumption in this setting.

There are three main threats to identification. The first is that regions that are measured to be more cyclically sensitive experience disproportionally positive or negative local shocks $\mathring{y}_{i,T+h} - \mathring{y}_{i,P}$ around the recession. This is the standard concern in Bartik settings. To address this concern we test the sensitivity of our estimates β_h to using variation from individual industries as well as overidentification restrictions implicit in the Bartik (Goldsmith-Pinkham, Sorkin, and Swift, 2020). Note that we do not take a stance on whether variation in the Bartik instrument is due to demand or supply shocks, but instead require that it reflects the aggregate business cycle shock (which may be demand or supply).

A second threat is that cyclical sensitivity is correlated with differential trends $\bar{y}_{i,T+h} - \bar{y}_{i,P}$. For example, locations with relatively large manufacturing shares are likely to be both more cyclically sensitive and on a relative downward trend. We control for these trends in two ways: with a Bartik trend control that combines local industry shares with our estimate of trend growth, $\sum_{g=1}^{G} s_{ig}(\tilde{\bar{y}}_{g,P} - \tilde{\bar{y}}_{g,T})$, and by controlling for average local growth sufficiently far in advance of the recession peak, $\Delta^T y_{i,P-h_{pre}} = y_{i,P-h_{pre}} - y_{i,P-h_{pre}-T}$. We cannot control for local growth immediately prior to the recession as this would also absorb any pre-recession cyclical component, such as the boom in a boom-bust cycle. Because any trend specification is likely subject to error, we examine the sensitivity of our estimates to different trend choices.

The third threat comes from trend mismeasurement $\sum_{g=1}^{G} s_{ig}(\varepsilon_{g,P} - \varepsilon_{g,T})$. Importantly, certain classes of measurement error do not pose an identification problem. Any time-varying

measurement error that is common across industries is absorbed by the intercepts in (12)-(13). Classical measurement error that is uncorrelated with local trends or idiosyncratic shocks also does not violate the exclusion restriction. Both the first stage and reduced form are proportionally attenuated so that β_h remains consistent. For the same reason, identification is not threatened if trend mismeasurement ε_{gt} is correlated with sectoral cyclical recession depth. Intuitively, our approach is robust to such mismeasurement, because it does not violate parallel trends (see Appendix F.1 for details).

The threat to identification is that measurement error in trend growth is correlated with trend growth itself. For example, suppose that industries are subject to trend growth shocks that we fail to measure, $Cov_g(\varepsilon_{g,T} - \varepsilon_{g,P}, \bar{y}_{g,T+h} - \bar{y}_{g,P}) < 0$. If this correlation translates to the local level, between local measurement error $\sum_{g=1}^{G} s_{ig}(\varepsilon_{g,T} - \varepsilon_{g,P})$ and local trend growth based on industry composition $\sum_{g=1}^{G} s_{ig}(\bar{y}_{g,T+h} - \bar{y}_{g,P})$, then we violate parallel trends by measuring greater recession exposure in regions with negative trend growth shocks. However, because each location aggregates over many industries, idiosyncratic measurement errors average out when regional industry shares are sufficiently diversified. Thus, if trend measurement error is idiosyncratic across industries, then $\lim_{G\to\infty} E\{Cov_i[\sum_{g=1}^{G} s_{ig}(\varepsilon_{g,T} - \varepsilon_{g,P}), \sum_{g=1}^{G} s_{ig}(\bar{y}_{g,T+h} - \bar{y}_{g,P})]\} = 0$. Consistent with this assumption, Foerster, Hornstein, Sarte, and Watson (2022) document that industry trend shocks have been almost entirely idiosyncratic.[8] In Appendix E we show that averaging measurement error works well in our setting with finite G, so that the bias from measurement error correlated with trend growth is small.

The ability of the cross-sectional Bartik approach to accommodate all of time-varying aggregate measurement error, classical measurement error across industries, and idiosyncratic non-classical measurement error across industries in the trend-cycle decomposition is a distinct strength. Statistical or model-based filtering procedures on aggregate data are vulnerable precisely to time-varying measurement error in the aggregate trend $\tilde{\bar{y}}_t$, as this would induce misidentification of the cyclical path $\{\hat{y}_t\}$ and therefore the ARIRF. And absent an IV strategy, cross-country comparisons are biased if country-specific trends are either classically or non-classically mismeasured.[9]

A second advantage of our approach is that the RRIRF is less likely to be unit specific, since we study the response of locations to the same aggregate shock. In cross-country studies, countries may be subject to different aggregate shocks with heterogeneous ARIRF.

[8] In principle, the same law-of-large numbers argument may be applied to the predicted recession depth instrument, which would then violate the relevance assumption. However, relevance is directly testable and remains asymptotically valid if there is regional clustering of cyclically sensitive sectors (e.g., manufacturing in the rust belt).

[9] Similarly a cross-industry comparison would also be biased by industry trend measurement error.

Conceptually, it is unclear how to define an object of interest for this case since there is no unique aggregate shock. Furthermore, cross-country comparisons would estimate the extent to which heterogeneity in ARIRF correlates with recession exposure, which may not be informative about any weighted average of the ARIRF across countries.[10]

The price we pay for these two advantages is that we require the Bartik instrument to satisfy the exclusion restriction, and that we cannot identify the ARIRF directly. We assess the plausibility of the exclusion restriction in data below. We next turn to the aggregation problem from the RRIRF to the ARIRF.

2.3 Aggregation A general setting under which the RRIRF isolated by our Bartik design aggregates exactly to the ARIRF is when relative local cyclical output satisfies

$$\hat{y}_{it} - \hat{y}_t = \theta_i \hat{y}_t + \nu_{it} \qquad \text{with} \qquad Cov_i(\tilde{\hat{y}}^b_{i,P,T}, \nu_{it}) = 0. \tag{17}$$

Here θ_i is the exposure of location i to the cyclical component of aggregate output and ν_{it} captures differential local dynamics that do not track the dynamics of the aggregate business cycle. Under the assumption that these local dynamics are uncorrelated with Bartik predicted recession depth, the IV estimate of the RRIRF aggregates as follows

$$\begin{aligned} \widehat{RRIRF}_h^{IV} &= \frac{Cov_i\Big(\tilde{\hat{y}}^b_{i,P,T}, \theta_i(\hat{y}_{T+h} - \hat{y}_P) + (\nu_{i,T+h} - \nu_{i,P})\Big)}{Cov_i\Big(\tilde{\hat{y}}^b_{i,P,T}, \theta_i(\hat{y}_P - \hat{y}_T) + (\nu_{i,P} - \nu_{i,T})\Big)} \\ &\xrightarrow{p} \frac{(\hat{y}_{T+h} - \hat{y}_P)}{(\hat{y}_P - \hat{y}_T)} \frac{Cov_i(\tilde{\hat{y}}^b_{i,P,T}, \theta_i)}{Cov_i(\tilde{\hat{y}}^b_{i,P,T}, \theta_i)} = ARIRF_h \end{aligned} \tag{18}$$

with $Cov_i\Big(\tilde{\hat{y}}^b_{i,P,T}, \theta_i\Big) \neq 0$ implied by the relevance condition in Proposition (1).

Importantly, (17) does not require all relative local cyclical dynamics to have a one-factor structure on the aggregate cycle. It only restricts the component of relative local cyclical output that is correlated with Bartik-predicted recession exposure. Any additional local dynamics ν_{it} may follow arbitrary paths as long as they are orthogonal to the Bartik instrument in the cross-section. Moreover, because both θ_i and the local dynamics may differ across recessions, (17) is still consistent with a multi-factor structure in the time series. Finally, (17) does not restrict local cycles through $\mathring{y}_{it}$, beyond what is already implied by the exclusion restriction in Proposition 1.

[10] Consider the case in which one recession is a -1% decline in output that lasts one year, and the other recession is a -2% decline in output that shrinks to -1% in the following year and then to zero. Comparing these recessions yields a RRIRF of -1% in the first year and -1% in the second year, and thus greater persistence than implied by either recession.

Because θ_i is unrestricted, the structure in (17) allows for arbitrary within-period interactions between regions and industries, so long as these interactions do not change over a given business cycle. Thus, (17) allows for complex trade networks, input-output structures, labor supply, migration, and pricing decisions across and within locations and industries. It also allows these interactions to vary across recessions, so θ_i can be recession-specific.

The main restriction embedded in (17) is that locations with high versus low Bartik-predicted recession depth do not exhibit differential output changes ν_{it} with distinct dynamics from the aggregate business cycle. A potential concern is then that local convergence dynamics are relatively more important in more exposed regions because they suffer a deeper recession, and that these dynamics are distinct from the aggregate cycle. In practice, studies show that the relationship between local cyclical output and the aggregate business cycle is well-represented by a one-factor structure and that any additional aggregate factors, such as convergence forces, add only small additional explanatory power (Carlino and Sill, 2001; Artis, Dreger, and Kholodilin, 2011). This is consistent with the evidence in Barro and Sala-i Martin (1991), Blanchard and Katz (1992), and Yagan (2019) that regional convergence forces are either very slow or absent. Finally, in Appendix D we show that a multi-region model calibrated to this evidence almost exactly aggregates in the sense that $RRIRF_h \approx ARIRF_h$. This is despite us taking a conservative approach by assuming that exposure to convergence forces is perfectly collinear with Bartik-predicted recession depth, which maximizes the scope for convergence forces to bias the RRIRF away from the ARIRF. The aggregation bias is small because regional convergence forces are slow and even more so in a forward-looking model with transitory shocks. In short, this evidence suggests that our estimates of the RRIRF closely correspond to the ARIRF, and can therefore be used to discriminate among theories for slow aggregate recoveries.

3 Data

Our identification strategy requires business-cycle frequency data that is disaggregated by both location and industry for the entire post-war period. We use the quarterly BEA state personal income data, which is disaggregated by industry since 1948. The industry detail increases from 12 SIC sectors from 1948-57 to 14 SIC sectors from 1957-2001 and to 24 NAICS sectors from 1997 onward. Our outcomes are quarterly state personal income, quarterly state unemployment rates from Fieldhouse, Munro, Koch, and Howard (2024), quarterly state employment data from the BLS,[11] and, post-2005, quarterly state-level GDP

[11] We seasonally adjust the raw monthly employment data using the Census X-13 algorithm up to February 2020. We then apply the last year of seasonal factors to the remaining sample. This procedure avoids

from the BEA. We deflate the nominal variables using the national GDP deflator, and convert all data (except unemployment) to per capita using state population estimates from the U.S. Census linearly interpolated to a quarterly frequency.

When using personal income data we aim to closely proxy for local private output. We therefore exclude the government sector from personal income. We also exclude the finance, insurance, and real estate (FIRE) sector and agriculture and mining. For these sectors we find the largest divergence between local GDP growth and local personal income growth when both data are available, likely because the geography of ownership and production are more often different in these sectors than in other sectors.

We follow the standard NBER dating for recession peaks and troughs, except that we combine the 1980 and 1981-2 recession into a single episode. We construct the cyclical component for national industries in the Bartik instrument (10) by removing the NBER-peak to NBER-peak growth rates. We also show robustness to different trend specifications below.[12]

4 Illustration of the Empirical Approach

Goldsmith-Pinkham, Sorkin, and Swift (2020) show that estimation using the Bartik instrument (10) is equivalent to weighted GMM using the shares s_{ig} as instruments.[13] We therefore illustrate the empirical approach based on the sectoral shares that receive the highest weight in the Bartik instrument. Specifically, in Section 5.4 we show that the Bartik instrument loads most strongly on the manufacturing share in pre-1990 recessions, on the construction share in the Great Recession, on the accommodation and food sector during COVID-19, and on IT and professional services in the 2001 recession.

Figure 3 plots the time path of average per-capita State personal income relative to the national recession peak, when sorted by the key sources of variation in the Bartik instrument. The top left panel sorts U.S. states by their manufacturing share in each recession, and then pools across all pre-1990 recessions.[14] The average recession begins five quarters before the

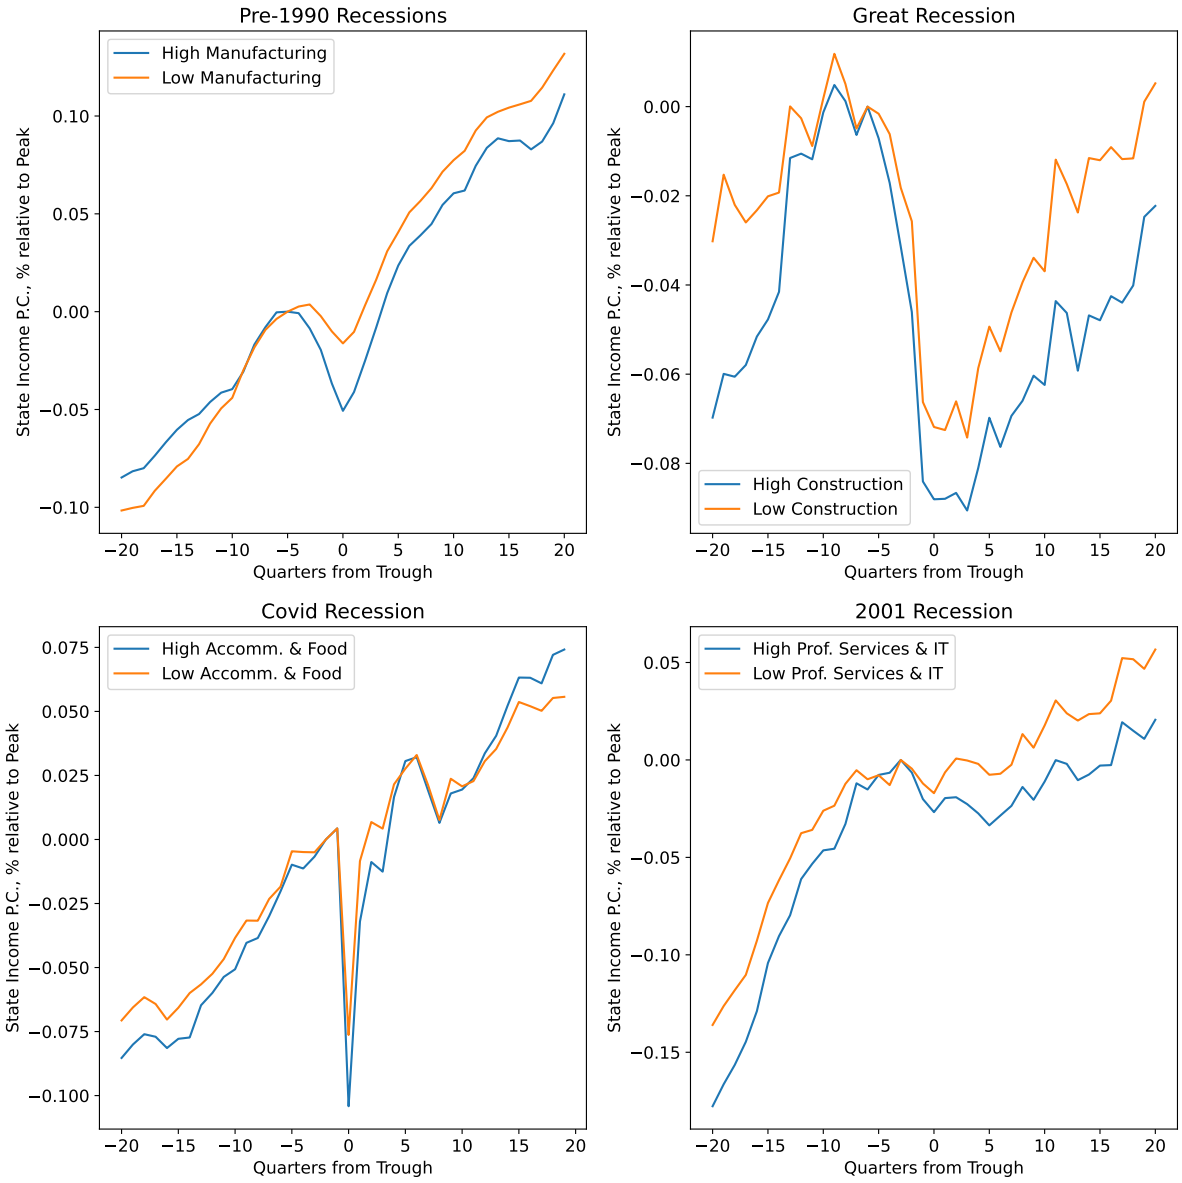
contamination of the seasonal factors by the sharp Covid recession. We then aggregate the seasonally adjusted data to a quarterly frequency.

[12] For the 1948-9 recession and for the SIC-NAICS switch in 1997 we do not have industry data from the prior peak, in which case we proxy the trend using the average growth rate from the available data up to the current peak. In 1957 the Personal Income files switch to a more detailed SIC classification. We aggregate these to the coarser historical classification for the purpose of constructing peak-to-peak growth rates for the 1957-8 recession.

[13] Adao, Kolesár, and Morales (2019) and Borusyak, Hull, and Jaravel (2022) also show that identification is also possible based on exogeneity of the sectoral shocks rather than the sectoral shares. However, their approach requires a large number of independent sectoral shocks, which is not applicable to our setting.

[14] We group states on whether they are above or below the mean share separately for each recession.

Figure 3: Outcomes by Key Sources of Variation of the Bartik Instrument



trough, when the high and low manufacturing groups cross at 0 on the y-axis. The difference between the two groups at the NBER trough (0 on the x-axis) therefore measures the relative recession depth. The high manufacturing states decline by 3.4% more during the recessions relative to the low manufacturing states. In other words, the high manufacturing states are more exposed to the national recession.

After the recession trough, high-manufacturing areas grow more quickly than low-manufacturing areas, which indicates that they catch up from a deeper recession. The high-manufacturing areas grow an additional 1.8% in the ten quarters of the recovery, making up more than half of the relative recession depth. The top left panel also shows that high-manufacturing states grow 0.11% less per quarter relative to low-manufacturing states before the recession. Accounting for this relative trend implies that the relative recovery is complete after 10 quarters.[15]

The pre-1990 recession plot also illustrates that our cross-sectional recession elasticities are robust to misspecification of the aggregate trend, unlike a time-series analysis. Both the low- and high-manufacturing states are 2.4% below their pre-recession trend 10 quarters after the recession trough. But this may reflect that the aggregate pre-recession trend was too high for both groups. Blanchard, Cerutti, and Summers (2015) show that time-series estimates of output gaps are very sensitive to assumptions about the aggregate pre-recession trend. In contrast, our analysis is robust to mismeasurement of the common trend.

In the top right panel of Figure 3 we show that U.S. states loading more on the construction sector in the Great Recession experience both a deeper recession and a weaker recovery. Specifically, high-construction exposure predicts a 1.6% deeper recession and a 0.9% weaker recovery in the first ten quarters after the NBER trough. These patterns are very different from the pre-1990 recession, in which more exposed states were catching up after the recession. Another difference is the pre-trend: The more exposed states in the Great Recession grew faster than the less exposed states. These patterns inform our cross-sectional regression that the Great Recession was a boom-bust cycle.

The COVID-19 recession shows a similar pattern to the pre-1990 recessions: States with a deeper recession due to greater exposure to accommodation and food also experience stronger recoveries. The catch-up in levels is nearly complete: the recession is 2.8% deeper in the states with a high accommodation and food share, but the recovery is also 2.6% stronger.

The final subplot on the bottom right shows that for the 2001 recession the patterns are qualitatively similar to the Great Recession: a deeper recession due to a greater loading on IT and professional services implies a weaker recovery. More exposed states also grow more

[15]The relative recovery adjusting for the trend is $1.8\% + 15 \times 0.11\% \approx 3.4\%$, which is equal to the 3.4% relative recession depth.

quickly before the recession, again indicating boom-bust dynamics across states.

In summary, the data suggests that the pre-1990 and the Covid recessions have V-shaped RRIRF, in which a deeper recession predicts a stronger recovery. In contrast, exposure to the Great Recession and the 2001 recession predicts an inverse-V shaped RRIRF: faster growth prior to the recession, followed by a steeper decline and weaker recovery. We next make this finding explicit using regression analysis.

5 Results

5.1 First Stage We pool our RRIRF estimates across the two recession groups: those with weak aggregate recoveries (1990-1, 2001, and 2007-9) and those with strong aggregate recoveries (1948-9, 1953-4, 1957-8, 1960-1, 1969-70, 1973-5, 1980-2, 2020). Pooling helps to statistically distinguish these recession groups, is consistent with the similar qualitative recession pattern above, and not rejected by the overidentification tests we report below.

A potential concern is that we sort recessions using the outcome we seek to explain—the strength of the aggregate recovery. Three features of the reserach design limit the scope for this choice to bias our results. First, the sorting variable is an aggregate object whereas our estimates are cross-sectional. As we emphasize in the introduction, theories for slow recoveries differ in their cross-sectional predictions, so that pooling imposes no mechanical relationship. Second, in Section 5.4 we report RRIRF for individual recessions for which we have sufficient power, and these are similar to the pooled estimates. Third, in Section 5.4 we also report overidentification tests that do not reject common dynamics within groups at almost all horizons, but do reject common recession dynamics across the two groups at most horizons, so the grouping is informative about the data.

We start with the pooled first stage,

$$\begin{aligned} y_{i,P_j} - y_{i,T_j} = {} & \gamma_1 \tilde{\hat{y}}^b_{i,P_j,T_j} I(P_j \in \{1990, 2001, 2007\}) + \gamma_2 \tilde{\hat{y}}^b_{i,P_j,T_j} I(P_j \in \{\text{Other Recessions}\}) \\ & + \delta_1 \tilde{\tilde{y}}^b_{i,P_j,T_j} I(P_j \in \{1990, 2001, 2007\}) + \delta_2 \tilde{\tilde{y}}^b_{i,P_j,T_j} I(P_j \in \{\text{Other Recessions}\}) \\ & + \zeta_1 \Delta^L y_{i,P_j - H_{pre}} I(P_j \in \{1990, 2001, 2007\}) + \zeta_2 \Delta^L y_{i,P_j - H_{pre}} I(P_j \in \{\text{Other Recessions}\}) \\ & + \theta_j + \varepsilon_{i,j,h} \end{aligned} \tag{19}$$

in which we regress state output growth over the recession $y_{i,P_j} - y_{i,T_j}$ on Bartik-predicted recession depth $\tilde{\hat{y}}^b_{i,P_j,T_j} = \sum_{g=1}^{G} s_{i,g,P_j} \tilde{\hat{y}}_{-i,g,P_j,T_j}$. We allow the relationship to vary across recession groups. γ_1 is the elasticity of output growth w.r.t. predicted recession depth in the 1990-1, 2001, and 2007-9 recessions. γ_2 captures the average elasticity for the other recessions.

Table 1: First Stage of Local Recession Depth on Bartik Predicted Recession Depth

	Incl. Earn.	All Earn.	Unempl.	Employ.	GDP
Coeff. on Predicted Recession Depth:					
90-1, 01, 07-9	2.795***	2.242***	0.585**	1.993***	4.068***
	(0.750)	(0.656)	(0.274)	(0.487)	(0.956)
Pre-1990 & 2020	1.487***	0.866***	0.651***	0.985***	1.078***
	(0.334)	(0.222)	(0.096)	(0.215)	(0.217)
Bartik Trend Control	Yes	Yes	Yes	Yes	Yes
Local Pre-Trend Control	Yes	Yes	Yes	Yes	Yes
Recession FE	Yes	Yes	Yes	Yes	Yes
F-stat (90-1, 01, 07-9)	13.9	11.7	4.5	16.7	18.1
F-stat (pre-1990/2020)	19.8	15.2	46.2	21.0	24.6
State Clusters	100	100	100	100	100
R^2	0.136	0.056	0.137	0.150	0.408
Observations	544	543	490	535	100

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Standard errors clustered by state.

We control for the annualized Bartik-predicted trend, $\tilde{\tilde{y}}^{b}_{i,P_j,T_j} = \frac{4}{T_j - P_j} \sum_{g=1}^{G} s_{i,g,P_j} \tilde{\tilde{y}}_{-i,g,P_j,T_j}$, separately for each recession group. This controls for differences in trends across more and less cyclically sensitive industries. We annualize this trend growth so we can consistently use the same control in the second stage. We also control for a state's average growth rate from 15 years before the recession to 5 years before the recession $\Delta^{L} y_{i,P_j - H_{pre}}$, again interacted with recession group. This control addresses any State-specific trends correlated with recession exposure but uncorrelated with industry trends (e.g., Fukui, Nakamura, and Steinsson, 2023). Figure 3 shows that differentially exposed states are on different pre-trends, which indicates either differences in underlying trend growth or a boom-bust cycle. We start the trend control five years prior to the recession so as to minimize capturing boom-bust dynamics, and we show that our results robust to extending this window. We do not need to restrict the timing of national industry trend growth in $\tilde{\tilde{y}}_{-i,g,P_j,T_j}$, which is constructed with data right up to the national recession peak. We include recession fixed effects θ_j so that our estimates exploit only within-recession variation. Within each Recession Group we cluster by state.[16]

[16]Double-clustering by both state and recession date and adjusting for small numbers of clusters following Imbens and Kolesar (2016) generally produces smaller standard errors than clustering by state alone. See Figure 16. This suggests that our Bartik approach largely captures cross-state correlations, consistent with Adao, Kolesár, and Morales (2019).

