Why Are Some Recoveries Weak and Others Strong?

Paula Donaldson
U.C. San Diego

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

March 27, 2026

Abstract

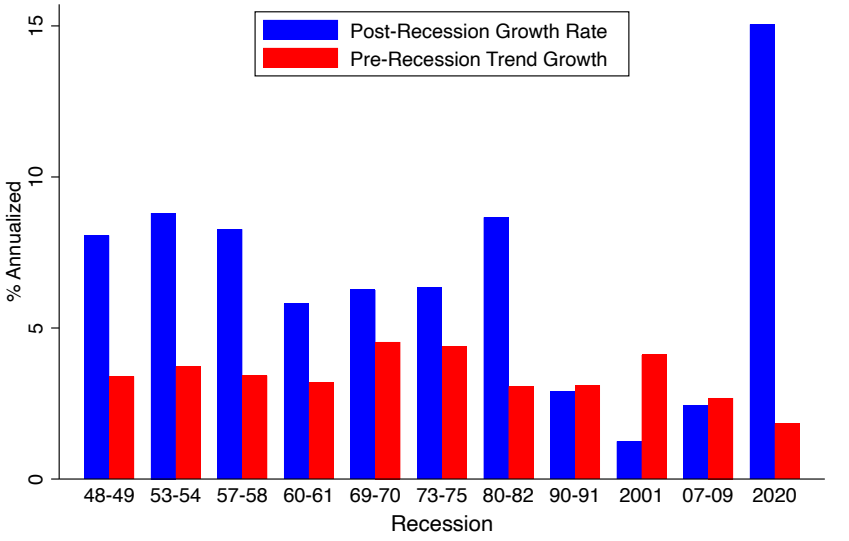
Why were the recoveries from the 1990-1, 2001, and 2007-9 recessions weak relative to other postwar recessions? Leveraging heterogeneous exposure to the national business cycle across U.S. States, we estimate the trajectory of more exposed U.S. States relative to less exposed U.S. States. For the 1990-1, 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. Our estimates support theories that these recessions are caused by different shocks. In a quantitative model matched to our cross-sectional estimates, boom-bust cycles persistently depress the natural rate of interest R^* in the recovery, whereas R^* is elevated following V-shaped recessions.

*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. We are grateful to Andrew Atkeson, Adrien Auclert, Olivier Coibion, James Hamilton, Joshua Hausman, Juan Herreño, and Valerie Ramey for helpful comments.

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 is also weak relative to the pre-recession trend, computed as the annualized growth rate from the last business cycle peak to the current business cycle peak. This contrasts with the strong recoveries from the other recessions, in which recovery growth exceeds the pre-recession trend. This paper asks why are the recoveries from the 1990-1, 2001, and 2007-9 weak—recovery growth is below the pre-recession trend—and why are recoveries from the other recessions 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.

A range of explanations have been put forward to explain why some recoveries are weak: Bad luck (Galí, Smets, and Wouters, 2012), investment boom-bust cycles (Beaudry, Galizia, and Portier, 2018; Rognlie, Shleifer, and Simsek, 2018), credit boom-bust cycles (Jordà,

[1]The sum of consumption expenditures and private investment.

Schularick, and Taylor, 2013), different sectoral incidences of shocks (Beraja and Wolf, 2022), changes in beliefs (Kozlowski, Veldkamp, and Venkateswaran, 2019, 2020), the secular decline in manufacturing (Leamer, 2021), hysteresis (Benigno and Fornaro, 2018; Garga and Singh, 2021; Reifschneider, Wascher, and Wilcox, 2015), and an exogenous slowdown in trend growth (Fernald, Inklaar, and Ruzic, 2025).

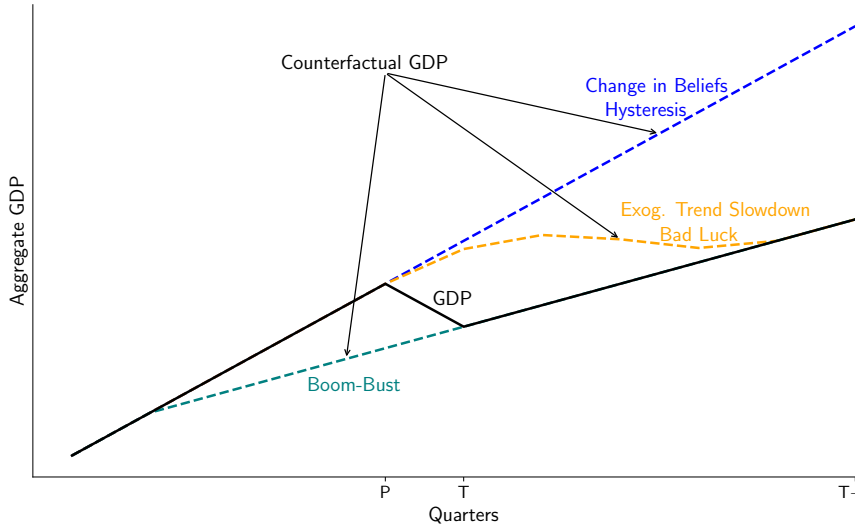
These theories imply different counterfactual paths for output. Figure 2a illustrates these counterfactuals around a hypothetical GDP path similar to the Great Recession. Hysteresis or changes in beliefs due to the recession imply a persistent gap between realized GDP and its counterfactual path absent the recession. An exogenous slowdown in trend growth instead implies that growth would have been weak even absent the recession. We also expect the counterfactual and realized GDP path to converge as bad luck eventually ends. And boom-bust explanations imply that GDP was high before the recession, and through the recession converges to counterfactual GDP.

Distinguishing between these theories is difficult because the counterfactual path of output is unobserved. Common strategies to recover aggregate 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 fixed the country and the recession, but must assume that the trend continues on its existing path. In practice, 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).