The pooled first stage estimates are shown in Table 1. In column (1), the outcome variable is local income from the sectors included in the instrument. Column (2) includes all state income. Both columns show that Bartik predicted recession depth strongly predicts State income growth over the national recession. For the 1990-1, 2001, and 2007-9 recessions the first stage coefficient for income included in the instrument is $\gamma_1 = 2.8$ (t-statistic 3.7). This means that a 1% deeper predicted recession translates into a 2.8% deeper local recession. For the other recessions the first stage coefficient is $\gamma_2 = 1.5$ (t-statistic 4.5). In columns (3) through (5) we measure State recession depth relative to trend with the unemployment rate, employment, and GDP. Our Bartik instrument is a strong predictor of the State recession depth in each case, except for unemployment in the 1990-1, 2001, and 2007-9 recessions.

5.2 Relative Recession Impulse Response Functions We next estimate RRIRFs across recession groups. Within each group, these capture the differential movement of income, employment, and output across more and less exposed states both before and after the national recession.

Specifically, we pool equation the second stage across recession groups,

$$\begin{aligned} y_{i,T_j+h} - y_{i,P_j} = {} & \beta_{1h}\widehat{(y_{i,P_j} - y_{i,T_j})} I(P_j \in \{1990, 2001, 2007\}) \\ & + \beta_{2h}\widehat{(y_{i,P_j} - y_{i,T_j})} I(P_j \in \{\text{Other Recessions}\}) \\ & + \delta_{h1}\tilde{\tilde{y}}^b_{i,P_j,T_j} I(P_j \in \{1990, 2001, 2007\}) + \delta_{h2}\tilde{\tilde{y}}^b_{i,P_j,T_j} I(P_j \in \{\text{Other Recessions}\}) \\ & + \zeta_{h1}\Delta^L y_{i,P_j-H_{pre}} I(P_j \in \{1990, 2001, 2007\}) + \zeta_{h2}\Delta^L y_{i,P_j-H_{pre}} I(P_j \in \{\text{Other Recessions}\}) \\ & + \theta_{j,h} + \varepsilon_{i,j,h} \end{aligned} \tag{20}$$

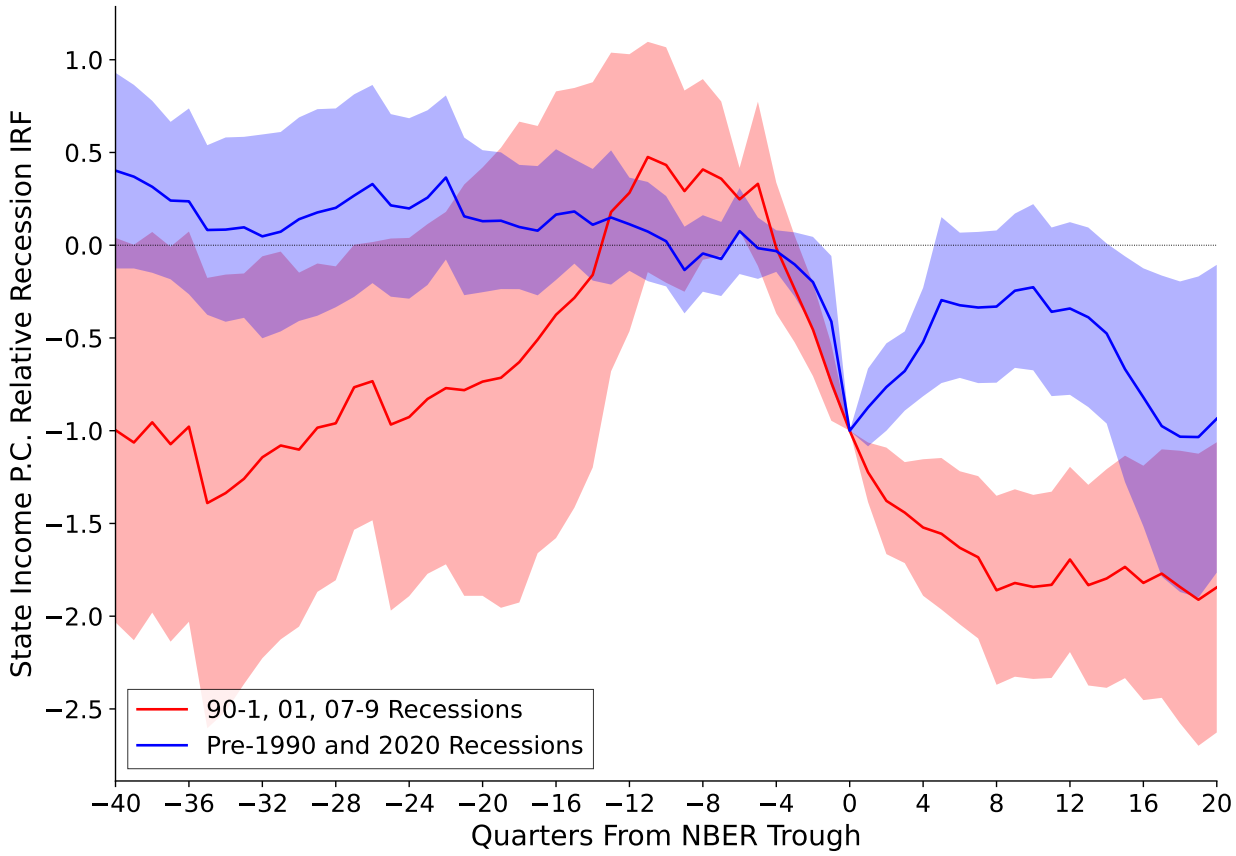
where $y_{i,T_j+h} - y_{i,P_j}$ is h-quarter growth in the outcome variable since the recession peak P_j. $\widehat{(y_{i,P_j} - y_{i,T_j})}$ is the instrumented growth in the outcome variable over the recession from the first stage (19).

Figure 4 plots our estimates for the relative recession impulse response functions, β_{1h} and β_{2h} for $h = -40$ quarters before the recession trough to $h = 20$ quarters after the recession trough. For $h = 0$, the estimate is equal to -1 for both recession groups by construction.

Other than the immediate period prior to the recession trough at $h = 0$, the two RRIRF are different. For the 1990-1, 2001, and 2007-9 recessions, a deeper relative recession predicts a weaker relative recovery.[17] This is reflected in the decreasing estimates of the relative recession effect for $h > 0$. Thus, states that are more exposed to the recession experience

[17] Within this group the Bartik estimator loads almost entirely on the 2001 ($\alpha = 0.32$) and 2007-9 ($\alpha = 0.68$) recessions. Local industry shares have little predictive power for local recession depth in 1990-1, which consequently receives a Rotemberg weight of $\alpha = 0$ (Section 5.4 and Appendix Figure 20c). The pooled estimates for this group are therefore identified by the 2001 and 2007-9 episodes.

Figure 4: Relative Recession IRF of State Personal Income P.C. by Recession Group



relatively weaker growth in the recovery compared to less exposed states.

From $h = -20$ to $h = -10$, the more exposed states grow relatively more quickly compared to the less exposed states in the 1990-1, 2001, and 2007-9 recessions. This is captured by the increase in β_{1h} over this window from roughly -1 to slightly above 0. The strong relative growth in more exposed states is a break from their prior behavior. Going from $h = -20$ back as far as $h = -40$, the more exposed states are growing at a similar pace as the less exposed states. The point estimates for the recovery period are somewhat below the pre-recession trend, which is indicative of modest hysteresis. However, the RRIRF is also consistent with no hysteresis effects in the 1990-1, 2001, and 2007-9 recessions: the test that the average pre-boom value from $h = -40$ to $h = -28$ is the same as the post-recovery value from $h = 8$ to $h = 12$ is not rejected with a p-value of 0.13.

The RRIRF for the pre-1990 and 2020 recessions is very different. Prior to the recession, the more exposed states are on a slight downward trend relative to the less exposed states, but the point estimates are not statistically different from zero. During the recession, the more exposed states experience a sharp relative decline, followed by an almost equally sharp relative recovery. Four quarters after the recession trough, we can no longer reject the null that the more exposed states have fully caught up to the less exposed states.[18] The point estimates for the recovery are again consistent with modest hysteresis. The second decline at $h > 12$ quarters reflects a subsequent recession: the peak of the 1973-5 recession follows twelve quarters after the trough of the 1969-70 recession.

Since the RRIRF plausibly aggregates to the ARIRF, our estimates for the 1990-1, 2001, and 2007-9 recessions are most consistent with a boom-bust cycle (Beaudry, Galizia, and Portier, 2018; Jordà, Schularick, and Taylor, 2013; Rognlie, Shleifer, and Simsek, 2018), in which more exposed states experience a stronger boom followed by a stronger bust. Recoveries from these recessions look weak based on their pre-recession trends, because the pre-trends are inflated by the boom.

The RRIRF for the pre-1990 and 2020 recessions is consistent with plucking models of the business cycle (Friedman, 1993; Dupraz, Nakamura, and Steinsson, 2025). The more exposed states are plucked down by a shock during the recession, but then return to their prior trend in the recovery. While we find some weak evidence for hysteresis effects, their magnitude is very similar across the two recession groups, which suggests that the difference in recovery strength is not driven by a hysteresis mechanism.[19] Instead our evidence suggests that the

[18] The equivalent test to the one we run on 1990-1, 2001, and 2007-9 recessions also does not reject that the pre-recession average from h=-40 to h=-28 is the same as the post-recovery average from h=8 to h=12 (p-value 0.20).

[19] We can calculate the relative hysteresis effects by differencing over the h=-40 to h=-28 average relative to the h=8 to h=12 average across the two recession types. The test that this difference is zero has a p-value

weak recoveries in the 1990-1, 2001, and 2007-9 recessions relative to the other recessions are due to different types of shocks.

5.3 Other Outcomes Our baseline outcome variable is state income, specifically the components of state income that were included in our Bartik instrument. We next show impulse response functions for four alternative outcomes: all state income (including all sectors, transfers, dividends, and rent), the state unemployment rate, state employment, and for the last two recessions, state GDP. The four impulse response functions are plotted in Figure 5. The first stage for these other outcomes is shown in Table 1.

The boom-bust pattern for all state income in the 1990-1, 2001, and 2007-9 recessions is essentially the same as in our baseline estimates, though less precisely estimated (Figure 5a). For the pre-1990 and 2020 recessions, the recovery is more complete based on all income and suggests a complete absence of hysteresis effects.

Because all state income includes transfers, whereas our baseline outcome measure in Figure 4 did not, the comparison of these Figures speaks to whether differences in countercyclical policy may explain differences in recovery strength. The fact that the estimates are very similar suggests that this is an unlikely explanation. We confirm this insight in Appendix Figure 18, in which we directly estimate the effect of predicted recession depth on the transfer share in total income by recession group. We find that the 1990-1, 2001, and 2007-9 recessions are followed by a persistently elevated transfer share, whereas it normalizes quickly following the pre-1990 and 2020 recessions. This shows that transfers are leaning against the differential recovery patterns we find rather than causing them.

The unemployment impulse response function in Figure 5b is noticeably different from the other outcomes. It displays convergence in unemployment rates across states for all recessions, consistent with prior work by Blanchard and Katz (1992) and Yagan (2019). But the speed is very different: For the 1990-1, 2001, and 2007-9 recessions convergence takes four years, whereas for the pre-1990 and 2020 recessions convergence is almost complete after one year.

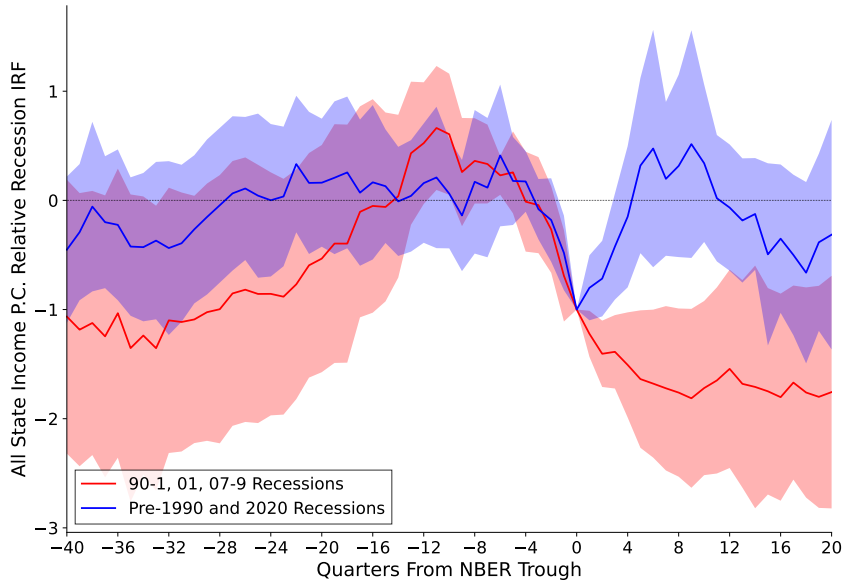
There are no noticeable pre-trends in unemployment for the pre-1990 and 2020 recessions. In contrast, unemployment in more exposed states was relatively low at the peak of the 1990-1, 2001, and 2007-9 recessions, consistent with a relative boom in the labor market. However, this boom is relatively small compared to the other outcomes. The relative boom in state income in Figure 4 in the five years prior to the recession is almost as large as the relative bust. The relative decline in unemployment prior to the recession is only one-quarter as large as the subsequent increase in unemployment. This suggests that boom-bust cycles may have

of 0.72.

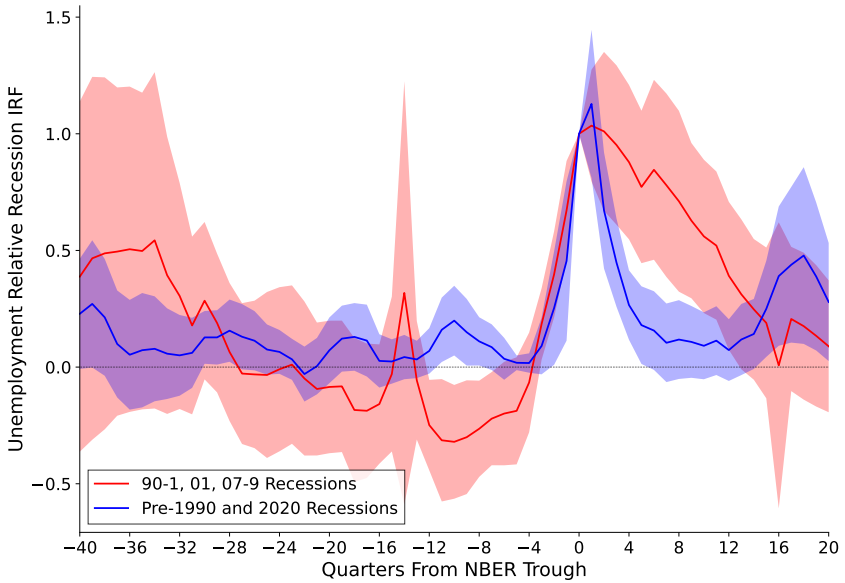
asymmetric effects on labor markets and that unemployment rates are less informative about the boom-bust nature of a recession than other outcomes.

Figure 5: Relative Recession IRF for Other Outcomes

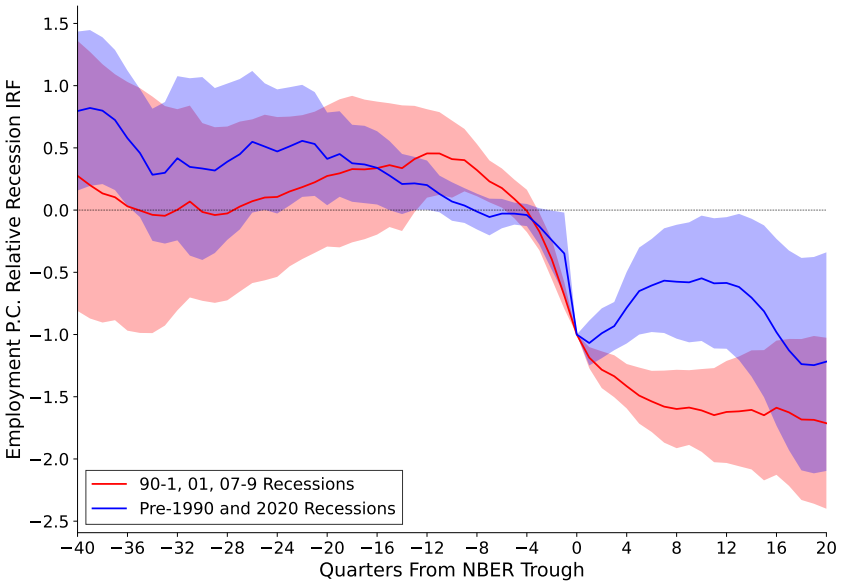
(a) All State Income P.C.



(b) Unemployment Rate



(c) Employment P.C.



(d) GDP P.C.

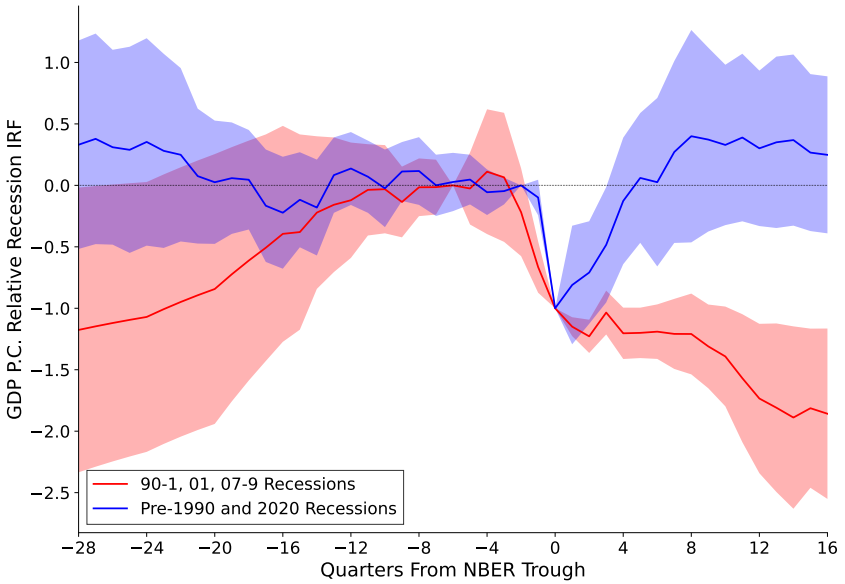


Figure 5c plots the relative recession impulse response function for employment to population ratio. The patterns are qualitatively similar to the income results: a sharp recovery in the pre-1990 and 2020 recessions, and a boom-bust cycle in the 1990-1, 2001, and 2007-9 recessions. But quantitatively there are two notable differences. First, there is weaker evidence that the more and less exposed states are on parallel trends. In particular, for the pre-1990 and 2020 recessions, the more exposed states are on a relative downward trend prior to the recession. For the 1990-1, 2001, and 2007-9 recessions, the pre-trend appears to be first negative before turning positive in the five years prior to the recession.

Second, there is stronger evidence for hysteresis effects in employment than in income. These hysteresis effects connect to the literature on jobless recoveries (Bernanke, 2003; Groshen and Potter, 2005; Bachmann, 2012; Jaimovich and Siu, 2020), which documents that employment recovered more slowly than output after the 1990-1, 2001, and 2007-9 recessions. We also find similar hysteresis effects for employment in the other recessions, which suggests that the mechanisms behind weak employment recoveries may not be unique to the 1990-1, 2001, and 2007-9 recessions. However, the failure of parallel trends does suggest this interpretation should be taken with some caution. The RRIRF for the total wage bill in Appendix Figure 19 does not show strong hysteresis effects, which suggests that the hysteresis effects for employment may be spurious or offset by per capita wages. Regardless of the interpretation, the hysteresis effects are not different across recession groups, and thus do not explain the differential recovery strength.

The GDP data is more limited since it is only available for the last two business cycles. Yet the patterns in Figure 5d are similar to the broader sample for state income. The Covid recession displays a fast recovery to the pre-recession trend. The Great Recession shows a boom-bust cycle. These outcomes corroborate our findings using state income.

For comparison, we plot the relative growth rates of population for the two recession groups in Appendix Figure 17. Population grows relatively more in the more exposed states in the 1990-1, 2001, and 2007-9 recessions, but relatively less in the other recessions. Population growth remains smooth throughout the recession, unlike the economic outcomes in Figures 4 and 5. Our results extend the findings of (Autor, Dorn, and Hanson, 2021, p. 428) for the Great Recession and support their conclusion that the "estimates . . . for population belie the expectation of a lasting downturn in the CZs with greater exposure to the Great Recession."

5.4 Decomposition and Overidentification of the RRIRF We decompose our RRIRF estimates to (1) show what variation is most important, (2) determine whether the changes across recessions come from different behavior within or across industries (or both), and (3)

test overidentification restrictions implicit in the Bartik instrument. Goldsmith-Pinkham, Sorkin, and Swift (2020) show that the Bartik estimator is equivalent to a weighted GMM estimator, in which the local industry shares s_{i,k,P_j} are the instruments and industry growth rates $\tilde{\hat{y}}_{-i,g,P_j,T_j}$ are the weights. When the industry growth rates are common across locations, then we can decompose the Bartik estimate for β_h in (20) as

$$\beta_h = \sum_{k,j} \alpha_{k,P_j} \beta_{h,k,P_j}$$

where β_{h,k,P_j} is the estimate for β_h from instrumenting (20) using the industry share s_{i,k,P_j} and α_{k,P_j} is its associated "Rotemberg weight" in the Bartik instrument. In our setting, this decomposition holds only asymptotically since we use leave-one-out growth rates to construct the Bartik instrument. In practice, the estimates from the leave-one-out Bartik and the common growth rate Bartik are very close, so the decomposition is informative about what sources of variation matter for our estimates.

We focus on the estimates for state income from Figure 4. In Figures 6a and 6b we show how these impulse response functions are informed by exposure to five broad sectoral aggregates that we can consistently define throughout the sample: Construction, manufacturing, transportation and utilities, wholesale and retail trade, and other services. Each estimated impulse response function uses the corresponding local income share of that sector as instrument for local recession depth. When a sector share is largely uninformative about local recession depth, then its corresponding Rotemberg weight α_{k,P_j} will be small and the standard error for β_{h,k,P_j} will be large. We therefore exclude sectors with a weight less than 0.05 from the figure. Appendix Figure 20 includes all sectors.

The three major sectors that inform recession depth in the 1990-1, 2001, and 2007-9 recession are construction ($\alpha = 0.3$), transportation and utilities ($\alpha = 0.08$), and other services ($\alpha = 0.54$). Figure 6a shows that recession exposure based on each of these major sectors predicts a boom-bust cycle. In almost all cases, the pooled estimate is within the 95% confidence band of each of the sectoral exposures. Figure 7a plots p-values for overidentification tests implicit in the Bartik instrument. Across major sectors—the green dashed-dotted line—we do not reject overidentification at almost any horizon, and in particular not the boom and bust period respectively.

Thus, both economically and statistically, the boom-bust nature of these recessions is broad-based rather than confined to a particular industry. It also supports the identification assumption that the boom-bust cycle is unlikely to be due to a correlation of sectoral exposure and local income shocks, as one then would expect the boom-bust cycle to be confined to sectors that are overrepresented in locations subject to those local shocks.

Leamer (2021) argues that weak recoveries in the manufacturing sector help explain weak aggregate recoveries in the 1990-1, 2001, and 2007-9 recessions based on cross-State data. While we also find that manufacturing exposure predicts weaker recoveries in the 1990-1, 2001, and 2007-9 recessions, exposure based on the other sectors yields a similar pattern. Furthermore, these other sectors matter more both in the cross-section and in the aggregate, based on their Rotemberg weights and aggregate sectoral shares. This suggests that an explanation for weaker recoveries based solely on the changing behavior of manufacturing is likely incomplete.

For the national recessions before 1990 and in 2020, exposure is primarily determined by the local income shares in manufacturing ($\alpha = 0.52$), wholesale and retail trade ($\alpha = 0.08$), and other services ($\alpha = 0.37$). Figure 6b shows that each of the three sectoral exposures predicts a rapid recovery from the NBER trough in these recessions. We again pass overidentification tests based on industry groups pre- and post-recession, which is shown by the green dashed-dotted line in Figure 7b.

There is no evidence of boom-bust cycles based on any of the main industry exposures. If anything, the more exposed states based on manufacturing and wholesale and retail trade are on a downward-trend relative to less exposed states. The recovery from the recession is consistent with a return to this downward-trend. Service sector exposure predicts no differential trend prior to the recession and a complete recovery thereafter.

Comparing Figures 6a and 6b, we find that all major sector exposures predict weak recoveries in the 1990-1, 2001, and 2007-9 recessions and strong recoveries in the other recessions. This implies that the weaker recovery dynamics are coming primarily from within-sector changes rather than changes in sectoral weights.

Figures 6c and 6d show the impulse response functions when the variation is restricted to particular recessions. For both the 2001 and 2007-9 recession the boom-bust pattern is evident and each contains the pooled estimate within their 95% confidence bands. For both the boom and the bust period, we pass the overidentification tests plotted by the orange line in Figure 7a.

The common recession patterns occur despite these recessions loading on different sectors: Cross-State variation in recession depth during the Great Recession is best explained using exposure to the construction and health sectors. For the 2001 recession, exposure to health and information and professional services matters most. The figure omits the 1990-1 recession in which local industry shares have little predictive power for local recession depth, which is reflected in both a low Rotemberg weight $\alpha = 0$ and wide confidence band (Appendix Figure 20c).

There is no evidence of a boom-bust cycle in the pre-1990 or the Covid recessions in

Figure 6d. Instead, both these recessions show a rapid recovery in the quarters after the NBER trough. States more exposed to the 2020 recession show a complete relative recovery within a year to their pre-recession relative trend. This pattern is primarily informed by exposure to the accommodation and food and the entertainment sectors: States like Nevada and Hawaii contracted sharply during Covid but then rapidly recovered as the economy re-opened. The more exposed states to pre-1990 recessions also grow relatively quickly after the recession trough. In this sense, the pre-1990 recessions look similar to the Covid recession and different from the 1990-1, 2001, and 2007-9 recessions. Statistically we also do not reject overidentification tests for this group of recessions, shown by the orange line in Figure 7b. Nevertheless, states exposed to the pre-1990 recessions and Covid recession differ somewhat in their pre-trends and thus in the extent to which their relative recovery looks complete in levels after the recession. For the pre-1990 recessions, the more exposed states are on a slight downward trend, whereas in the Covid recession the more exposed states are on a slight upward trend prior to the recession, as the trend controls in (20) only absorbs the average trend across recessions.

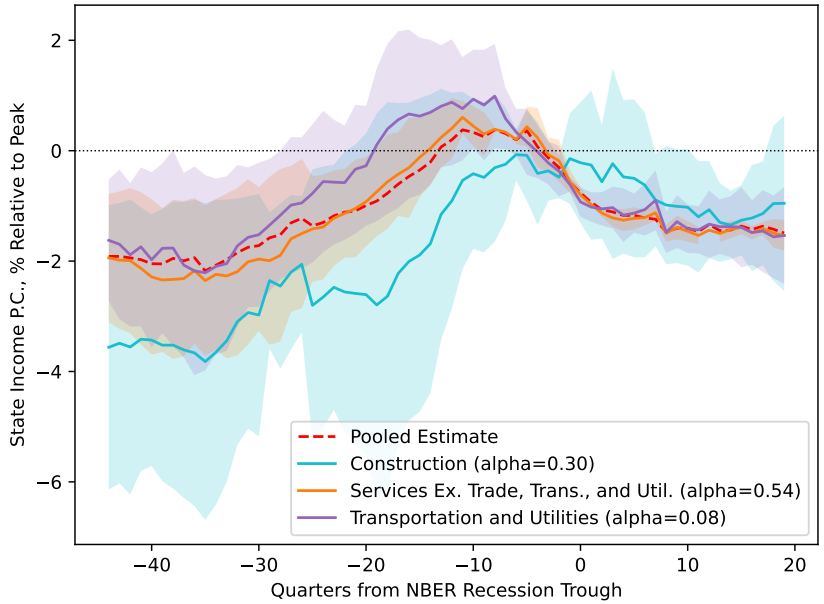
Because idiosyncratic industry trends remain after orthogonalizing the Bartik with trend controls, the overidentification tests pass more likely with aggregated industry-recession variation compared to using the individual instruments (Figure 7). The aggregated industry and recession groups include all of the top-5 instruments, but smooth over idiosyncratic trends. Note that our identification assumptions only require the Bartik instrument to be valid conditional on controls, and not each individual industry-recession share, on which the broad industry and recession groups are more informative.

That we do not reject overidentification, particularly after aggregating to broad groups, is not due to a lack of power. In Appendix Figure 22 we plot the p-values for the overidentification tests when we pool the RRIRF estimates across all recessions.[20] Even though we conservatively allow the trend controls to still vary over the original recession groups and the composition of instruments is biased towards the larger subsample of "other recessions", we reject overidentification at almost all horizons. In particular, the data strongly reject the view that recession dynamics are identical prior to the recession and during the recovery phase.

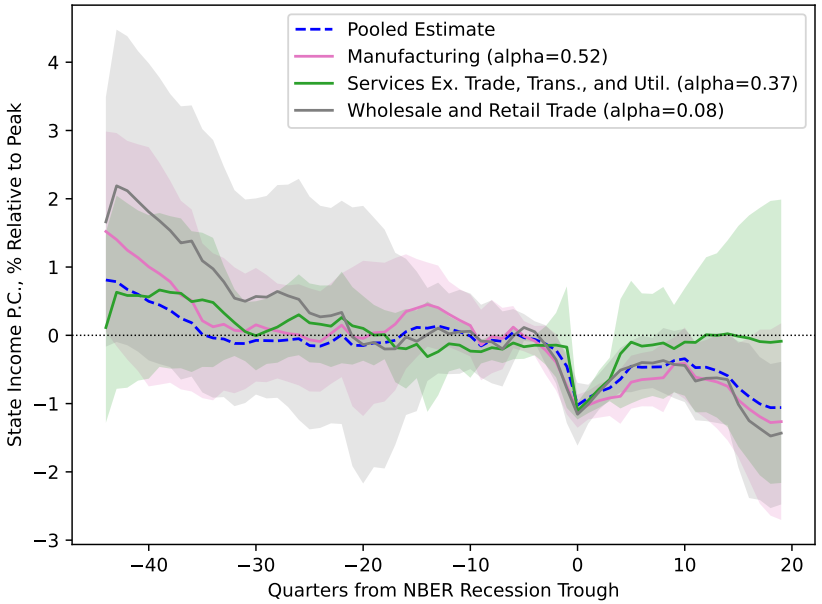
[20]Because the Rotemberg weights add to one in the pooled sample and in each of the two subsamples, we set the threshold to $\alpha = 0.025$ in the pooled sample so as to include a comparable set of instruments.

Figure 6: Impulse Response Functions of State Income P.C. by Industry and Recession Variation

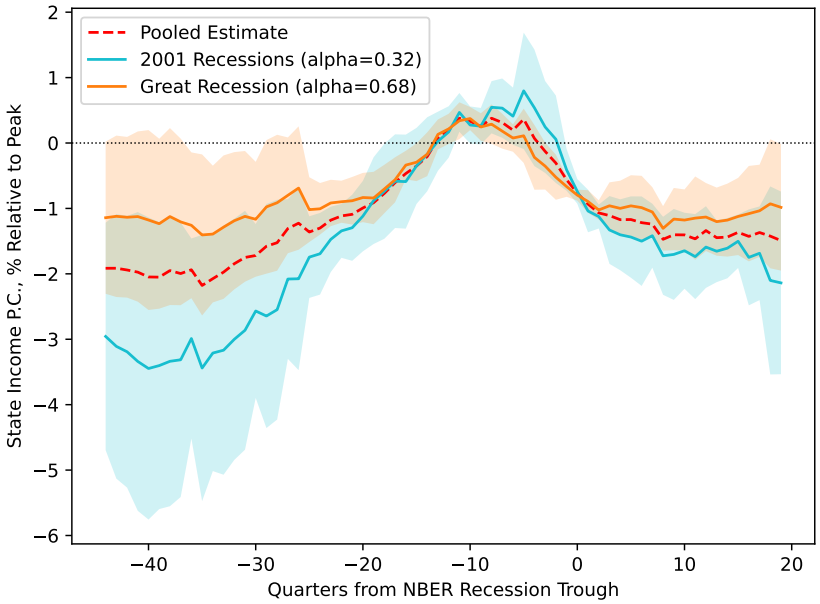
(a) By Industry: 1990-1, 2001, 2007-9



(b) By Industry: Other Recessions



(c) By Recessions: 1990-1, 2001, 2007-9



(d) By Recessions: Other Recessions

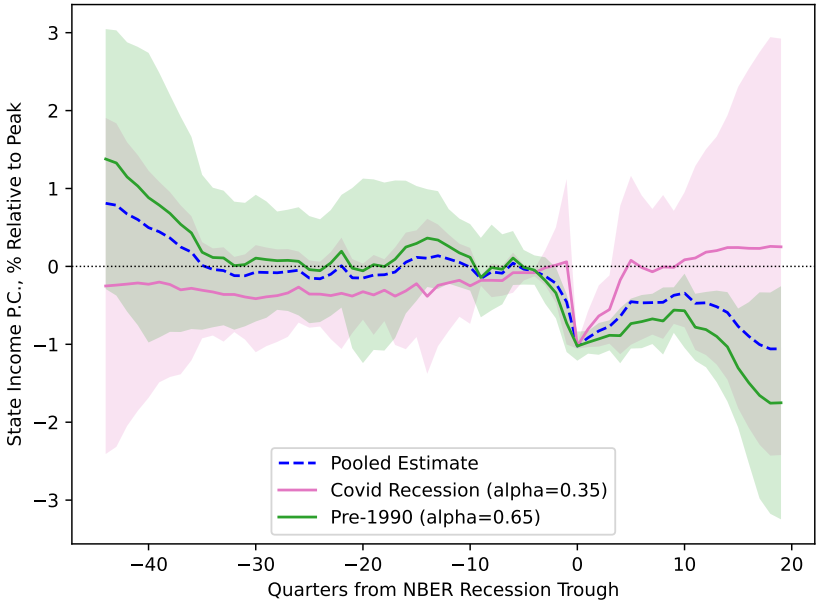
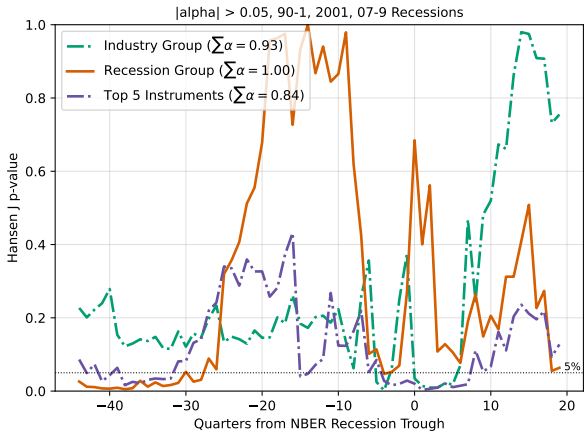
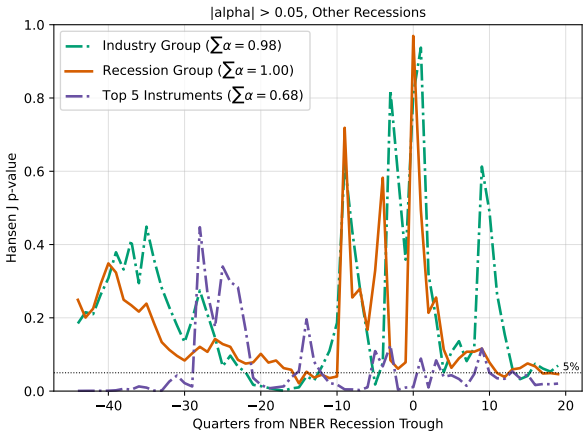


Figure 7: P-Values for Overidentification Tests

(a) 1990-1, 2001, 2007-9 Recessions



(b) Other Recessions



In short, weaker recoveries following a boom-bust cycle are a broad-based phenomenon across sectors and recessions for the 1990-1, 2001, and 2007-9 recessions. In contrast, the strong recoveries from the pre-1990 and Covid recessions are also broad-based, and they do not display any boom-bust dynamics. This evidence is most consistent with the view that the recovery dynamics are different across these recessions because they are caused by different shocks.

5.5 Robustness We conclude our empirical analysis by showing that our main results are robust to a variety of changes in the construction of the instrument, the set of fixed effects, and the sample. We present these robustness checks in Figures 8a to 8f. Each figure contains six panels. The left column shows the IRFs for the 1990-1, 2001, and 2007-9 recessions and the right column for the other recessions. The top panel shows how the IRFs change when we vary the instrument (Figures 8a and 8b), the middle panel varies the fixed effects (Figures 8c and 8d), and the bottom panel changes the sample (Figures 8e and 8f). In each case, we re-estimate the first stage and second stage for each variation.

Figures 8a and 8b show how the IRFs for the 1990-1, 2001, 2007-9 recessions and the other recessions change when we construct the instrument differently. In (2) we construct the peak-to-peak national industry growth rates two quarters prior to the NBER peak and in (3) we construct it using five-year pretrend from the NBER peak. In (4) we omit detrending industry growth in the Bartik instrument entirely. Neither of these changes materially affects the impulse response functions. In (5) we push back the local trend estimate to end seven years prior to the recession peak, which results in minimal changes. And in (6) we include the government and government and finance sectors in local income shares. This means

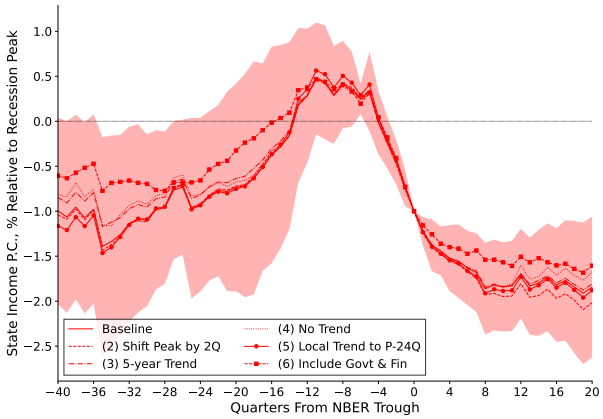
that as well as comparing across private industries we also compare more cyclical private industries with the less cyclical government sector. The main change is that the boom-bust pattern for the 1990-1, 2001, 2007-9 recessions starts earlier.

In Figures 8c and 8d we vary the set of fixed effects to proxy for local shocks correlated with national recession exposure. In (2) we include State fixed effects which absorb any permanent difference across states such as permanently different trends correlated with recession exposure. We find that State fixed effects primarily affect the IRF for the fast recoveries by flattening the pretrend and making the recovery look more complete. In (3) and (4) we use Census region and Census division fixed effects interacted with recession fixed effects. These fixed effects capture local shocks that are regional in nature even if they vary across recessions. They primarily change the IRF for the 1990-1, 2001, 2007-9 recessions by increasing the size of both the boom and the bust.

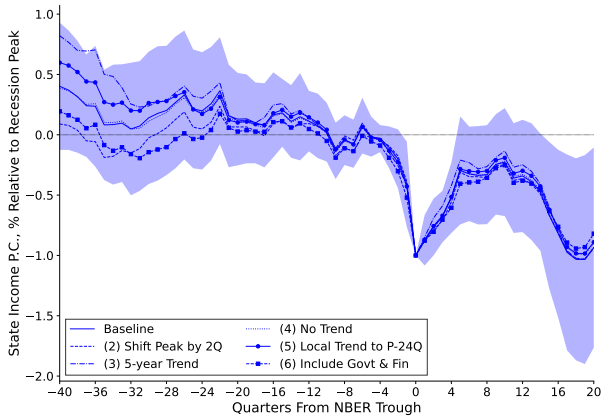
In Figures 8e and 8f we show the sensitivity to the sample. In (2) we drop the three most mining-intensive states and in (3) the three most finance-intensive states, the two categories in which we are least confident that the location of income corresponds to production. The differences from our baseline are minimal. (4) drops the Covid recession, which makes the downward-trend in the other recessions more pronounced. (5) and (6) show only the Great Recession and the 2001 recession, respectively. Both display the boom-bust pattern of our baseline.

Figure 8: RRIRF Robustness

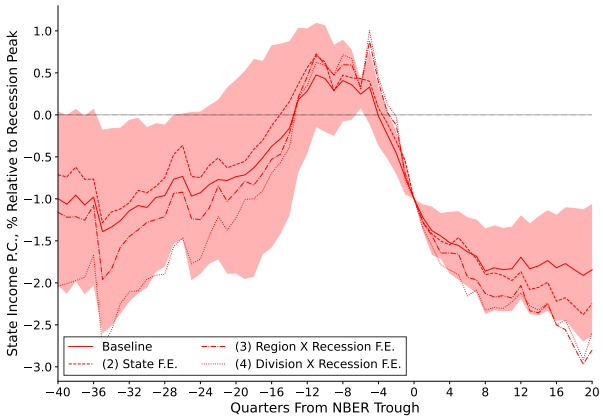
(a) Instrument Variation: 90-1, 01, 07-9



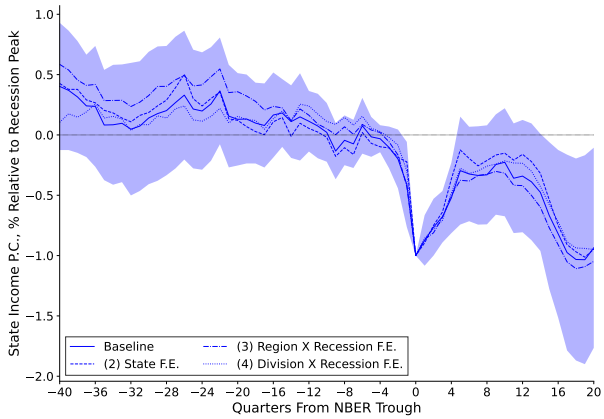
(b) Instrument Variation: Other Recessions



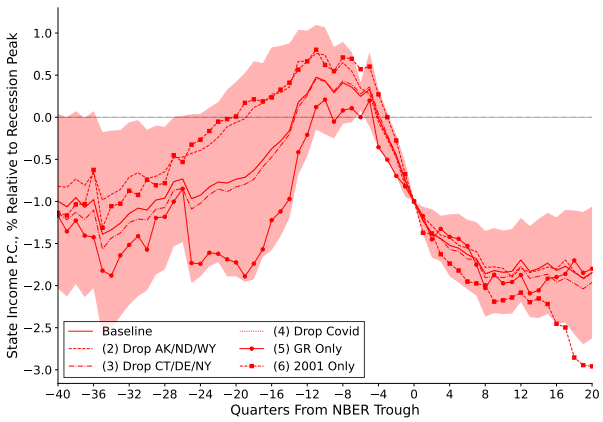
(c) Fixed Effects: 90-1, 01, 07-9



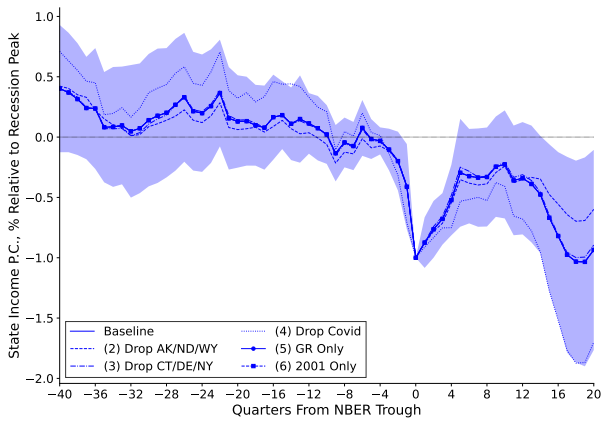
(d) Fixed Effects: Other Recessions



(e) Sample Variation: 90-1, 01, 07-9



(f) Sample Variation: Other Recessions



6 Implications for Countercyclical Policy Version

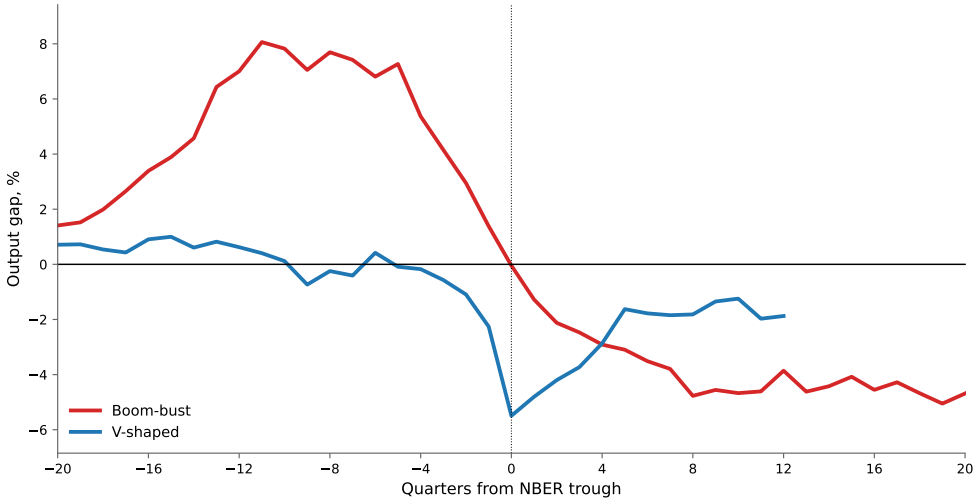
We next show that distinguishing between the two types of recessions we identify matters for policy. Specifically, we compare the real interest rate paths that close the output gap of each recession type. We find that in a V-shaped recession, the long-run real rate needs to be substantially lower primarily at the trough but then normalizes quickly. In contrast, a boom-bust cycle of the same depth calls for substantially higher long-run real rates during the boom and substantially lower long-run real interest well into the recovery phase. Thus, monetary policy response depends importantly on a correct diagnosis of the recession.

We construct these policy counterfactuals following the semi-structural approach of Barnichon and Mesters (2023), McKay and Wolf (2023), and Caravello, McKay, and Wolf (2024). Two ingredients are necessary: first, we must map our RRIRF estimates into a path for the output gap for each recession type. Second, we determine what counterfactual real rate path zeros out this path using an estimated monetary policy causal effects matrix, which we take from Caravello, McKay, and Wolf (2024). McKay and Wolf (2023) and Caravello, McKay, and Wolf (2024) argue this procedure yields valid counterfactuals in a wide class of standard models. While we focus on monetary policy due to readily available monetary policy causal effects matrices from this literature, the approach in principle extends to any policy instrument. An advantage of the semi-structural approach is that, like in our empirical work, we do not need to take a stance on what causes the different types of recessions. We also believe that this approach makes our assumptions transparent, which may be less clear if we imposed more structure.

The structure we need is how the RRIRF map to the output gap. As we discuss in Section 2.3 and our aggregation results in Appendix D, the RRIRF plausibly maps directly into the ARIRF for empirically reasonable cross-state convergence rates. However, the ARIRF is not an output gap path for two reasons: first, it is scaled and normalized; and second, it is expressed in terms of output not the output gap. Our first assumption is that both recession types do not affect potential output, so that the ARIRF for output is the same as that for the output gap. As we showed in Section 5, we can statistically reject hysteresis for each recession type, and to the extent that it is present there is little evidence that it is different across recessions. Thus, we do not view this assumption as first order important when comparing real rate paths across recessions. Perhaps more important is that some theories of boom-bust cycles, such as those that involve an investment boom, imply that potential output is elevated throughout the boom-bust cycle due to an excess capital stock. This would mean we overstate the degree of tightening needed in the boom and understate the degree of accommodation needed in the bust.

We scale the ARIRF to a 5.6% decline based on the average peak-to-trough drop in private domestic demand relative to the pre-recession trend.[21] The thought experiment thus of a policy maker that sees two recessions of the same depth (by this standard metric) but of different type. We assume that the policy maker recognizes that the boom-bust depth is overstated because the pre-trend is inflated by the boom and corrects for this upward bias.[22] Finally, we normalize the ARIRF for the V-shape recession to a zero output gap at the peak, and the ARIRF for the boom-bust cycle to a zero output gap 5-10 years prior to the trough. Figure 9 plots the implied output gap by recession type. Note that we terminate the output gap path for V-shaped recessions at $h = 12$ quarters so as not capture a second recession.

Figure 9: Implied Output Gap Paths from Estimated RRIRFs



Denote these output gap paths by $\mathcal{Y}^{\text{V-shape}}$ and $\mathcal{Y}^{\text{boom-bust}}$. From Caravello, McKay, and Wolf (2024) we obtain monetary policy causal effects matrices for the output gap $\mathcal{M}^{\mathcal{Y}}$ and the long-run real rate $\mathcal{M}^{\mathcal{R}}$. We also impose a quadratic penalty $\lambda||dR||^2$ on changes in the real rate to keep the problem well behaved.[23] Starting at any horizon H, the path for the

[21] Based on this metric, the two recession groups have very similar depths: 5.67% in V-shape and 5.42% in boom-bust cycles.

[22] The correct output gap at the peak is $B_j = (\beta_{P_j} - \beta_{-40:-20})d_j$ where d_j is the true depth, β_{P_j} the estimate of the RRIRF at the peak (which we normalize to 0), and $\beta_{-40:-20}$ the estimate of the RRIRF before the boom. The trend is inflated by $\frac{B_j - B_{j-1}}{T_j}$, where T_j the number of quarters over which the trend is constructed and B_{j-1} is the output gap peak in the previous recession. The true depth can be solved for as $d_j = \frac{d_j^{\text{raw}} + \frac{L_j}{T_j} B_{j-1}}{1 + \frac{L_j}{T_j}\beta_{-40:-20}}$, where L_j are peak-to-trough quarters.

[23] For our benchmark value of $\lambda = 0.5$ (following Caravello, McKay, and Wolf (2024)), the solution closes 99% of squared output gaps.

implied long-run real rate that closes the output gap going forward is,

$$\mathcal{R}^* = \mathcal{R} - \left[\mathcal{M}^{\mathcal{Y}'}\mathcal{M}^{\mathcal{Y}} + \lambda \mathcal{M}^{\mathcal{R}'}\mathcal{M}^{\mathcal{R}}\right]^{-1}\mathcal{M}^{\mathcal{Y}'}\mathcal{Y}_{h\geq H} \tag{21}$$

We call this counterfactual path "R-star."[24]

We need to make two additional assumptions to implement this approach. First, because the $\mathcal{M}$ matrices have a 50-year dimension, we need to make assumptions about how the output gap behaves beyond the terminal date in Figure 9. We assume that the output gap decays at rate $1-\rho$, and choose $\rho = 0.9$ as baseline. As we show in Appendix Figure 23, our exact choice matters relatively little for our headline conclusions. Second, we need a baseline path for the long-run real rate. We focus on a ten year window, and use the ten-year nominal treasury yield for the nominal rate. For inflation expectations we use the Federal Reserve Board ten-year inflation expectation from 1968 onward, and before that the five-year trailing average of the Livingston median one-year inflation expectations.[25] We detrend this real rate using peak-to-peak changes prior to each recession and normalize it to zero at each recession peak.

Figure 10 plots for each recession type the path of the realized ten-year real interest rate (dashed lines) and the counterfactual ten-year real interest rate that would close the output gap at any point going forward. In V-shaped recessions, the largest discrepancy is right at the trough, where our counterfactuals call for about one percentage point lower ten-year real rates. This difference shrinks quickly after the trough as the output gap closes so there is less need for additional accommodation.

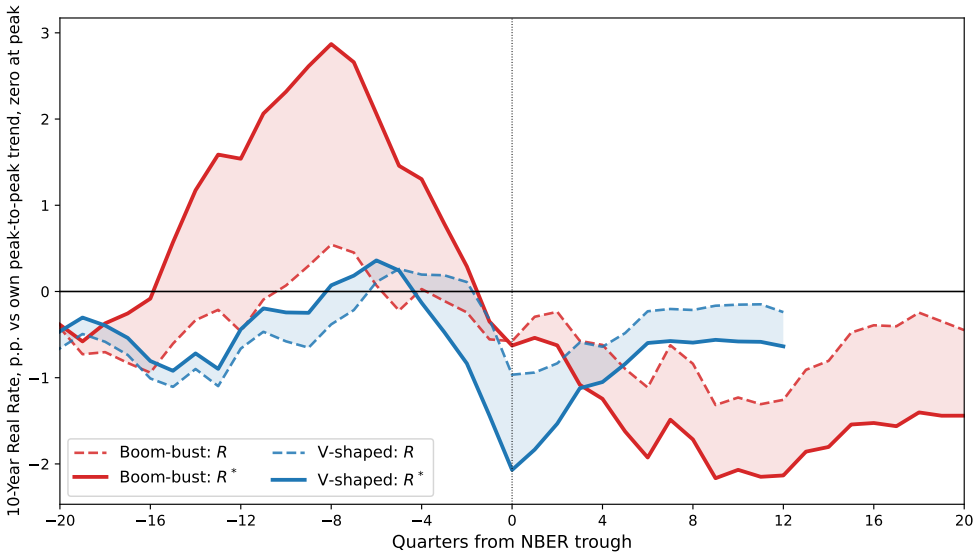
In boom-bust cycles, the ten-year R-star is two percentage points above the realized real rate during the boom phase, and one percentage point below during the recovery phase. This reflects the change from positive to negative output gaps in Figure 9. Notably, despite the same recession depth, the ten-year R-star eight quarters into the recovery phase is 1.5 percentage lower in a boom-bust cycle than in a V-shaped recession. Thus, by this metric, boom-bust recessions of the same depth require persistently more accommodation than V-shaped recessions.

More generally, our results imply that countercyclical monetary policy needs to look very different to close the output gap across these two recession types. This suggests that appropriate countercyclical policy relies on a correct diagnosis of the recession.

[24] Our concept of R-star coincides with the natural rate of interest if the "divine coincidence" holds conditional on the shocks that drive our estimated RRIRF (Galí, 2008).

[25] 10-year rate is FRED series GS10. Prior to April 1953 we splice it with the long-term bond series LTGOVTBD. 10-year inflation expectation is PTR from the FRBUS history database.

Figure 10: $\mathcal{R}^*$ Path for V-shaped and Boom-Bust Recession



7 Conclusion

The recoveries from the 1990-1, 2001, and 2007-9 recessions have been weak relative to other postwar business cycles. We provide new cross-state evidence that these recessions are caused by different types of shocks: for the 1990-1, 2001, and 2007-9 recessions we estimate that more exposed states are subject to a boom-bust cycle, whereas for the pre-1990 and 2020 recessions we estimate that more exposed states are subject to a deeper V-shaped recessions. We argue that these cross-sectional estimates plausibly aggregate, so that weak aggregate recoveries from the 1990-1, 2001, and 2007-9 recessions reflect the boom-bust nature of the cycle compared to the V-shaped pre-1990 and Covid recessions. Closing the output gap requires very different monetary policy responses across these recession types, so that appropriate countercyclical policy depends importantly on a correct diagnosis of the recession.

References

Adao, R., M. Kolesár, and E. Morales (2019): "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 134(4), 1949–2010.

Amior, M., and A. Manning (2018): "The persistence of local joblessness," *American Economic Review*, 108(7), 1942–1970.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-cycle anatomy," *American Economic Review*, 110(10), 3030–3070.

Artis, M., C. Dreger, and K. Kholodilin (2011): "What drives regional business cycles? The role of common and spatial components," *The Manchester School*, 79(5), 1035–1044.

Autor, D., D. Dorn, and G. H. Hanson (2021): "On the Persistence of the China Shock," *Brookings Papers on Economic Activity*, 52(2), 381–447.

Bachmann, R. (2012): "Understanding the jobless recoveries after 1991 and 2001," *Unpublished Manuscript, University of Notre Dame*.

Ball, L. (2015): "Comment on 'Inflation and activity' by Olivier Blanchard, Eugenio Cerutti and Lawrence Summers," *ECB, Inflation and Unemployment in Europe, Frankfurt am Main, October*, pp. 47–52.

Barnichon, R., and G. Mesters (2023): "A sufficient statistics approach for macro policy," *American Economic Review*, 113(11), 2809–2845.

Barro, R. J., and X. Sala-i Martin (1991): "Convergence across states and regions," *Brookings Papers on Economic Activity*, 1991(1), 107–182.

Beaudry, P., D. Galizia, and F. Portier (2018): "Reconciling Hayek's and Keynes' views of recessions," *The Review of Economic Studies*, 85(1), 119–156.

Benigno, G., and L. Fornaro (2018): "Stagnation traps," *The Review of Economic Studies*, 85(3), 1425–1470.

Beraja, M., E. Hurst, and J. Ospina (2019): "The aggregate implications of regional business cycles," *Econometrica*, 87(6), 1789–1833.

Beraja, M., and C. K. Wolf (2022): "Demand Composition and the Strength of Recoveries," Working Paper 29304, National Bureau of Economic Research.

Bernanke, B. S. (2003): "The jobless recovery," in *Remarks by Governor Ben S. Bernanke at the Global Economic and Investment Outlook Conference, Carnegie Mellon University, Pittsburgh Pennsylvania*.

Blanchard, O., E. Cerutti, and L. Summers (2015): "Inflation and activity–two explorations and their monetary policy implications," Discussion paper, National Bureau of Economic Research.

Blanchard, O. J., and L. F. Katz (1992): "Regional evolutions," *Brookings Papers on Economic Activity*, 1992(1), 1–75.

Bordalo, P., N. Gennaioli, A. Shleifer, and S. J. Terry (2026): "Real credit cycles," *American Economic Review*, 116(4), 1274–1308.

Borusyak, K., P. Hull, and X. Jaravel (2022): "Quasi-Experimental Shift-Share Research Designs," *The Review of economic studies*, 89(1), 181–213.

Caravello, T. E., A. McKay, and C. K. Wolf (2024): "Evaluating Monetary Policy Counterfactuals:(When) Do We Need Structural Models?," Discussion paper, National Bureau of Economic Research.

Carlino, G., and K. Sill (2001): "Regional income fluctuations: Common trends and common cycles," *Review of Economics and Statistics*, 83(3), 446–456.

Carvalho, V. M., and A. C. Harvey (2005): "Growth, cycles and convergence in US regional time series," *International Journal of Forecasting*, 21(4), 667–686.

Cerra, V., A. Fatás, and S. C. Saxena (2023): "Hysteresis and business cycles," *Journal of Economic Literature*, 61(1), 181–225.

Donaldson, P. (2026): "Cross-sectional Identification with Heterogeneous Exposure to General Equilibrium Effects," Ph.D. thesis.

Dupraz, S., E. Nakamura, and J. Steinsson (2025): "A plucking model of business cycles," *Journal of Monetary Economics*, 152, 103766.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2019): "A model of secular stagnation: Theory and quantitative evaluation," *American Economic Journal: Macroeconomics*, 11(1), 1–48.

Fernald, J., R. Inklaar, and D. Ruzic (2025): "The productivity slowdown in advanced economies: common shocks or common trends?," *Review of Income and Wealth*, 71(1), e12690.

Fernald, J. G., R. E. Hall, J. H. Stock, and M. W. Watson (2017): "The disappointing recovery of output after 2009," Discussion paper, National Bureau of Economic Research.

Fieldhouse, A. J., D. Munro, C. Koch, and S. Howard (2024): "The emergence of a uniform business cycle in the United States: Evidence from new claims-based unemployment data," *Brookings Papers on Economic Activity*, 55(1), 265–342.

Foerster, A. T., A. Hornstein, P.-D. G. Sarte, and M. W. Watson (2022): "Aggregate implications of changing sectoral trends," *Journal of Political Economy*, 130(12), 3286–3333.

Friedman, M. (1993): "The "plucking model" of business fluctuations revisited," *Economic Inquiry*, 31(2), 171–177.

Fukui, M., E. Nakamura, and J. Steinsson (2023): "Women, wealth effects, and slow recoveries," *American Economic Journal: Macroeconomics*, 15(1), 269–313.

Galí, J. (2008): *Monetary policy, inflation, and the business cycle*, vol. 2. Princeton University Press Princeton, NJ.

Galí, J., F. Smets, and R. Wouters (2012): "Slow recoveries: A structural interpretation," *Journal of Money, Credit and Banking*, 44, 9–30.

Garga, V., and S. R. Singh (2021): "Output hysteresis and optimal monetary policy," *Journal of Monetary Economics*, 117, 871–886.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik instruments: What, when, why, and how," *American Economic Review*, 110(8), 2586–2624.

Groshen, E. L., and S. Potter (2005): "Has structural change contributed to a jobless recovery?," *Available at SSRN 683258*.

Hershbein, B., and B. A. Stuart (2024): "The evolution of local labor markets after recessions," *American Economic Journal: Applied Economics*, 16(3), 399–435.

House, C. L., C. Proebsting, and L. L. Tesar (2026): "Quantifying the benefits of labour mobility in a currency union," *Review of Economic Studies*, 93(2), 1038–1076.

Howard, G., R. Martin, and B. A. Wilson (2011): "Are recoveries from banking and financial crises really so different?," International Finance Discussion Paper 1037, Board of Governors of the Federal Reserve System.

Imbens, G. W., and M. Kolesar (2016): "Robust standard errors in small samples: Some practical advice," *Review of Economics and Statistics*, 98(4), 701–712.

Jaimovich, N., and H. E. Siu (2020): "Job polarization and jobless recoveries," *Review of Economics and Statistics*, 102(1), 129–147.

Jordà, Ò., M. Schularick, and A. M. Taylor (2013): "When credit bites back," *Journal of money, credit and banking*, 45(s2), 3–28.

Kehoe, P. J., V. Midrigan, and E. Pastorino (2018): "Evolution of modern business cycle models: Accounting for the great recession," *Journal of Economic Perspectives*, 32(3), 141–166.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2019): "The tail that keeps the riskless rate low," *NBER Macroeconomics Annual*, 33(1), 253–283.

——— (2020): "The tail that wags the economy: Beliefs and persistent stagnation," *Journal of Political Economy*, 128(8), 2839–2879.

Leamer, E. E. (2021): "Why Are Some Recoveries Short and Others Long?," Discussion paper, National Bureau of Economic Research.

McKay, A., and C. K. Wolf (2023): "What can time-series regressions tell us about policy counterfactuals?," *Econometrica*, 91(5), 1695–1725.

Mian, A., A. Sufi, and E. Verner (2017): "Household debt and business cycles worldwide," *The Quarterly Journal of Economics*, 132(4), 1755–1817.

Quah, D. (1993): "Galton's fallacy and tests of the convergence hypothesis," *The Scandinavian Journal of Economics*, 95(4), 427–443.

Reifschneider, D., W. Wascher, and D. Wilcox (2015): "Aggregate supply in the United States: recent developments and implications for the conduct of monetary policy," *IMF Economic Review*, 63(1), 71–109.

Rognlie, M., A. Shleifer, and A. Simsek (2018): "Investment hangover and the great recession," *American Economic Journal: Macroeconomics*, 10(2), 113–153.

Schaal, E., and M. Taschereau-Dumouchel (2023): "Herding through booms and busts," *Journal of Economic Theory*, 210, 105669.

Stock, J. H., and M. W. Watson (1999): "Business cycle fluctuations in US macroeconomic time series," *Handbook of macroeconomics*, 1, 3–64.

Yagan, D. (2019): "Employment hysteresis from the great recession," *Journal of Political Economy*, 127(5), 2505–2558.

A Model of Regional Exposure to National Recessions

We develop a multi-region, multi-good new Keynesian model, in which regions are heterogeneously exposed to the aggregate business cycle that is nested by the DGP in section 2. We go through the simple model in detail, and at the end we describe extensions that the DGP also nests.

A.1 Households A representative household maximizes utility

$$\max_{\{C_{g,t}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t} \beta_s \right) \left[\sum_{g=1}^{G} \alpha_{g,t} \frac{C_{g,t}^{1-\sigma_g}}{1-\sigma_g} - \chi \frac{L_t^{1+\eta}}{1+\eta} \right].$$

subject to the budget constraint,

$$\sum_{g=1}^{G} P_{g,t} C_{g,t} + B_t = W_t L_t + (1 + i_{t-1}) B_{t-1},$$

where $C_{g,t}$ is consumption of good g at time t, $P_{g,t}$ is its price, B_t are bond holdings, L_t is labor, W_t is the nominal wage, and i_t is the nominal interest rate.

We assume that the household only chooses consumption goods. Labor is not a choice variable but determined by labor demand due to sticky wages. In Section A.8 we discuss an extension of the model with hand-to-mouth consumers.

The first order conditions for the consumption goods are,

$$\alpha_{g,t} C_{g,t}^{-\sigma_g} = \lambda_t \frac{P_{g,t}}{P_t}, \qquad g = 1, ..., G \tag{22}$$

$$\lambda_t = \beta_t (1 + r_t) \mathbb{E}_t \lambda_{t+1} \tag{23}$$

where λ_t is the marginal utility of consumption, P_t is an aggregate price index, and $1 + r_t = (1 + i_t) \frac{P_t}{P_{t+1}}$ is the real interest rate.

Heterogeneity in utility curvature σ_g makes goods differentially sensitive to aggregate shocks. Heterogeneity in curvature σ_g also affects the balanced growth path. In this model we neutralize the latter impact by setting

$$\alpha_{g,t} = \alpha_g Z_{g,t}^{\sigma_g - 1}. \tag{24}$$

where $Z_{g,t}$ is good-level productivity. This implies that in our model output of good g will grow linearly in $Z_{g,t}$ irrespective of the curvature σ_g. We make this assumption to create a separation between business cycle sensitivity and long-run growth across sectors.

A.2 Regional Industry Demand We assume a standard constant-elasticity of substitution demand system across regions for each good,

$$C_{g,t} = \left(\sum_{i=1}^{I} \alpha_{igt}^{\frac{1}{\eta_g}} C_{i,g,t}^{\frac{\eta_g - 1}{\eta_g}} \right)^{\frac{\eta_g}{\eta_g - 1}}$$

which implies that the demand for good g in region i is given by,

$$C_{i,g,t} = \alpha_{igt} \left(\frac{P_{i,g,t}}{P_{g,t}} \right)^{-\eta_g} C_{g,t}. \tag{25}$$

We allow the shares α_{igt} to vary over time across regions,

$$\alpha_{igt} = \alpha_{ig} e^{\hat{\alpha}_{it}}, \tag{26}$$

and we discuss the relevance of the region-specific demand shocks $\hat{\alpha}_{it}$ for identification below.

A.3 Firms, Technology, Prices, and Market Clearing Output of good g in region i, $Y_{i,g,t}$, is produced by a representative firm using linear technology,

$$Y_{i,g,t} = Z_{g,t} L_{i,g,t}$$

where $Z_{g,t}$ is good-level productivity and $L_{i,g,t}$ is labor used in region i for good g.

Sectoral productivity $Z_{g,t}$ grows according to an ideas production function,

$$Z_{g,t} = \theta_g X_{g,t} Z_{g,t-1}^{1-\xi} Y_{g,t-1}^{\xi}, \tag{27}$$

where $X_{g,t}$ is an exogenous productivity trend, the term $Y_{g,t-1}^{\xi}$ captures hysteresis effects of output on productivity when $\xi > 0$, and θ_g is set such that $X_{g,t}$ is the gross productivity growth rate when there are no output fluctuations. We define the exogenous productivity path with

$$\tilde{Z}_{g,t} \equiv X_{g,t} \tilde{Z}_{g,t-1}.$$

Firms are perfectly competitive and set prices equal to marginal cost,

$$P_{i,g,t} = \frac{W_{i,g,t}}{Z_{g,t}}$$

We assume wages are indexed to an aggregate productivity index,

$$Z_t = \prod_{g=1}^{G} Z_{g,t}^{\alpha_g}, \tag{28}$$

and aggregate prices P_t, but are otherwise insensitive to the business cycle,

$$W_{i,g,t} = \bar{W} P_t Z_t$$

This assumption fulfills two objectives. First, over the long-run wages are consistent with a balanced growth path with full labor mobility across sectors. Second, in the short-run, prices and wages insensitive to the business cycle,

$$\frac{P_{i,g,t}}{P_t} = \frac{P_{g,t}}{P_t} = \bar{W} \frac{Z_t}{Z_{g,t}}. \tag{29}$$

In Section A.8 we discuss an extension in which prices respond endogenously to the local business cycle.

Regional goods market clear when output equals consumption,

$$Y_{i,g,t} = C_{i,g,t} \tag{30}$$

A.4 Equilibrium Given exogenous known sequences for the real interest rate $\{r_t\}_{t=0}^{\infty}$, the discount factor $\{\beta_t\}_{t=0}^{\infty}$, expectations of long-run marginal utility wealth, $\{\mathbb{E}_t \lambda_\infty\}_{t=0}^{\infty}$, regional demand $\{\hat{\alpha}_{it}\}_{t=0}^{\infty}$, and sectoral productivity $\{\{X_{g,t}\}_{g=1}^{G}\}_{t=0}^{\infty}$ the equilibrium is a sequence of $\{\lambda_t, Z_t, Z_{g,t}, C_{g,t}, \frac{P_{g,t}}{P_t}, C_{i,g,t}, Y_{i,g,t}, \frac{P_{i,g,t}}{P_t}\}$ that satisfy (22)-(30).

We log-linearize this equilibrium around a "balanced growth" equilibrium with the same productivity paths $\{\{Z_{g,t}\}_{g=1}^{G}\}_{t=0}^{\infty}$ and all other sequences at steady state. We denote the outcomes of the balanced growth equilibrium by $\bar{X}_{igt}$ for any variable X, lower case letters are logs $x = \ln X$, and the outcomes in the log-linearized equilibrium by $\hat{x}_{igt} = x_{igt} - \bar{x}_{igt} = \ln X_{igt} - \ln \bar{X}_{igt}$.

A.5 Aggregate Business Cycle The ARIRF is the sum of the deviation from the balanced growth path and any hysteresis effects to the balanced growth path,

$$\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s = -\left(\sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g} \right) \left[\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right] \tag{31}$$

The Lagrange multiplier can be solely expressed in terms of the exogenous sequences of

the equilibrium,

$$\hat{\lambda}_t = \frac{1}{1-\xi\sum_{g=1}^{G}\frac{\alpha_g}{\sigma_g}}\sum_{s=0}^{\infty}\left(\frac{1}{1-\xi\sum_{g=1}^{G}\frac{\alpha_g}{\sigma_g}}\right)^{s}(\hat{\beta}_{t+s}+\hat{r}_{t+s}) + \lim_{T\to\infty}\left(\frac{1}{1-\xi\sum_{g=1}^{G}\frac{\alpha_g}{\sigma_g}}\right)^{T}\mathbb{E}_t\hat{\lambda}_T.$$

The growth rate of output is the sum of the growth rate of the business cycle and trend growth,

$$\Delta^h y_{t+h} = \underbrace{\Delta^h \hat{y}_{t+h} + \xi\sum_{s=0}^{h-1}\hat{y}_{t+s}}_{\text{ARIRF}} + \underbrace{\Delta^h \tilde{z}_{t+h}}_{\text{Trend Growth}}. \tag{32}$$

where Δ^h is an h-period difference. This equation is nested by our specification of aggregate output as the sum of a cyclical and a trend component (1).

A.6 Regional Bartik Structure Combining market clearing (30), regional demand (25), and sticky relative prices (29) implies that regional output of good g is proportional to aggregate output of good g and regional demand shifters $\hat{\alpha}_{it}$,

$$Y_{i,g,t} = \alpha_{igt}\left(\frac{P_{i,g,t}}{P_{g,t}}\right)^{-\eta_g}Y_{g,t} = \alpha_{ig}e^{\hat{\alpha}_{it}}Y_{g,t}.$$

This setting therefore implies a Bartik structure for local real output growth,

$$\underbrace{\Delta^h y_{i,t+h}}_{\text{Local Output Growth}} = \underbrace{\Delta^h \hat{\alpha}_{i,t+h}}_{\text{Local Demand Shifter}} + \underbrace{\sum_{g=1}^{G}\mu_{i,g,t}\Delta^h y_{g,t+h}}_{\text{Bartik Predicted Growth, } y^b_{i,t,t+h}} \tag{33}$$

where $\mu_{i,g,t} = \frac{P_{i,g,t}Y_{i,g,t}}{P_{i,t}Y_{i,t}}$ is the nominal output share of good g in region i at t.

We can further decompose industry growth into a cyclical and a trend component,

$$\Delta^h y_{g,t+h} = \frac{1}{\sigma_g}\left(\Delta^h \hat{y}_{t+h} + \xi\sum_{s=0}^{h-1}\hat{y}_{t+s}\right) + \Delta^h \tilde{z}_{g,t+h} \tag{34}$$

which is nested by the DGP for industry growth (7).

This allows us to decompose Bartik predicted growth $y^b_{i,t,t+h}$ into a business cycle $\hat{y}^b_{i,t,t+h}$ and a trend component $\bar{y}^b_{i,t,t+h}$,

$$y^b_{i,t,t+h} \equiv \hat{y}^b_{i,t,t+h} + \bar{y}^b_{i,t,t+h} \tag{35}$$

where

$$\hat{y}^b_{i,t,t+h} = \underbrace{\frac{\sum_{g=1}^{G} \mu_{i,g,t} \frac{1}{\sigma_g}}{\sum_{g=1}^{G} \mu_{g,t} \frac{1}{\sigma_g}}}_{\text{Relative Regional Business Cycle Exposure}} \times \underbrace{\left(\Delta^h \hat{y}_{t+h} + \xi \sum_{s=0}^{h-1} \hat{y}_{t+s} \right)}_{\text{ARIRF}} \tag{36}$$

$$\bar{y}^b_{i,t,t+h} = \sum_{g=1}^{G} \underbrace{\mu_{i,g,t}}_{\text{Regional Trend Exposure}} \times \underbrace{\Delta^h \tilde{z}_{g,t+h}}_{\text{Sectoral Trend}}$$

This decomposed local output growth,

$$\Delta^h y_{i,t+h} = \Delta^h \hat{\alpha}_{i,t+h} + \hat{y}^b_{i,t,t+h} + \bar{y}^b_{i,t,t+h}, \tag{37}$$

is nested by the DGP for local output (3), which allows for idiosyncratic, cyclical, and trend components.

A.7 Conclusion Since this model's aggregate, regional, and industry output is nested by the DGP in Section 2, the same identification argument in Proposition 1 applies.

A.8 Extensions In Appendix B, we extend the model to include regional hand-to-mouth consumers, prices that respond to the local recession, and endogenous monetary policy. We show that in each case Proposition 1 continues to hold unchanged. Intuitively, each of these extensions either scales up or down the effect of the aggregate business cycle shock on regions, which in turn scales the first-stage and reduced-form estimates, or it rescales the aggregate business cycle shock itself. But since our IV estimator normalizes the scale to a 1% relative recession, we recover exactly the same (scaled) aggregate recession impulse response function.

B Simple Model Extension

B.1 Hand-to-Mouth Households We assume a fraction μ of all households are hand-to-mouth. Hand-to-mouth households are distributed across regions proportional to output across regions. They earn income a fraction μ of local income and only spend on local goods.

Market clearing for each good is,

$$Y_{i,g,t} = (1-\mu) C^o_{i,g,t} + \mu C^m_{i,g,t} \tag{38}$$

where $C^o_{i,g,t}$ is per capita consumption by the optimizing representative household and $C^m_{i,g,t}$ is per capita consumption by hand-to-mouth households.

Hand-to-mouth consumers have consumption shares α_{gt} for each good g, similar to optimizing households,

$$C^m_{i,g,t} = \alpha_{gt} \left(\frac{P_{i,g,t}}{P_{it}} \right)^{-\sigma_g} C^m_{it}. \tag{39}$$

Given optimal prices (29), we get

$$C^m_{g,t} = \alpha_g Y_t \frac{Z_{g,t}}{Z_t} \tag{40}$$

$$C^o_{g,t} = \alpha_g \left(\frac{\lambda_t}{Z_t} \right)^{-\frac{1}{\sigma_g}} Z_{g,t} \tag{41}$$

Market clearing implies,

$$Y_{g,t} = (1-\mu)\alpha_g \left(\frac{\lambda_t}{Z_t} \right)^{-\frac{1}{\sigma_g}} Z_{g,t} + \alpha_g \mu Y_t \frac{Z_{g,t}}{Z_t} \tag{42}$$

Total local output is,

$$Y_{i,t} = \sum_{g=1}^{G} (1-\mu)\alpha_{ig} e^{\hat{\alpha_{it}}} \left(\alpha_g \left(\frac{\lambda_t}{Z_t} \right)^{-\frac{1}{\sigma_g}} Z_{g,t} \right) + \mu C^m_{i,t} \tag{43}$$

$$= e^{\hat{\alpha_{it}}} \sum_{g=1}^{G} \alpha_{ig} Y_{g,t} - e^{\hat{\alpha_{it}}} \mu Y_t \frac{Z_{it}}{Z_t} + \mu Y_{i,t} \tag{44}$$

$$= e^{\hat{\alpha_{it}}} \left[\frac{1}{1-\mu} \sum_{g=1}^{G} \alpha_{ig} Y_{g,t} - \frac{\mu}{1-\mu} Y_t \frac{Z_{it}}{Z_t} \right] \tag{45}$$

The log-linear approximation around the trend growth path is,

$$\hat{y}_{g,t} = -\frac{1}{\sigma_g}(1-\mu)\hat{\lambda}_t + \mu \hat{y}_t \tag{46}$$

$$\hat{y}_t = -\sum_{g=1}^{G} \left(\alpha_g \frac{1}{\sigma_g} \right) \hat{\lambda}_t \tag{47}$$

Thus, both industry and aggregate output can be decomposed into a cyclical and a trend piece as in Section 2.

Local output growth is

$$\Delta^h y_{i,t,t+h} \approx \hat{\alpha_{it}} + \frac{1}{1-\mu} \sum_{g=1}^{G} \mu_{i,g,t} \Delta^h y_{g,t+h} + \frac{\mu}{1-\mu} \Delta^h y_{t+h} \tag{48}$$

which using the industry and aggregate output decomposition is again nested by (3).

B.2 Local Prices Respond to Recession We assume that the relative price of good g in region i relative to the industry price is set based on the relative recession in region i with elasticity κ_I,

$$\frac{P_{i,g,t}}{P_{g,t}} = \frac{\bar{W}_{ig}}{\bar{W}_g} \left(\frac{Y_{i,t}/Z_{i,t}}{Y_t/Z_t} \right)^{\kappa_I} \tag{49}$$

and the industry price relative to the aggregate price is based on the goods-level relative recession with elasticity $\kappa_G < \min_g \sigma_g$,

$$\frac{P_{g,t}}{P_t} = \frac{\bar{W}_g}{\bar{W}} \left(\frac{Y_{g,t}/Z_{g,t}}{Y_t/Z_t} \right)^{\kappa_G} \frac{Z_t}{Z_{g,t}} \tag{50}$$

where $\kappa_G = \kappa_I = 0$ (acyclical relative prices) corresponds to the baseline model.

Then local output of good g is equal to

$$Y_{i,g,t} = \alpha_{igt} \left(\frac{\bar{W}_g}{\bar{W}} \right)^{-\eta} \left(\frac{Y_{i,t}/Z_{i,t}}{Y_t/Z_t} \right)^{-\eta\kappa_I} Y_{g,t} \tag{51}$$

and aggregate output of good g is equal to,

$$Y_{g,t} = Z_{g,t} \left[\alpha_g \left(\frac{\bar{W}_g}{\bar{W}} \right)^{-1} \left(\frac{Y_t}{Z_t} \right)^{\kappa_G} \lambda_t^{-1} Z_t \right]^{\frac{1}{\sigma_g - \kappa_G}} \tag{52}$$

Local output growth has a Bartik structure

$$\underbrace{\Delta^h y_{i,t+h}}_{\text{Local Output Growth}} = \frac{1}{1+\eta\kappa_I} \underbrace{\Delta^h \hat{\alpha}_{i,t+h}}_{\text{Local Demand Shifter}} + \frac{1}{1+\eta\kappa_I} \underbrace{\sum_{g=1}^{G} \mu_{i,g,t} \Delta^h y_{g,t+h}}_{\text{Bartik Predicted Growth, } y^B_{i,t,t+h}} + \frac{\eta\kappa_I}{1+\eta\kappa_I} \underbrace{\Delta^h y_{t+h}}_{\text{Aggregate Growth}} \tag{53}$$

Define $\tilde{\sigma}_g \equiv \sigma_g - \kappa_G > 0$. The cyclical component of industry growth is,

$$\hat{y}_{g,t} + \xi \sum_{s=0}^{t-1} \hat{y}_{g,s} = \frac{\kappa_G}{\tilde{\sigma}_g} \left(\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s \right) - \frac{1}{\tilde{\sigma}_g} \left(\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right) \tag{54}$$

The cyclical component of output is,

$$\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s = - \left(\frac{\sum_{g=1}^{G} \mu_{gt} \frac{1}{\tilde{\sigma}_g}}{1 - \kappa_G \sum_{g=1}^{G} \mu_{gt} \frac{1}{\tilde{\sigma}_g}} \right) \left(\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right) \tag{55}$$

which implies,

$$\hat{y}_{g,t} + \xi \sum_{s=0}^{t-1} \hat{y}_{g,s} = \frac{1}{\tilde{\sigma}_g} \left[\kappa_G + \frac{1 - \kappa_G \sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}}{\sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}} \right] \left(\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s \right) \tag{56}$$

The cyclical component of Bartik predicted growth is,

$$\hat{y}^b_{i,t,t+h} = \underbrace{\left(\frac{\sum_{g=1}^{G} \mu_{i,g,t} \frac{1}{\tilde{\sigma}_g}}{\sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}} \right)}_{\text{Relative Regional Business Cycle Exposure}} \times \underbrace{\left(\Delta^h \hat{y}_{t+h} + \xi \sum_{s=0}^{h-1} \hat{y}_{t+s} \right)}_{\text{Aggregate Business Cycle}} \tag{57}$$

Thus, the model is again nested by the DGP in Section 2 and Proposition 1.

B.3 Endogenous Monetary Policy Response Suppose the central bank follows a Taylor rule in which the real interest rate responds to the cyclical component of output,

$$\hat{r}_t = \phi \hat{y}_t + \hat{\nu}_t \tag{58}$$

where $\hat{\nu}_t$ is a monetary policy shock.

Then the only difference from the baseline model is the solution for $\hat{\lambda}_t$,

$$\hat{\lambda}_t = \frac{1}{1 + (\phi - \xi) \sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}} \sum_{s=0}^{\infty} \left(\frac{1}{1 + (\phi - \xi) \sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}} \right)^s (\hat{\beta}_{t+s} + \hat{\nu}_{t+s}) \tag{59}$$

$$+ \lim_{T \to \infty} \left(\frac{1}{1 + (\phi - \xi) \sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}} \right)^T \mathbb{E}_t \lambda_T. \tag{60}$$

Since no other equation changes, this model is nested by the general DGP in Section 2 and Proposition 1 goes through unchanged.

C Proof of Proposition 1

Let

$$Z_i = \tilde{\hat{y}}^b_{i,P,T}, \qquad D_i = y_{iP} - y_{iT}, \qquad Y_{ih} = y_{iT+h} - y_{iP}. \tag{61}$$

Let $x_i^{\perp}$ denote the residual from projecting x_i on a constant and the controls $X_{i,P}$. By the Frisch-Waugh-Lovell theorem, the 2SLS estimator can then be written as

$$\beta_h^{IV} = \frac{I^{-1} \sum_i Z_i^{\perp} Y_{ih}}{I^{-1} \sum_i Z_i^{\perp} D_i}. \tag{62}$$

Using the decomposition of local output, we can write observed regional output growth as

$$Y_{ih} = \hat{y}_{i,T+h} - \hat{y}_{i,P} + N_{i,h}, \tag{63}$$

where $N_{i,h} = \bar{y}_{i,T+h} - \bar{y}_{i,P} + \mathring{y}_{i,T+h} - \mathring{y}_{i,P}$ denotes the non-cyclical component of local output growth. Replacing (63) in (62) yields

$$\beta_h^{IV} = \frac{I^{-1}\sum_i Z_i^{\perp}(\hat{y}_{i,T+h} - \hat{y}_{i,P} + N_{i,h})}{I^{-1}\sum_i Z_i^{\perp}(\hat{y}_{i,P} - \hat{y}_{i,T} - N_{i,0})}. \tag{64}$$

Under standard regularity conditions that guarantee that cross-sectional sample moments satisfy a law of large number, the numerator in (64) converges to[26]

$$I^{-1}\sum_i Z_i^{\perp}(\hat{y}_{i,T+h} - \hat{y}_{i,P} + N_{i,h}) \xrightarrow{p} Cov_i[Z_i^{\perp}, (\hat{y}_{i,T+h} - \hat{y}_{i,P})] + Cov_i[Z_i^{\perp}, N_{i,h}]. \tag{65}$$

By the exclusion restriction $Cov_i[Z_i^{\perp}, N_{i,h}] = 0$, thus

$$I^{-1}\sum_i Z_i^{\perp}(\hat{y}_{i,T+h} - \hat{y}_{i,P} + N_{i,h}) \xrightarrow{p} Cov_i[Z_i^{\perp}, (\hat{y}_{i,T+h} - \hat{y}_{i,P})]. \tag{66}$$

By the same argument, the denominator in in (64) converges to

$$I^{-1}\sum_i Z_i^{\perp}(\hat{y}_{i,P} - \hat{y}_{i,T} - N_{i,0}) \xrightarrow{p} Cov_i[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})] \neq 0, \tag{67}$$

where the last inequality follows from the relevance condition. By the continuous mapping theorem it follows that

$$\beta_h^{IV} \xrightarrow{p} \frac{Cov_i[Z_i^{\perp}, (\hat{y}_{i,T+h} - \hat{y}_{i,P})]}{Cov_i[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})]}. \tag{68}$$

Using the definition of the RRIRF, we get

$$\beta_h^{IV} \xrightarrow{p} \frac{Cov_i\Big[Z_i^{\perp}, RRIRF_{ih}(\hat{y}_{i,P} - \hat{y}_{i,T} - (\hat{y}_P - \hat{y}_T)) + \hat{y}_h\Big]}{Cov_i[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})]}. \tag{69}$$

Dropping out aggregate terms and rearranging yields,

$$\beta_h^{IV} \xrightarrow{p} \frac{Cov_i\Big[Z_i^{\perp}, RRIRF_{ih}(\hat{y}_{i,P} - \hat{y}_{i,T}) - RRIRF_{ih}(\hat{y}_P - \hat{y}_T)\Big]}{Cov_i[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})]}. \tag{70}$$

[26] We assume finite second moments exist for Z_i, Y_{ih}, and D_i, and sufficiently weak cross-sectional dependence such that the sample covariances converge in probability to their population counterparts.

Under the assumption of homogeneous RRIRF, $RRIRF_{ih} = RRIRF_h$, and

$$\beta_h^{IV} \xrightarrow{p} RRIRF_h \frac{Cov_i\left[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})\right]}{Cov_i[Z_i^{\perp}, (\hat{y}_{i,P} - \hat{y}_{i,T})]} = RRIRF_h. \tag{71}$$

D Regional Convergence and Aggregation

In this appendix we calculate the plausible aggregation errors when convergence forces operate across regions so that the one-factor structure (17) does not hold. In the following subsections we first consider a statistical error correction model (ECM) and analytically derive and quantify the aggregation bias in this framework. In subsection D.5 we solve a two-region model with forward-looking migration decisions. This model yields the same bias as the ECM for permanent shocks. It also shows that for transitory shocks, the endogenous convergence forces and consequently the aggregation are much smaller.

D.1 Error Correction Model To capture regional convergence forces we extend the single factor structure in (17) to include an error correction term. Let $d_{it} = \hat{y}_{it} - \hat{y}_t$, then the error correction model is,

$$\Delta d_{it} = \theta_i \Delta \hat{y}_t - (1 - \rho) d_{i,t-1}, \tag{72}$$

where ρ determines the persistence of gaps between a region and the aggregate. $\rho = 1$ corresponds to the baseline case of no independent regional convergence beyond that from the aggregate shock.

There are two important caveats in studying the error correction model (72). First, it does not allow for permanent hysteresis effects in the cross-section as for instance documented in Hershbein and Stuart (2024) and Yagan (2019). To allow for these while still maintaining some convergence in the ECM one would have to redefine d_{it} as net of any differential permanent effect. Then the permanent effects aggregate, and the aggregation error only comes from the transitory component that is subject to convergence forces. Of course, that same evidence can be interpreted as an absence of convergence ($\rho = 1$) in which case the transitory component also aggregates without error.

Second, the ECM is a statistical model and does not capture that the speed of convergence ρ likely depends on the underlying shock. Migration and pricing decisions that generate convergence are forward-looking. We should therefore expect that regional convergence forces are weaker for transitory shocks than for permanent shocks, as households and firms can wait out transitory shocks and avoid paying the fixed costs associated with relocation or price changes. In subsection D.5 below, we show that convergence is much weaker and the

aggregation error is much smaller in a simple model with forward-looking migration decisions calibrated to match evidence on permanent shocks.

D.2 Bias Formula We next solve for the aggregation bias from the RRIRF to the ARIRF given the ECM (72). Iterating backward yields,

$$d_{it} = \theta_i \sum_{k=0}^{\infty} \rho^k \Delta \hat{y}_{t-k} \tag{73}$$

$$= \theta_i \left(\hat{y}_t - (1-\rho) \sum_{k=1}^{\infty} \rho^{k-1} \hat{y}_{t-k} \right), \tag{74}$$

so that regional differences load on a distributed lag of aggregate cyclical output x_t rather than only on its contemporaneous value.

Define,

$$x_t \equiv \hat{y}_t + (1-\rho) C_t \tag{75}$$

where $C_t \equiv -\sum_{k=1}^{\infty} \rho^{k-1} \hat{y}_{t-k}$.

The RRIRF is then,

$$RRIRF_h = \frac{\theta_i (x_{T+h} - x_P)}{\theta_i (x_P - x_T)} = \frac{\hat{y}_{T+h} - \hat{y}_P + (1-\rho)(C_{T+h} - C_P)}{\hat{y}_P - \hat{y}_T + (1-\rho)(C_P - C_T)}. \tag{76}$$

Given that the ECM model only applies to a transitory component of cyclical output and that recessions are of relatively short duration (though recoveries may be long), we approximate the RRIRF using $C_P - C_T \approx 0$,

$$RRIRF_h \approx \frac{\hat{y}_{T+h} - \hat{y}_P + (1-\rho)(C_{T+h} - C_P)}{\hat{y}_P - \hat{y}_T} \tag{77}$$

$$= ARIRF_h + (1-\rho) \frac{C_{T+h} - C_P}{\hat{y}_P - \hat{y}_T} \tag{78}$$

$$= ARIRF_h - (1-\rho) \frac{\sum_{k=1}^{\infty} \rho^{k-1} (\hat{y}_{T+h-k} - \hat{y}_{P-k})}{\hat{y}_P - \hat{y}_T} \tag{79}$$

$$= ARIRF_h - (1-\rho) \frac{\sum_{k=1}^{\infty} \rho^{k-1} [(\hat{y}_{T+h-k} - \hat{y}_P) - (\hat{y}_{P-k} - \hat{y}_P)]}{\hat{y}_P - \hat{y}_T} \tag{80}$$

$$= ARIRF_h - (1-\rho) \sum_{k=1}^{\infty} \rho^{k-1} [ARIRF_{h-k} - ARIRF_{P-T-k}] \tag{81}$$

This shows that the RRIRF can be expressed solely as a function of the ARIRF and the convergence rate $1-\rho$. The scale of the recession is irrelevant. The dynamics of the recession matter as they determine the ARIRF.

The aggregation bias is then approximately,

$$E_h = RRIRF_h - ARIRF_h \approx -(1-\rho)\sum_{k=1}^{\infty}\rho^{k-1}[ARIRF_{h-k} - ARIRF_{P-T-k}] \tag{82}$$

We can use this result in two ways. First, given an ARIRF and ρ we can calculate the aggregation bias. Second, we can use the formula to invert the RRIRF estimates and uncover the true ARIRF given an estimate for ρ.

As a first illustration of the plausible aggregation bias, consider the case where the ARIRF is 0 at and prior to the peak and then follows an AR(1) with persistence ζ. The bias is then,

$$RRIRF_h - ARIRF_h = (1-\rho)\sum_{k=1}^{h}\rho^{k-1}\zeta^{h-k} = (1-\rho)\frac{\zeta^h - \rho^h}{\zeta - \rho}. \tag{83}$$

Given that the ARIRF is generally negative, this formula shows that the RRIRF is upward-biased except at $h = 0$. This means the RRIRF will show faster convergence than the ARIRF. It may even overshoot, depending on parameters.

D.3 Calibrated Example The empirically relevant case is when the transitory component of the recession decays faster than regional differences $\zeta < \rho$ (House, Proebsting, and Tesar, 2026). Estimates suggest that regional convergence in output and employment rates are low. Amior and Manning (2018) estimate a decadal persistence of 0.68, Carvalho and Harvey (2005) above 0.8, and Barro and Sala-i Martin (1991) an equivalent decadal persistence of 0.817.[27] Meanwhile Hershbein and Stuart (2024) and Yagan (2019) document an absence of regional convergence, and Quah (1993) argues that estimates are biased towards finding convergence. In addition, these estimates all load at least in part on permanent shocks, which would strengthen any fundamental convergence force. As a baseline we set $\rho = 0.817^{1/40}$ at quarterly frequency, which we interpret as biased towards convergence given our reading of the evidence.

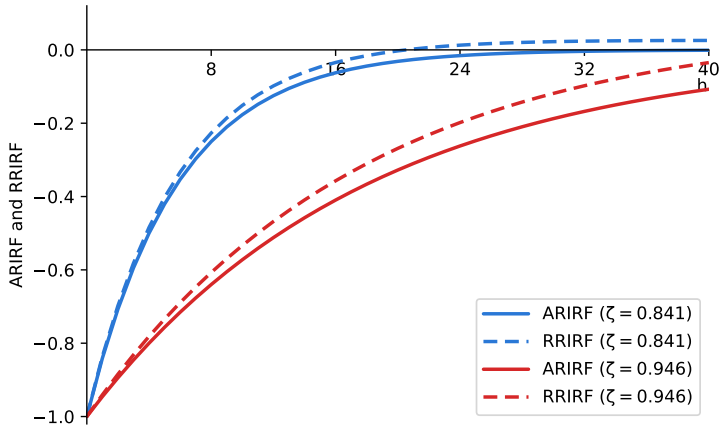
Figure 11 plots the ARIRF and the RRIRF for a recession with $\zeta = 0.5^{\frac{1}{4}}$—a fast recovery in which half the gap is closed every year—and a much slower recovery with $\zeta = 0.75^{\frac{1}{4}}$, consistent with the persistence of unemployment in the recovery from the Great Recession.[28]

For fast recoveries the bias is tiny. The Great Recession calibration is biased towards finding more convergence, but the bias remains negligible. Note that both examples display

[27] Amior and Manning (2018) argue that the faster convergence rates in Blanchard and Katz (1992) are likely due to measurement error.

[28] The increase in unemployment was 5.6 percentage points from May 2007 to October 2009. Five years later in October 2014 the unemployment rate was 1.3 percentage points above its May 2007 level. This yields an annual persistence of $(1.3/5.6)^{1/5} = 0.75$.

Figure 11: Aggregation Bias in Recessions with Fast and Slow Recoveries

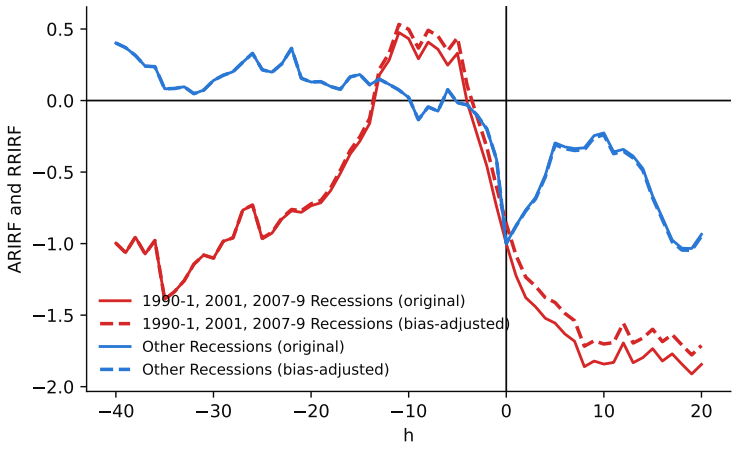


overshooting in regional comparisons (though for the Great Recession calibration this happens after 10 years). We do not estimate overshooting in the data even for fast recoveries consistent with House, Proebsting, and Tesar (2026). This suggests that this calibration may still overstate the importance of regional convergence.

D.4 Bias-Correcting the Estimates We next invert the bias formula to recover the implied ARIRF from our empirical estimates in Figure 4. For this exercise we use the exact bias formula, not the approximation. We conservatively treat the entire estimated RRIRF as transitory, since permanent components of the RRIRF aggregate without error. Since both RRIRF estimates are slightly trending down (though not statistically significantly so) we fit a linear trend through the $h \in [-40, -36]$ average and the $h \in [16, 20]$ average for the 1990-1, 2001, and 2007-9 recessions. For the other recessions we stop the window over $h \in [8, 12]$ since subsequent horizons cover a second recession. We then adjust the end point of this linear trend such that the bias-recovered ARIRF is exactly 0 at its end point, consistent with a transitory aggregate shock. Finally we set $C_{-28} = 0$ to initialize the algorithm. For the other recessions, the implied RRIRF is close to zero prior to $h = -6$ so we set all those values to zero. As before we set $\rho = 0.817^{1/40}$.

Figure 12 shows that for the recessions with fast recoveries, there is essentially no bias. The recovery is so fast that little additional convergence is taking place. There is a slightly larger adjustment for the recessions with boom-bust cycles. According to the model, our estimates show too much divergence because they are still biased by convergence forces from the boom period (the lagged terms in C). It is clear that neither bias adjustment causes a

Figure 12: Aggregation Bias-Corrected RRIRF Estimates



meaningful difference between the RRIRF and the ARIRF.

An objection to this analysis is that in boom-bust cycles it makes little sense for convergence to operate in accordance with the prior boom when the bust has already occurred. In essence, this is a critique of the backward-looking structure of the ECM (72). In the next subsections we calculate the plausible aggregation bias in a forward-looking model.

D.5 Model with Forward-looking Migration In this subsection we lay out a two-region RBC model with forward-looking migration decision. The purpose of the model is to quantify the aggregation bias in response to transitory shocks when convergence is based on forward-looking decisions and disciplined by the convergence estimates in the literature that primarily cover permanent shocks. The key message from the model is that convergence forces and thus the aggregation bias are weaker in response to transitory shocks than implied by the ECM model (72). While the model only explicitly models migration, the same logic applies to other convergence forces, such as relative price adjustment. Furthermore, broad convergence estimates such as Barro and Sala-i Martin (1991) measure all of these mechanisms jointly, which implies that we account for them by calibrating to this evidence.

There are two locations $l \in \{1, 2\}$. Each produces a single freely traded good, so the law of one price holds and all prices are in units of the consumption good. Output in region l is produced by competitive firms using labor with a DRS production function,

$$Y_{lt} = Z_{lt} N_{lt}^{\alpha}, \qquad \alpha \in (0, 1), \tag{84}$$

where N_{lt} is regional population and Z_{lt} is regional productivity.

Perfect competition implies wages and profits are

$$w_{lt} = \alpha Z_{lt} N_{lt}^{\alpha-1} \tag{85}$$

$$\Pi_{lt} = (1-\alpha) Z_{lt} N_{lt}. \tag{86}$$

Households belong to a continuum of identical families of measure 1. Each family is atomistic, holds a representative cross-section of members in both regions, pools all labor income, profits, and moving-cost realizations, and therefore assigns every member the same consumption c_t. Families take wages, rents, and asset prices as given. Because families are identical, no assets are traded in equilibrium. Total population is fixed at $N_{1t} + N_{2t} = 1$.

Each household supplies a fixed unit of labor to its region. Households may change their labor supply location by migrating across regions. Migration of households is costly and staggered. Each quarter a household receives a migration opportunity with Calvo probability λ. They then draw an iid moving cost $\phi \sim U[0, \bar{k}]$ and decide whether to move or not. This setting allows for a upward-sloping labor supply across regions in short-run and perfectly elastic labor supply across regions in the long-run. A perfectly elastic long-run labor supply curve implies that wages will converge across regions in the long-run even in response to permanent shocks.

Given an allocation of households across regions, the maximization problem of the household is,

$$\max_{\{a_t, c_t\}} \sum_{t=0}^{\infty} \beta^t \ln c_t \tag{87}$$

$$\text{s.t. } a_t = w_{1t} N_{1t} + w_{2t} N_{2t} + \Pi_{1t} + \Pi_{2t} - c_t - \text{migration costs} + (1 + r_t) a_{t-1} \tag{88}$$

The first order conditions are standard,

$$1 = \Lambda_{t,t+1}(1 + r_{t+1}) \tag{89}$$

$$\Lambda_{t,t+1} = \beta \frac{c_t}{c_{t+1}}. \tag{90}$$

The migration decision is the optimal choice of $N_{1,t+1} = 1 - N_{2,t+1}$, which takes place after the choice c_t. Let V_{lt} denote the family's value of a member in region l and $D_t \equiv V_{1t} - V_{2t}$ the value gap. A member with an opportunity moves from region 2 to region 1 if and only if $\phi \leq \Lambda_{t,t+1} D_{t+1}$. Therefore, the flow from 2 to 1 is,

$$N_{1,t+1} - N_{1t} = \lambda N_{2t} \frac{\max\{\min\{\Lambda_{t,t+1} D_{t+1}, \bar{k}\}, 0\}}{\bar{k}} - \lambda N_{1t} \frac{\max\{\min\{-\Lambda_{t,t+1} D_{t+1}, \bar{k}\}, 0\}}{\bar{k}}. \tag{91}$$

The value gap satisfies

$$D_t = (w_{1t} - w_{2t}) + \Lambda_{t,t+1} D_{t+1} + \lambda \left[\Omega(-\Lambda_{t,t+1} D_{t+1}) - \Omega(\Lambda_{t,t+1} D_{t+1}) \right], \tag{92}$$

where $\Omega(x) = \mathbb{E} \max\{0, x - \phi\}$ collects the option value of future opportunities. Expectations are taken over the idiosyncratic fixed cost of moving.

Let $\nu_t \equiv \ln N_{1t} - \ln N_{2t}$, $z_t \equiv \ln Z_{1t} - \ln Z_{2t}$, and $\omega_t \equiv \ln w_{1t} - \ln w_{2t} = z_t - (1-\alpha)\nu_t$ denote log regional gaps for labor, productivity, and wages. The migration decisions linearized around the symmetric steady state are then

$$D_t = \sum_{s=0}^{\infty} \beta^s \bar{w} \omega_{t+s} \tag{93}$$

$$\nu_{t+1} - \nu_t = 4(N_{1,t+1} - N_{1t}) \tag{94}$$

$$N_{1,t+1} - N_{1,t} = \frac{\lambda}{2\bar{k}} \beta D_{t+1} \tag{95}$$

which yields

$$\nu_{t+1} - \nu_t = \psi \sum_{s=1}^{\infty} \beta^s [z_{t+s} - (1-\alpha)\nu_{t+s}], \qquad \psi \equiv \frac{2\lambda\bar{w}}{\bar{k}} \tag{96}$$

Equation (96) captures the entire endogenous dynamics of the economy given a relative productivity path $\{z_{t+s}\}$. In particular, given a permanent shock to z as in much of the empirical work, then relative wage gaps ω_t converge at rate ρ, where $\rho \in (0,1)$ is the stable root of

$$(1-\rho)(1-\beta\rho) = \psi(1-\alpha)\beta\rho. \tag{97}$$

Conversely, a target convergence rate ρ together with λ pins down $\bar{k} = 2\lambda\bar{w}(1-\alpha)\beta\rho/[(1-\rho)(1-\beta\rho)]$.

D.6 Aggregation Bias with Forward-Looking Migration Now we run the aggregation experiment in the model. The recession shock is an aggregate AR(1) productivity shock with persistence ζ and differential regional loadings, $z_{lt} = \theta_l \zeta^t \varepsilon_0$ with $\theta_1 > \theta_2$.

Up to the first order around the symmetric steady state, aggregate output tracks aggregate productivity,

$$ARIRF_h = -\zeta^h \tag{98}$$

To solve for relative regional output, we first solve (96) using the relative productivity

shock $z_t = (\theta_1 - \theta_2)\zeta^t \varepsilon_0$,

$$\nu_t = \frac{\psi\beta\zeta}{(1-\beta\zeta)(\zeta-1)+\psi\beta\zeta(1-\alpha)}(\theta_1-\theta_2)(\zeta^t-\rho^t)\varepsilon_0. \tag{99}$$

Since the per capita income gap is $z_t - (1-\alpha)\nu_t$, the relative recession IRF is,

$$RRIRF_h = -\zeta^h + \frac{\psi\beta\zeta(1-\alpha)}{(1-\beta\zeta)(\zeta-1)+\psi\beta\zeta(1-\alpha)}(\zeta^h-\rho^h) \tag{100}$$

Therefore the aggregation bias is,

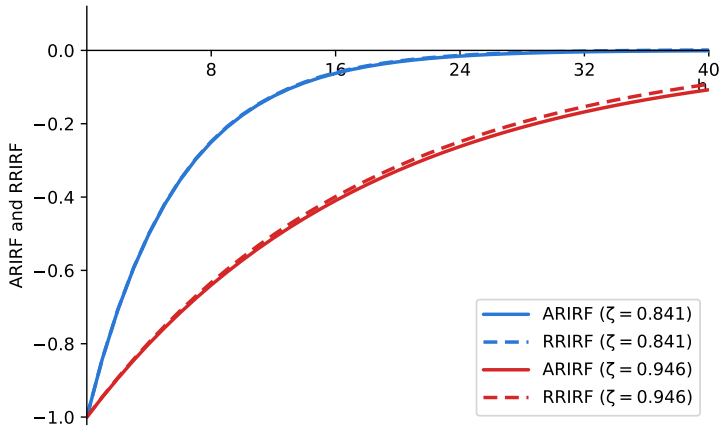
$$RRIRF_h - ARIRF_h = \frac{\psi\beta\zeta(1-\alpha)}{(1-\beta\zeta)(\zeta-1)+\psi\beta\zeta(1-\alpha)}(\zeta^h-\rho^h) \tag{101}$$

We emphasize two main points from this result. First, for $\zeta = 1$ the aggregation bias is identical to the error correction model (83) even though in this model we microfound regional convergence with forward-looking migration and the ECM is a purely statistical model. This is because when shocks are permanent, the income gap—the error correction term—is the only state variable of the model.

Second, for transitory business cycle shocks the aggregation bias is negligible. The ARIRF and RRIRF for the model with endogenous migration in Figure 13 are almost indistinguishable. The bias is smaller compared to the ECM in Figure 11 because the convergence forces are weaker with transitory shock. Intuitively, for a transitory shock there is a greater incentive to wait rather than pay the fixed costs twice (once migrating out and then back in).

D.7 Conclusion For realistic convergence forces across regions, the aggregation bias is negligible and we can treat the RRIRF as approximately equal to the ARIRF.

Figure 13: Aggregation Bias in Recessions with Fast and Slow Recoveries with Endogenous Migration



E Trend Measurement Error: Monte Carlo

We next show that trend measurement error induces only a small bias in what we consider an adversarial DGP. In particular, we allow for very volatile permanent shocks to industry output growth with standard deviation of 1.3 p.p. per quarter. Fitting a linear trend as we do in the data does not capture these trend shocks, and the trend measurement error is systematically correlated with the true trend. However, because we average over many industries in the construction of the Bartik instrument, the bias from this idiosyncratic measurement error in our RRIRF estimates is small. In addition, the Monte Carlo simulation shows that any other biases from our procedure, such as finite-sample bias in fitting the linear trend, do not meaningfully impact our estimated RRIRF.

E.1 Data Generating Process There are $i = 1, \ldots, I$ regions and $g = 1, \ldots, G$ industries. The recession peak is date $t = 0$ and the trough is date $t = 1$. Regression horizons are $h \in \{-8, \ldots, 16\}$, and horizon h corresponds to calendar date $t(h) = 1 + h$. The linear trend window is $t \in [-60, 0]$ and the pre-trend control window is $t \in [-60, -20]$ (i.e. fifteen to five years before the peak).

Each region has industry shares $\mu_{ig} \geq 0$ with $\sum_g \mu_{ig} = 1$, generated by normalizing independent Gamma draws,

$$\tilde{\mu}_{ig} \sim \text{Gamma}(\kappa, 1), \qquad \mu_{ig} = \frac{\tilde{\mu}_{ig}}{\sum_{g'} \tilde{\mu}_{ig'}}.$$

The concentration κ is chosen to match a target average Herfindahl–Hirschman index of $\mathrm{HHI}^\star = 0.0934$, which corresponds to the average regional HHI at the peak of the Great Recession:

$$\mathbb{E}\Big[\sum_g \mu_{ig}^2\Big] = \frac{\kappa+1}{G\kappa+1} \quad \Longrightarrow \quad \kappa = \frac{1-\mathrm{HHI}^\star}{G\,\mathrm{HHI}^\star - 1} = 1.834 \quad (G=16).$$

Each industry receives a positive recession loading from a log-normal family, $b_g = \exp(\sigma_b z_g)$, $z_g \sim \mathcal{N}(0,1)$, normalized to a cycle profile with unit average aggregate exposure,

$$\frac{1}{\sigma_g} = \frac{b_g}{\sum_g \bar{\mu}_g b_g}, \qquad \bar{\mu}_g = \frac{1}{I}\sum_i \mu_{ig}, \qquad x_i = \sum_g \mu_{ig}\frac{1}{\sigma_g}.$$

The dispersion σ_b is solved by bisection in every replication so that the Bartik instrument has cross-sectional standard deviation $\mathrm{sd}_i(R\,x_i) = 0.0082$ for the current recession draw R. x_i denotes regional exposure to the aggregate business cycle.

The aggregate recession size is a truncated-normal draw $R \sim \mathcal{N}(0.071, 0.0255^2)$, and redrawn until $R \geq 0.02$. The industry cyclical level is $\hat{y}_{g,t} = -R\frac{1}{\sigma_g}m(t)$ with case-specific path $m(t)$:

$$m_{\mathrm{AR1}}(t) = \begin{cases} 0, & t \leq 0, \\ \rho^{t-1}, & t \geq 1, \end{cases} \qquad m_{\mathrm{H}}(t) = \begin{cases} 0, & t \leq 0, \\ \rho^{t-1} + \xi\dfrac{1-\rho^{t-1}}{1-\rho}, & t \geq 1, \end{cases}$$

and, in the boom–bust case with boom length T_B,

$$m_{\mathrm{BB}}(t) = \begin{cases} \min\big(1,\, -t/T_B\big), & t < 0, \\ 0, & t = 0, \\ 1, & t \geq 1, \end{cases}$$

so output rises linearly to the peak over T_B quarters and the bust occurs in one quarter. The implied true coefficients are

$$\beta_h^\star = \begin{cases} 0, & h < 0 \text{ (AR(1) and hysteresis cases)}, \\ -\min\big(1,\, -(1+h)/T_B\big), & h < 0 \text{ (boom–bust)}, \\ -\rho^h, & h \geq 0 \text{ (AR(1))}, \\ -1, & h \geq 0 \text{ (boom–bust)}, \\ -\Big(\rho^h + \xi\frac{1-\rho^h}{1-\rho}\Big), & h \geq 0 \text{ (hysteresis)}. \end{cases}$$

The true industry trend level is

$$\bar{y}_{g,t} = \gamma_g t + \sigma_u u_{g,t} + \sigma_a u_t, \qquad \gamma_g \sim \mathcal{N}(0, \sigma_\gamma^2),$$

where $u_{g,t}$ and u_t are independent random walk processes with standard-normal innovations, scaled by $\sigma_u = 0.013$ and $\sigma_a = 0.013$. The industry value matches the residual standard deviation of industry annual TFP growth after industry and time fixed effects, converted to quarterly frequency from Foerster, Hornstein, Sarte, and Watson (2022). We do not have a corresponding calibration target for the aggregate trend shocks. Since our approach is robust to aggregate trend shocks we simply set the standard deviation to the same value of 1.3% per quarter. Industries additionally carry an idiosyncratic stationary AR(1) level component $\mathring{y}_{g,t}$ with persistence 0.75 and unconditional standard deviation 0.005. Observed industry output is

$$y_{g,t} = \bar{y}_{g,t} + \hat{y}_{g,t} + \mathring{y}_{g,t}.$$

Each region has an idiosyncratic AR(1) level process $\mathring{y}_{i,t}$ with persistence $\rho_\alpha = 0.5$ and level standard deviation calibrated so that a four-quarter sum has standard deviation 0.028 (the residual standard deviation of the first stage for post-1990 recessions), implying $\sigma_\alpha = 0.00822$. With growth measured relative to the peak, $\Delta y_{g,t} = y_{g,t} - y_{g,0}$ and $\Delta \mathring{y}_{i,t} = \mathring{y}_{i,t} - \mathring{y}_{i,0}$, regional output growth is

$$\Delta y_{i,t} = \sum_g \mu_{ig} \, \Delta y_{g,t} + \Delta \mathring{y}_{i,t}.$$

E.2 Trend-Fitting For each industry, the econometrician runs OLS of observed industry output on a constant and a linear time trend over the pre-recession window,

$$y_{g,t} = a_g + \gamma_g t + e_{g,t}, \qquad t = -60, \ldots, 0,$$

and extrapolates the fitted line $\tilde{\bar{y}}_{g,t} = \hat{a}_g + \hat{\gamma}_g t$ to all dates. The measured cyclical component is $\tilde{\hat{y}}_{g,t} = y_{g,t} - \tilde{\bar{y}}_{g,t}$. The implied trend measurement error $\varepsilon_{g,t} = \tilde{\bar{y}}_{g,t} - \bar{y}_{g,t}$ consists of three components:

1. Estimation and approximation error of a straight line fitted to a unit-root trend: the fitted slope picks up the realized drift of $\gamma_g t + \sigma_u u_{g,t} + \sigma_a u_t$ over the window. In addition, the post-peak random-walk innovations at both the industry and aggregate level are unpredictable from any pre-period fit and are thus part of the measurement error.

2. Cyclical contamination: any cyclical variation inside the fit window tilts the industry trend. In the boom-bust scenario the pre-period cyclical component is non-zero within T_B quarters of the peak, which contaminates the fitted trend.

3. Projection of industry idiosyncratic noise $\overset{\circ}{y}_{g,t}$ onto the time trend, which is idiosyncratic.

E.3 Estimated Regressions Regional recession depth, the measured Bartik instrument, the fitted Bartik trend control, and the pre-trend control are

$$D_i = -\Delta y_{i,1}, \qquad Z_i = \sum_g \mu_{ig}\big(-\Delta \tilde{\tilde{y}}^{C}_{g,1}\big), \qquad W_{i,t} = \sum_g \mu_{ig}\,\Delta \tilde{\tilde{y}}_{g,t}, \qquad P_i = y_{i,-20} - y_{i,-60}.$$

For each horizon h (date $t = 1 + h$) the script estimates $\Delta y_{i,t} = \alpha_h + \beta_h D_i + \text{controls} + \varepsilon_{i,h}$ by OLS (no instrument, benchmark) and by IV with fitted Bartik trend control $W_{i,t}$ and pre-trend control P_i.

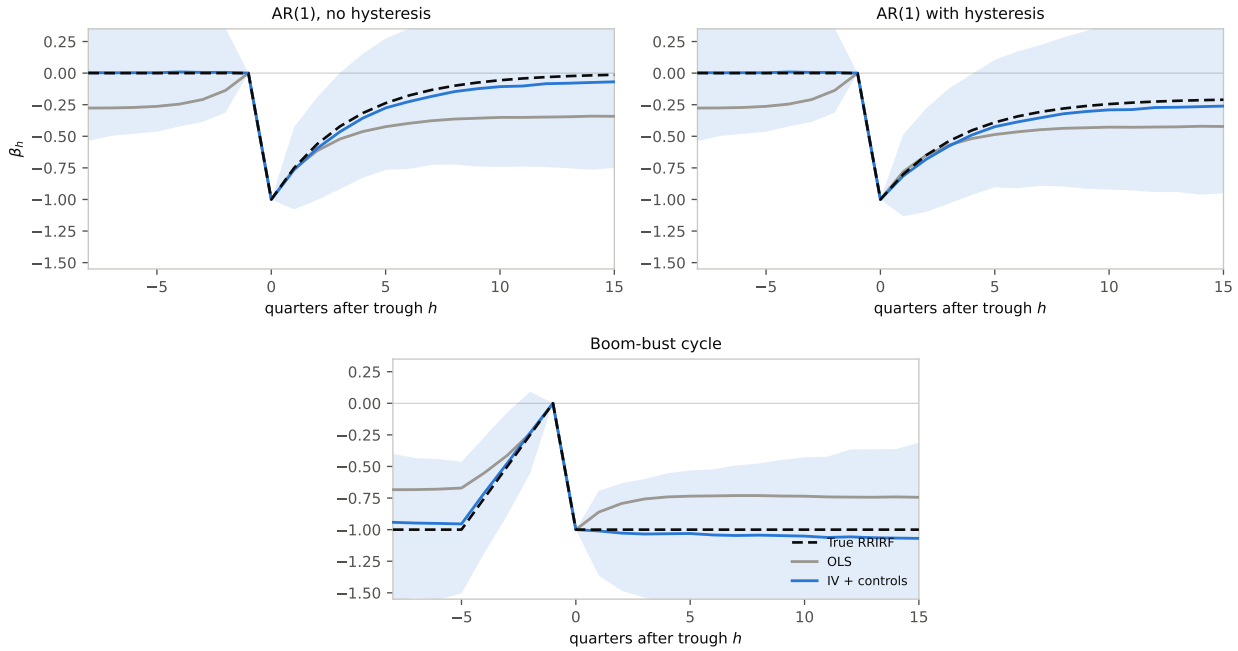
E.4 Calibration Table 2 summarizes the parameter values and their targets.

Table 2: Calibration of DGP.

Object	Symbol	Value	Target / meaning
Regions	I	50	Number of cross-sectional units
Industries	G	16	Number of industries
Replications	–	2000	Monte Carlo draws
Recession mean	$\mathbb{E}[R]$	0.071	Mean aggregate recession size
Recession sd	$\mathrm{sd}(R)$	0.0255	Dispersion of aggregate recession size
AR(1) cycle persistence	ρ	0.75	AR(1) cycle cases
Hysteresis parameter	ξ	0.05	Long-run level effect
Boom length	T_B	4	Quarters of linear boom before the peak
Implied share concentration	κ	1.834	Target average HHI 0.0934
Sd cyclical industry exposure	σ_b	varies	Target Bartik cross-sectional sd 0.0082 in each simulation
Industry trend slope sd	σ_γ	0.003	Dispersion of γ_g
Industry trend shock sd	σ_u	0.013	Residual sd of industry TFP growth (quarterly)
Aggregate trend shock sd	σ_a	0.013	Aggregate unit-root trend component
Industry idiosyncratic sd	–	0.005	Sd of $\overset{\circ}{y}_{g,t}$ innovation
Industry idiosyncratic persistence	–	0.75	Persistence of $\overset{\circ}{y}_{g,t}$
Local noise persistence	ρ_α	0.5	Persistence of $\overset{\circ}{y}_{i,t}$
Implied local level sd	σ_α	0.00822	Target four-quarter $y_{i,t}$ sd 0.028

E.5 Results Figure 14 shows that under all three aggregate recessions, the IV estimator with controls closely tracks the true RRIRF. There is a small bias towards estimating hysteresis because not all of the trend measurement error vanishes with finite number of sectors: this means that a small component of predicted recession depth reflects differential trend shocks. We note that the 90% simulation bands are very wide, which reflects the very high standard deviation of industry trend shocks. The figures also show that OLS is more significantly biased. This is because OLS also loads on the local shocks, which have different dynamics compared to the aggregate recession.

Figure 14: Monte Carlo Simulation of IV Approach with Mismeasured Trends for Three Aggregate Shocks



E.6 Conclusion Our approach performs well even under an adversarial DGP with large trend measurement error within and across industries.

F Trend Measurement Error: Analytics

F.1 Trend Mismeasurement in the Cross-Region Approach In order to construct Bartik predicted recession depth we need to an estimate of the cyclical component of sectoral output.

We obtain the estimated cyclical component of sectoral output $\tilde{\hat{y}}_{g,t}$ from

$$\tilde{\hat{y}}_{g,t} = y_{g,t} - \tilde{\bar{y}}_{g,t} \tag{102}$$

where $\tilde{\bar{y}}_{g,t}$ is an estimate for the trend-component of sectoral output. Assume we model the true sectoral trend as

$$\bar{y}_{g,t} = \gamma_g t + u_{g,t} + u_t \tag{103}$$

where $\gamma_g t$ captures a deterministic trend, $u_{g,t}$ is a sector-specific stochastic component, and u_t is an aggregate trend shock. We assume that both processes are mean-zero, $E(u_{g,t}) = E(u_t) = 0$.

Throughout, we assume that sector-specific shocks are (i) independent across sectors, and (ii) independent from the aggregate shock. Sectoral trend mismeasurement $\varepsilon_{g,t}$ after fitting a linear trend is

$$\varepsilon_{g,t} = (\hat{\gamma}_g - \gamma_g)t + u_{g,t} + u_t \tag{104}$$

where $\hat{\gamma}_g$ is the OLS estimate for the deterministic trend coefficient. As $T \to \infty$, $\hat{\gamma}_g - \gamma_g \xrightarrow{p} 0$ and trend mismeasurement converges to

$$\varepsilon_{g,t} \xrightarrow{p} u_{g,t} + \phi_g u_t + u_t. \tag{105}$$

The estimated sectoral component of output is

$$\tilde{\hat{y}}_{g,t} = \hat{y}_{g,t} + \varepsilon_{g,t}. \tag{106}$$

Then, the Bartik instrument $\tilde{\hat{y}}^b_{i,P,T}$ can be constructed as

$$\tilde{\hat{y}}^b_{i,P,T} = \underbrace{\sum_g s_{i,g}(\hat{y}_{g,P} - \hat{y}_{g,T})}_{\tilde{\hat{y}}^{bc}_{i,P,T}} + \underbrace{\sum_g s_{i,g}(\varepsilon_{g,P} - \varepsilon_{g,T})}_{\tilde{\hat{y}}^{be}_{i,P,T}}, \tag{107}$$

where $\tilde{\hat{y}}^{bc}_{i,P,T}$ and $\tilde{\hat{y}}^{be}_{i,P,T}$ denote the cycle and measurement error components of the Bartik instrument. The IV estimator of the *relative recession impulse response function* is

$$\hat{\beta}^{IV}_h = \frac{Cov_i\left(\tilde{\hat{y}}^b_{i,P,T}, y_{i,T+h} - y_{i,P}\right)}{Cov_i\left(\tilde{\hat{y}}^b_{i,P,T}, y_{i,P} - y_{i,T}\right)}. \tag{108}$$

Next, we show how different forms of measurement error in sectoral trend affect our estimates. Throughout we use the notation $x_{i,t,s}$ to denote the change in x_i between periods t and s.

First, it is straightforward to show that any measurement error that is constant across sectors is differenced-out in the cross-section:

$$\hat{\beta}_h^{IV} = \frac{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,T+h,P}\right) + Cov_i\left(\tilde{\hat{y}}^{be}_{i,P,T}, \sum_g s_{i,g} u_{T+h,P}\right)}{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,P,T}\right) + Cov_i\left(\tilde{\hat{y}}^{be}_{i,P,T}, \sum_g s_{i,g} u_{P,T}\right)} \tag{109}$$

$$= \frac{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,T+h,P}\right) + u_{T+h,P} Cov_i\left(\tilde{\hat{y}}^{be}_{i,P,T}, 1\right)}{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,P,T}\right) + u_{P,T} Cov_i\left(\tilde{\hat{y}}^{be}_{i,P,T}, 1\right)} \tag{110}$$

$$= \frac{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,T+h,P}\right)}{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,P,T}\right)} \tag{111}$$

Using the definition of the RRIRF together with the homogeneity assumption implies $\hat{\beta}_h^{IV} \xrightarrow{p} RRIRF_h$. Second, classical measurement error that is uncorrelated with the non-cyclical components of local output proportionally scales both the first stage and reduced form coefficients, leaving the 2SLS estimator unchanged. A similar argument can be used to show that correlation between sectoral measurement error and sectoral cyclical recession depth does not represent an identification threat. Suppose that the sectoral measurement error satisfies

$$\varepsilon_{g,t} = \kappa_g \hat{y}_{g,t} + \tilde{\varepsilon}_{g,t} \tag{112}$$

where $\tilde{\varepsilon}_{g,t}$ represents an idiosyncratic component of sectoral measurement error that is uncorrelated with the non-cyclical component of local output. We can re-write the Bartik instrument as

$$\tilde{\hat{y}}^{b}_{i,P,T} = \underbrace{\sum_g s_{i,g}(\hat{y}_{g,P,T} + \kappa_g \hat{y}_{g,P,T})}_{\tilde{\hat{y}}^{bc}_{i,P,T}} + \sum_g s_{i,g} \tilde{\varepsilon}_{g,P,T} \tag{113}$$

Plugging this expression in the IV estimator yields:

$$\hat{\beta}_h^{IV} = \frac{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, y_{i,T+h,P}\right)}{Cov_i\left(\tilde{\hat{y}}^{bc}_{i,P,T}, y_{i,P,T}\right)}. \tag{114}$$

Using the definition of the RRIRF together with the homogeneity assumption implies $\hat{\beta}_h^{IV} \xrightarrow{p} RRIRF_h$. Lastly, we discuss non-classical measurement error between the true sectoral trend and sectoral measurement error. Throughout we maintain the assumption that sectoral

measurement error is idiosyncratic across sectors. Using the decomposition of the Bartik instrument we have

$$\hat{\beta}_h^{IV} = \frac{Cov_i\Big(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,T+h,P}\Big) + Cov_i\Big(\tilde{\hat{y}}^{be}_{i,P,T}, \sum_g s_{i,g}\bar{y}_{g,T+h,P}\Big)}{Cov_i\Big(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,P,T}\Big) + Cov_i\Big(\tilde{\hat{y}}^{be}_{i,P,T}, \sum_g s_{i,g}\bar{y}_{g,P,T}\Big)} \tag{115}$$

$$= \frac{Cov_i\Big(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,T+h,P}\Big) + Cov_i\Big(\sum_g s_{i,g}u_{g,P,T}, \sum_g s_{i,g}u_{g,T+h,P}\Big)}{Cov_i\Big(\tilde{\hat{y}}^{bc}_{i,P,T}, \hat{y}_{i,P,T}\Big) + Cov_i\Big(\sum_g s_{i,g}u_{g,P,T}, \sum_g s_{i,g}u_{g,P,T}\Big)} \tag{116}$$

With idiosyncratic trend shocks across sectors we get:

$$E(C_g) = E\Big(Cov_i\Big(\sum_g s_{i,g}u_{g,P,T}, \sum_g s_{i,g}u_{g,T+h,P}\Big)\Big) \tag{117}$$

$$= \sum_g E(u_{g,P,T}u_{g,T+h,P})Var_i\Big(s_{i,g}\Big) \tag{118}$$

where the expectation is taking over the distribution of sectoral trend shocks. Then, if

$$\sup_g |E(u_{g,P,T}u_{g,T+h,P})| \leq K < \infty. \tag{119}$$

We get

$$|E(C_G)| \leq K\sum_g \text{Var}_i(s_{ig}) \tag{120}$$

$$\leq K\sum_g E(s_{ig}^2) \tag{121}$$

$$= KE_i\Big(\sum_g s_{ig}^2\Big). \tag{122}$$

Then, if

$$E_i\Big(\sum_g s_{ig}^2\Big) \xrightarrow{p} 0, \tag{123}$$

as $G \to \infty$, it follows that $E(C_g) \xrightarrow{p} 0$ and

$$E(\hat{\beta}_h^{IV}) \xrightarrow{p} RRIRF_h. \tag{124}$$

Condition (123) can be interpreted as a requirement on regional diversification. If regions are sufficiently diversified across sectors, then mismeasurement in sectoral trend that is correlated with sectoral trend shocks washes-out in the cross-section.

F.2 Trend Mismeasurement in a Time-Series Approach We observe J recession episodes indexed by j, each with peak P_j and trough T_j. Consider a local projection for each horizon h:

$$\underbrace{y_{T_j+h} - y_{P_j}}_{\Delta^{T_j+h-P_j} y_{T_j+h}} = \alpha + \beta_h^{ts}(\hat{y}_{P_j} - \hat{y}_{T_j}) + \varepsilon_j$$

where $\hat{y}_{P_j} - \hat{y}_{T_j}$ measures the national recession depth relative to trend. $\{\beta_h^{ts}\}_h$ is the aggregate recession impulse response function and measures the behavior of aggregate output to a 1% deeper recession. We assume that each recession episode j is driven by an i.i.d shock $\beta_{0,j} \sim (0, \sigma_\beta^2)$. This generates the cross-episode variation in recession depth required to identify $\{\beta_h^{ts}\}_h$ using a time-series approach.

Constructing recession depth relative to trend requires an estimate of the aggregate trend. The estimated trend component of aggregate output is

$$\bar{y}_t^e = \bar{y}_t + \bar{e}_t \tag{125}$$

where $\bar{y}_t$ is the true trend component, $\bar{e}_t$ is measurement error and the superscript e is used to denote *estimated* components of output. The estimated cyclical component is then given by

$$\hat{y}_t^e = y_t - \bar{y}_t^e \tag{126}$$

$$= \hat{y}_t - \bar{e}_t \tag{127}$$

where $\hat{y}_t$ is the true cyclical component. The estimated h-period difference in the cyclical component is

$$\Delta^h \hat{y}_{t+h}^e = \hat{y}_{t+h}^e - \hat{y}_t^e = \Delta^h \hat{y}_{t+h} - \Delta^h \bar{e}_{t+h} \tag{128}$$

The OLS estimator of the aggregate recession impulse response function converges to

$$\hat{\beta}_h^{ts} = \frac{Cov_j(y_{T_j+h} - y_{P_j}, \hat{y}_{P_j}^e - \hat{y}_{T_j}^e)}{Var_j(\hat{y}_{P_j}^e - \hat{y}_{T_j}^e)} \tag{129}$$

$$= -\frac{Cov_j(\Delta^{T_j+h-P_j}\hat{y}_{T_j+h} + \Delta^{T_j+h-P_j}\bar{y}_{T_j+h}, \Delta^{T_j-P_j}\hat{y}_{T_j} - \Delta^{T_j-P_j}\bar{e}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}_{T_j}^e)} \tag{130}$$

Further assuming that the trend and cyclical component of aggregate output are orthogonal

to each other we get:

$$\hat{\beta}_h^{ts} = -\frac{Cov_j(\Delta^{T_j+h-P_j}\hat{y}_{T_j+h}, \Delta^{T_j-P_j}\hat{y}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} + \frac{Cov_j(\Delta^{T_j+h-P_j}y_{T_j+h}, \Delta^{T_j-P_j}\bar{e}_{T_j}\Big)}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} \tag{131}$$

Proposition 2

$$\beta_h^{ts} \to -\frac{Cov_j(\Delta^{T_j+h-P_j}\hat{y}_{T_j+h}, \Delta^{T_j-P_j}\hat{y}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} = \beta_h^{ts} \times \mathcal{A} \tag{132}$$

where $\mathcal{A} = \frac{Var_j(\Delta^{P_j-T_J}\hat{y}_{P_j})}{Var_j(\Delta^{P_j-T_J}\hat{y}^e_{P_j}))}$ *measures the degree of attenuation bias.*

The assumption that $Cov_j(\Delta^{T_j+h-P_j}y_{T_j+h}, \Delta^{T_j-P_j}\bar{e}_{T_j}\Big) = 0$ is strong. Even if the estimated trend $\bar{y}^e_t$ is orthogonal to the cycle, the covariance remains non-zero whenever the aggregate trend has a stochastic component. This is because trend growth during the recession window $[P_j, T_j]$ and trend growth over the recovery horizon $[P_j, T_j + h]$ share a common component, generating a non-zero covariance across episodes j. We show this formally next when the trend has a deterministic and stochastic component.

Let the aggregate trend $\bar{y}_t$ evolve according to

$$\bar{y}_t = \gamma t + \upsilon_t \tag{133}$$

where $\upsilon_t = \rho_\upsilon \upsilon_{t-1} + \varepsilon^\upsilon_t$ and $\varepsilon^\upsilon_t \sim N(0, \sigma^2_{\varepsilon^\upsilon})$. The OLS estimator converges to

$$\bar{y}^e_t \to \hat{\gamma} t, \tag{134}$$

Measurement error is given by $\bar{e}_t = u_t + \nu_t$ where $\nu_t = (\hat{\gamma} - \gamma)t$ is finite-sample estimation noise (i.e., $\nu_t \to 0$ as $T \to \infty$). Under $\bar{y}^e_t \perp \{\hat{y}_s\}_s$, the covariance between output growth over the recovery and measurement error reduces to

$$Cov_j(\Delta^{D_j+h}\bar{y}_{T_j+h}, \Delta^{D_j}\bar{e}_{T_j}\Big) = Cov_j(\gamma(D_j + h) + \Delta^{D_j+h}\upsilon_{T_j+h}, (\hat{\gamma} - \gamma)(D_j) + \Delta^{D_j}\upsilon_{T_j}\Big) \tag{135}$$

$$= Cov_j(\gamma(D_j + h) + \Delta^{D_j+h}\upsilon_{T_j+h}, (\hat{\gamma} - \gamma)D_j + \Delta^{D_j}\upsilon_{T_j}\Big) \tag{136}$$

$$= \gamma(\hat{\gamma} - \gamma)Var_j(D_j) + Cov_j(\Delta^{D_j+h}\upsilon_{T_j+h}, \Delta^{D_j}\upsilon_{T_j}\Big) \tag{137}$$

where $D_j = T_j - P_j$ is recession duration. The first covariance term collapses to zero as

$T \to \infty$. The second term is the main driver of identification failure. Further noting that

$$\Delta^{D_j+h}\upsilon_{T_j+h} = (\rho_{\upsilon}^{D_j+h} - 1)\upsilon_{P_j} + \sum_{s=1}^{D_j+h} \rho_{\upsilon}^{D_j+h-s}\varepsilon_{P_j+s}^{\upsilon}, \tag{138}$$

$$\Delta^{D_j}\upsilon_{T_j} = (\rho_{\upsilon}^{D_j} - 1)\upsilon_{P_j} + \sum_{s=1}^{D_j} \rho_{\upsilon}^{D_j-s}\varepsilon_{P_j+s}^{\upsilon} \tag{139}$$

It follows that $Cov_j(\Delta^{D_j+h}\upsilon_{T_j+h}, \Delta^{D_j}\upsilon_{T_j}\Big) \neq 0$ whenever $\sigma_{\varepsilon^{\upsilon}}^2 \neq 0$.

G Figures

Figure 15: First Stage Monotonicity of Important Industry Shares

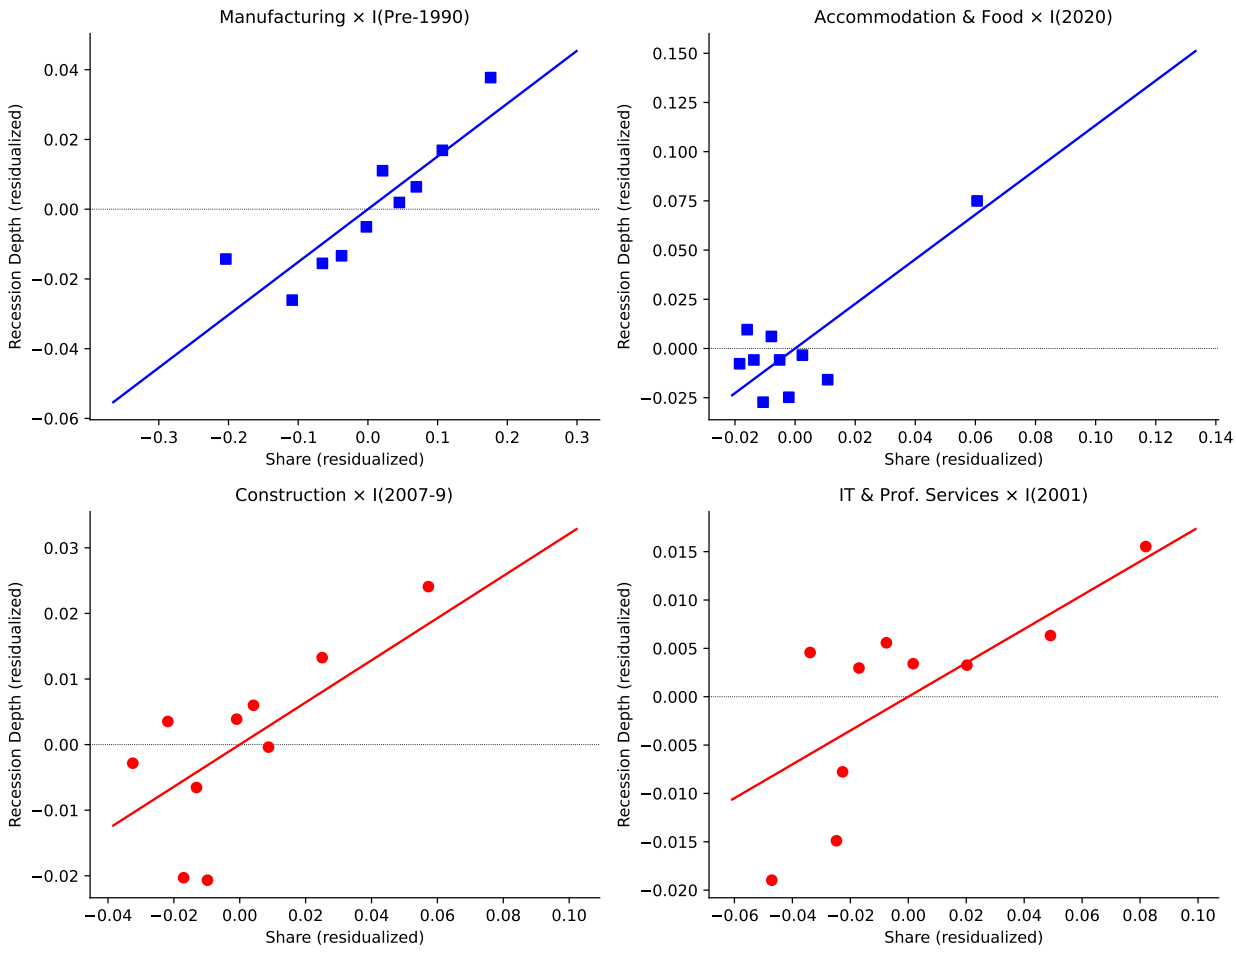


Figure 16: RRIRF of State Personal Income by Recession Group Double-Clustering by State and Recession

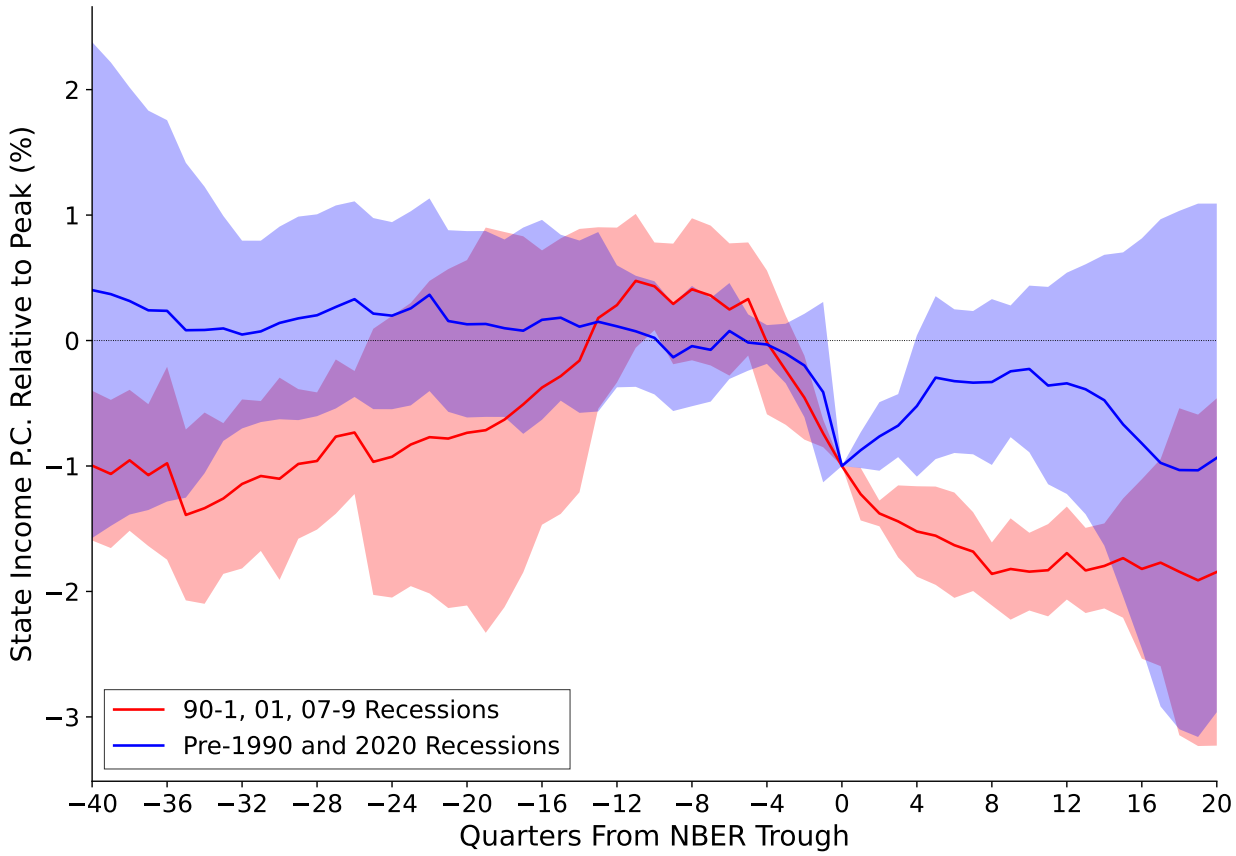


Figure 17: Reduced Form Impulse Response Function for Population

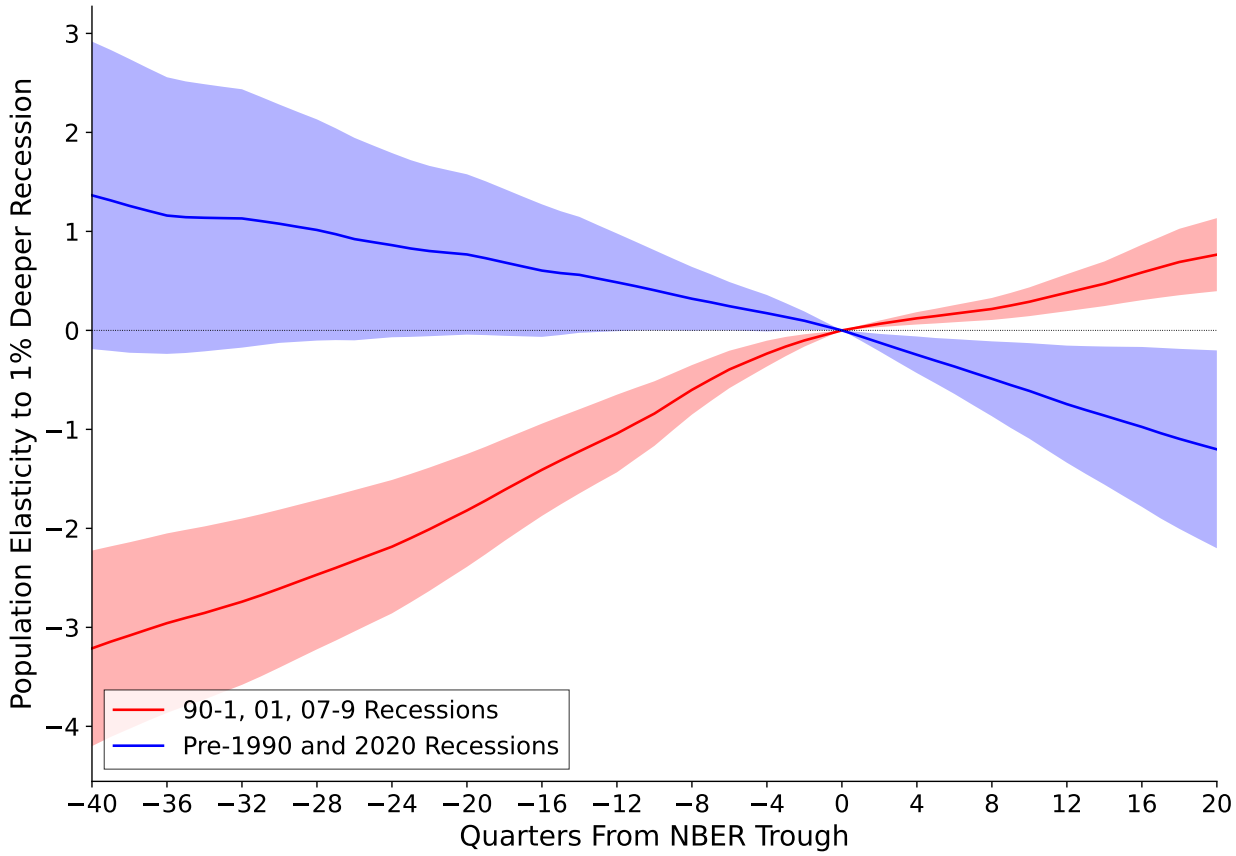


Figure 18: Reduced Form Impulse Response Function for Transfer Share in State Income

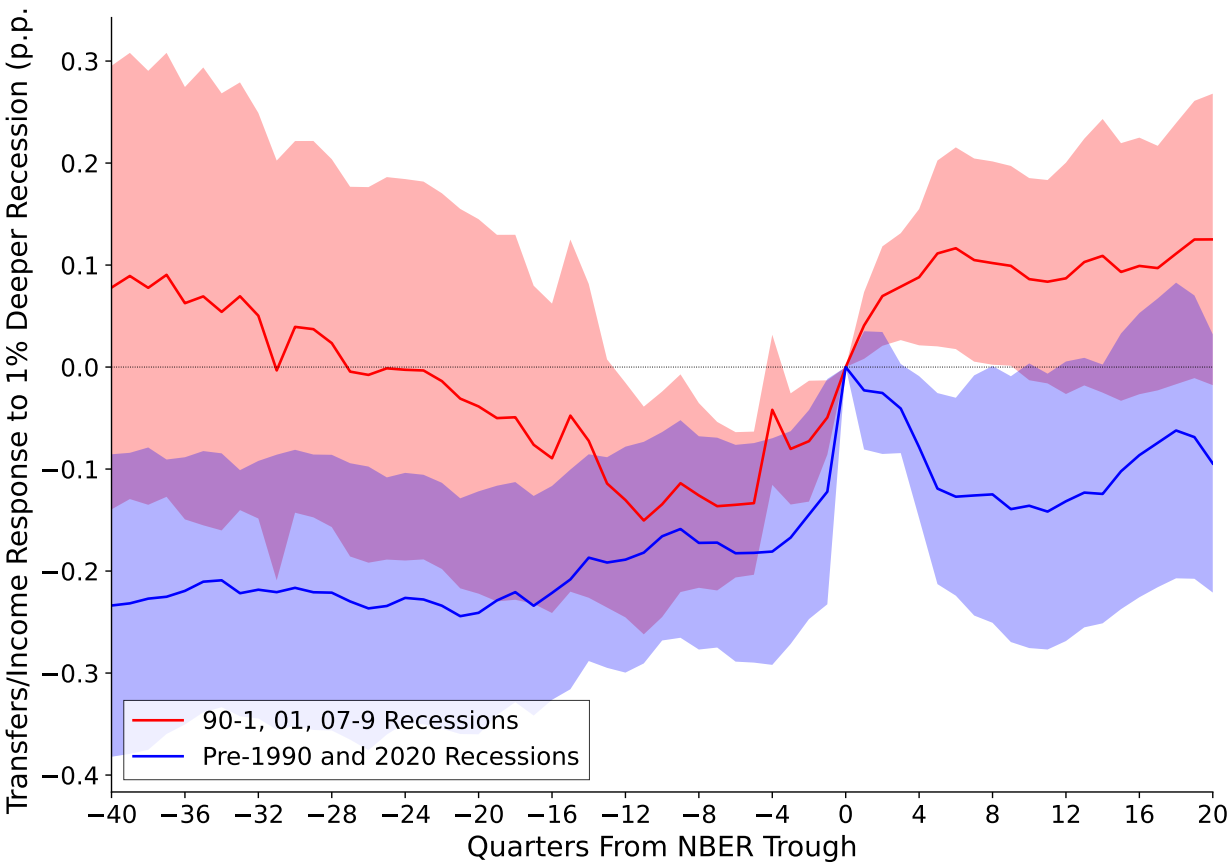
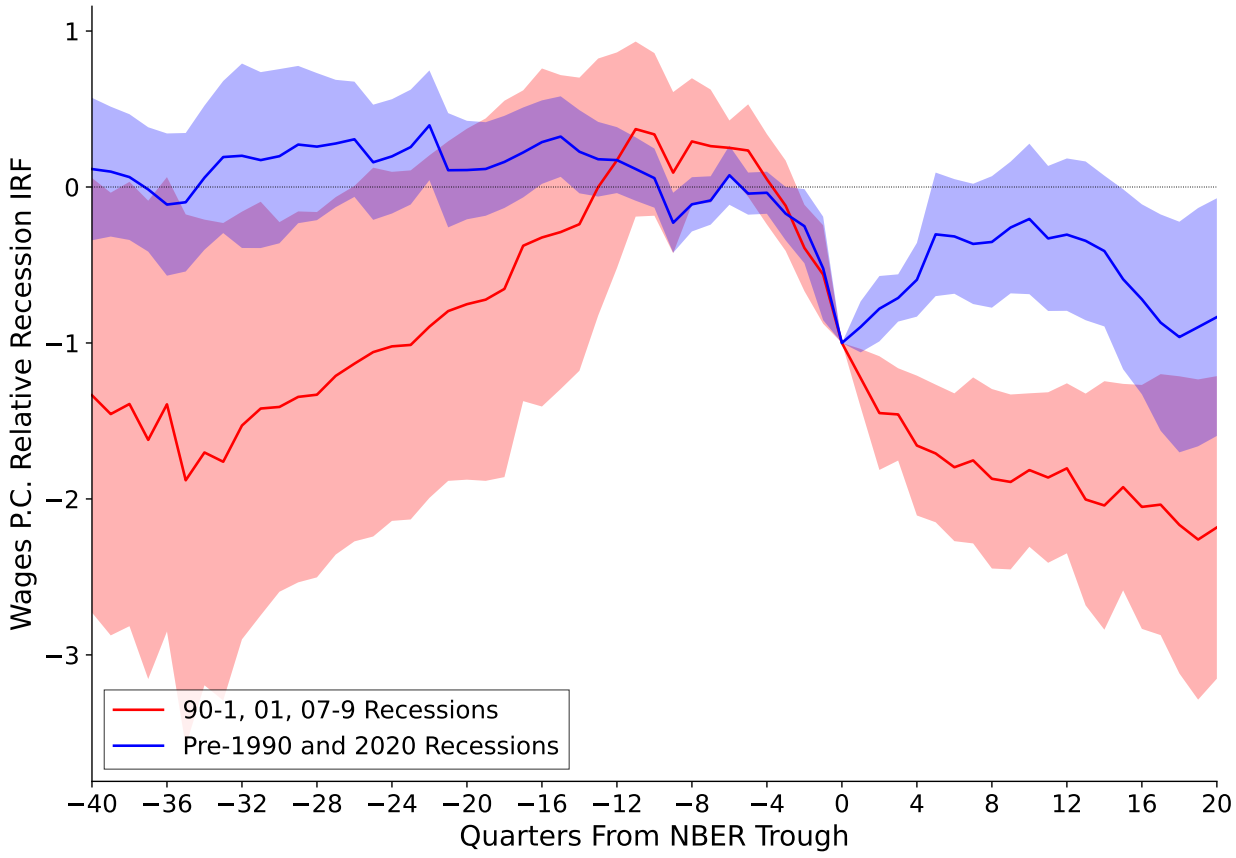


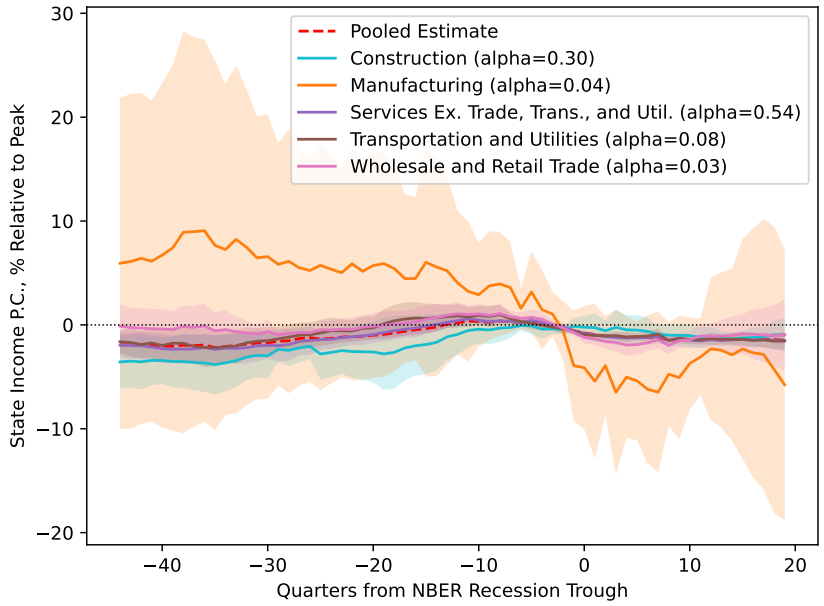
Figure 19: RRIRF for Total Wage Bill



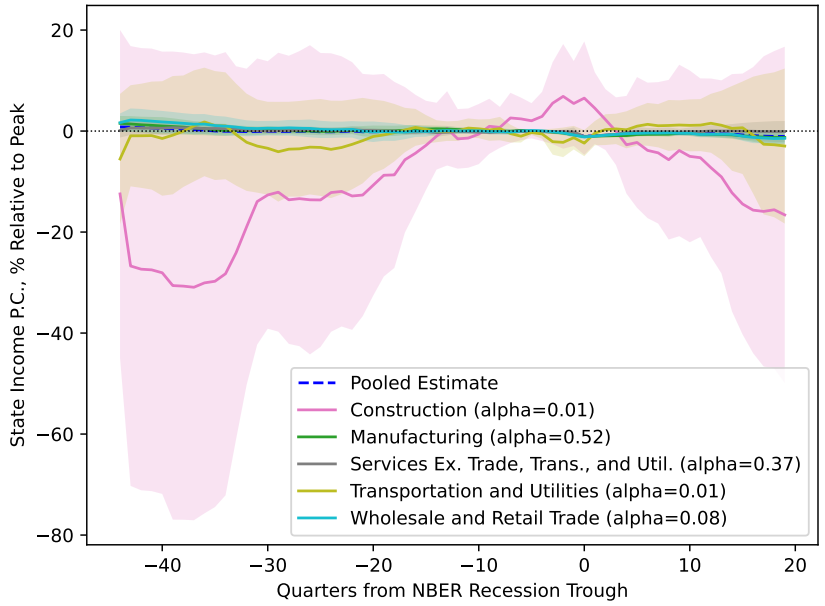
H Rotemberg Weights

Figure 20: Impulse Response Functions based on Major Industry and Recession Variation

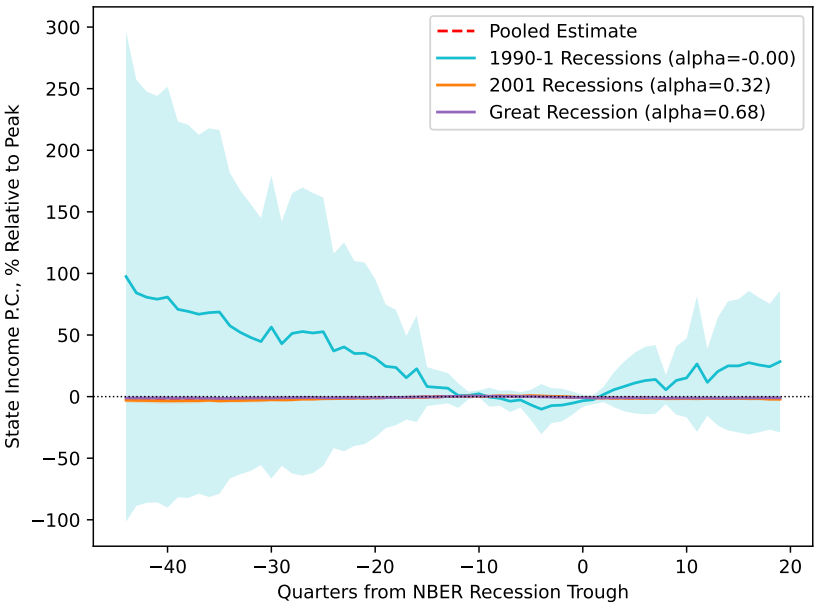
(a) Industry: 1990-1, 2001, 2007-9



(b) Industry: Other Recessions



(c) Recessions: 1990-1, 2001, 2007-9



(d) Recessions: Other Recessions

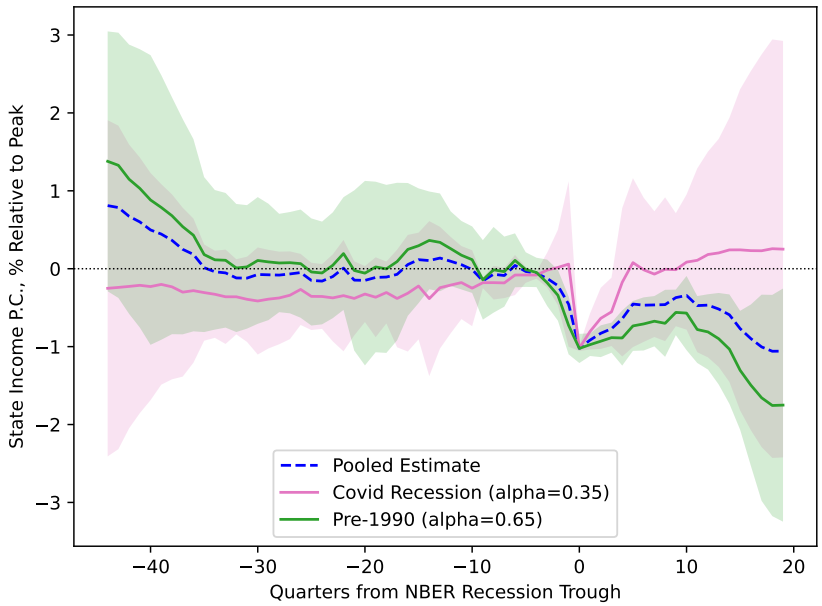
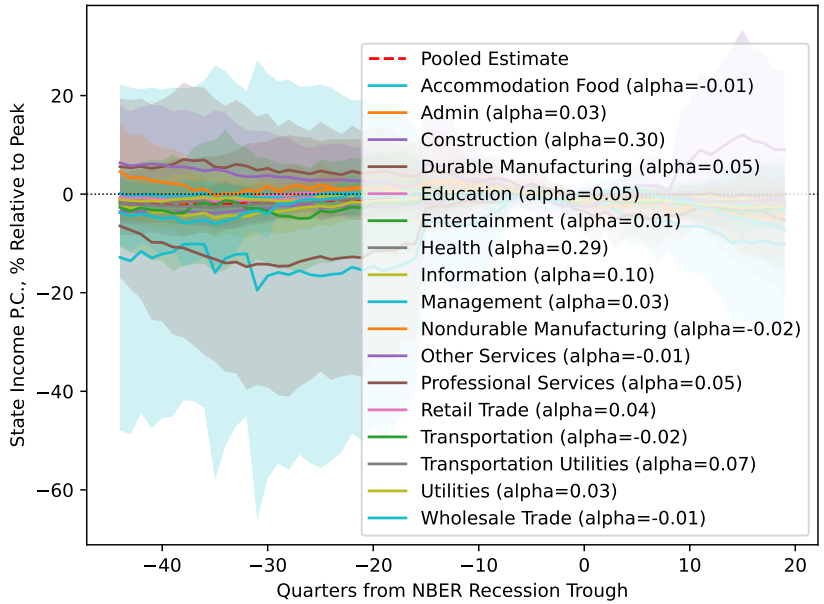
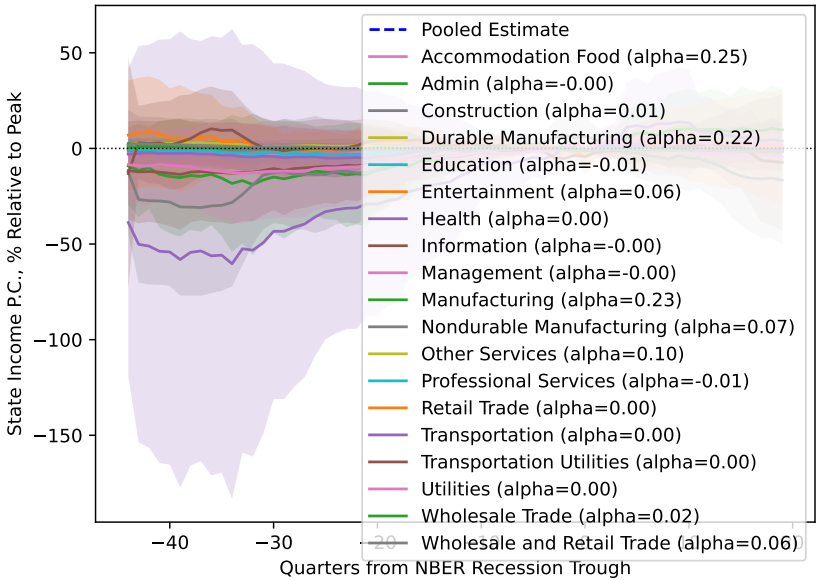


Figure 21: Impulse Response Functions based on Industry and Recession Variation

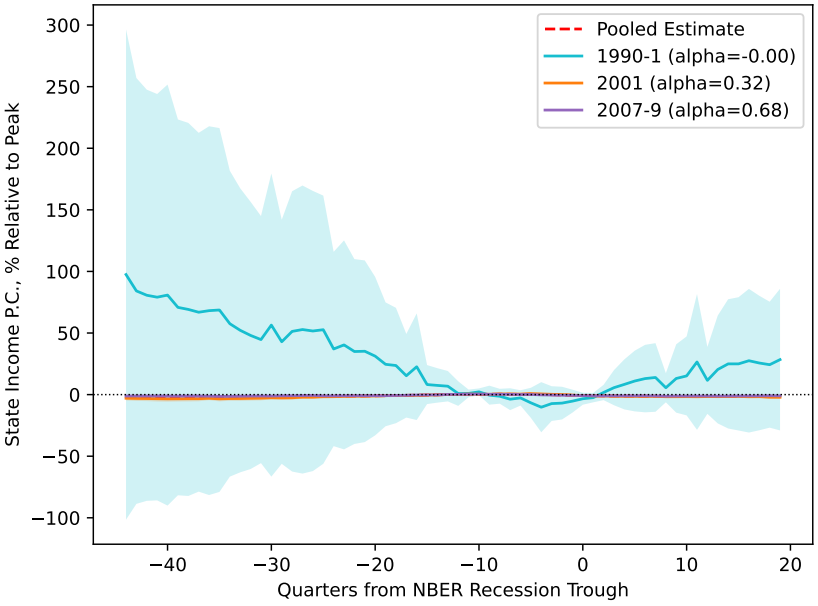
(a) Industry: 1990-1, 2001, 2007-9



(b) Industry: Other Recessions



(c) Recessions: 1990-1, 2001, 2007-9



(d) Recessions: Other Recessions

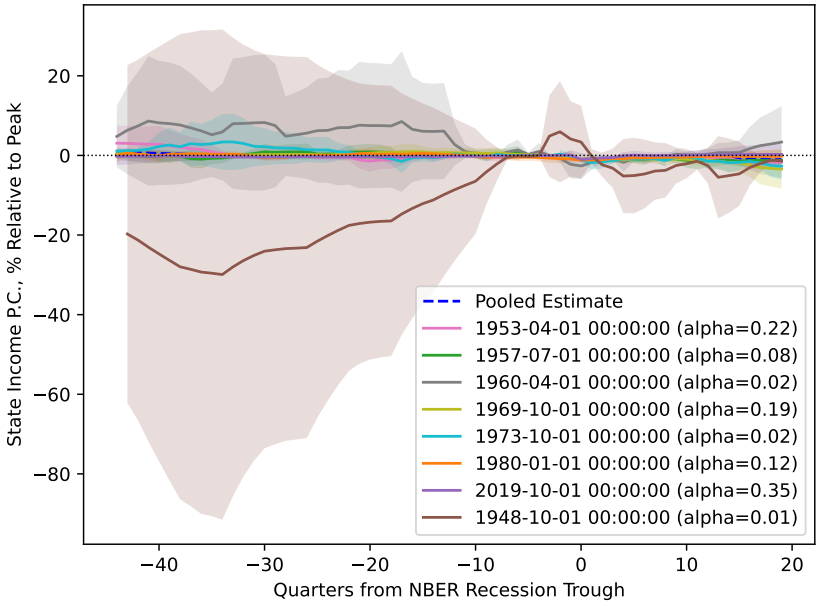
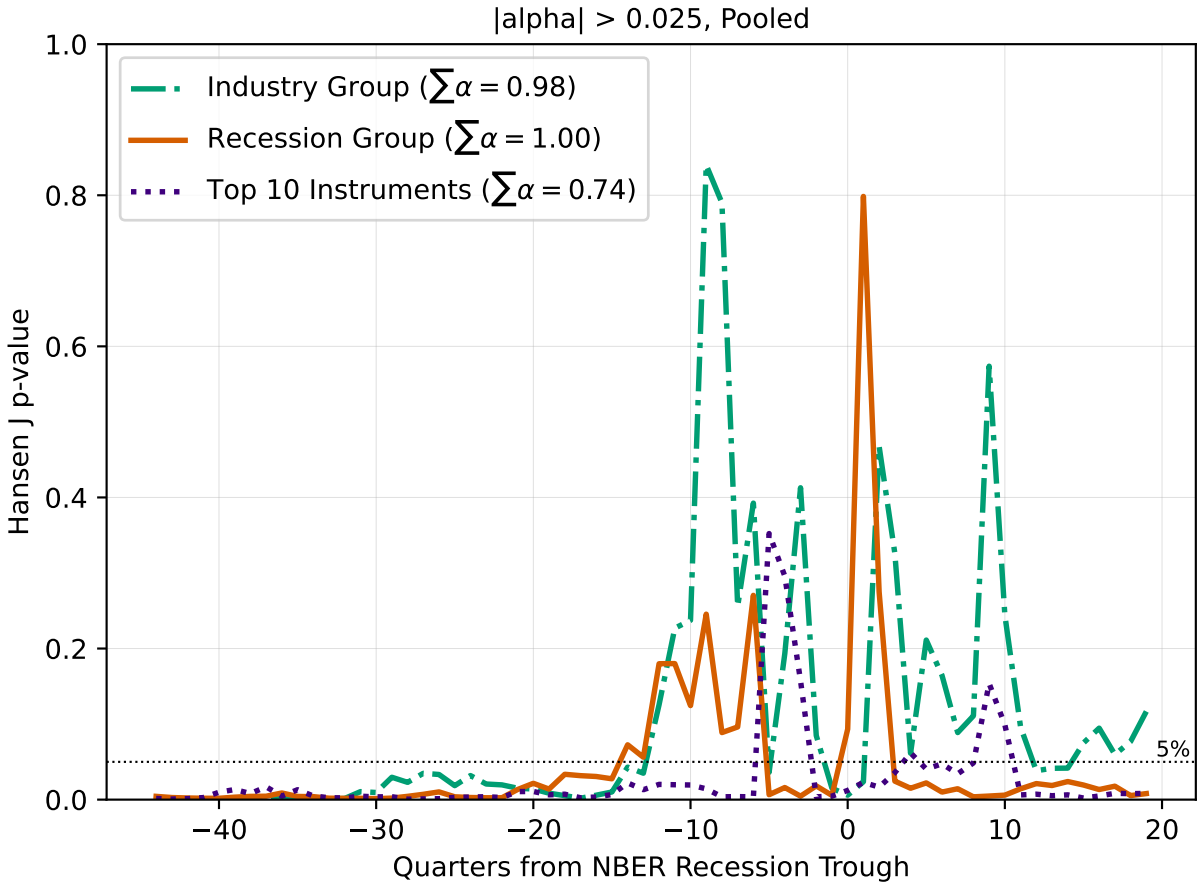


Figure 22: P-values for Overidentification Tests: Pooled Estimates



I Policy Figures

Figure 23: $R*$ Path for V-shaped and Boom-Bust Recession by ρ

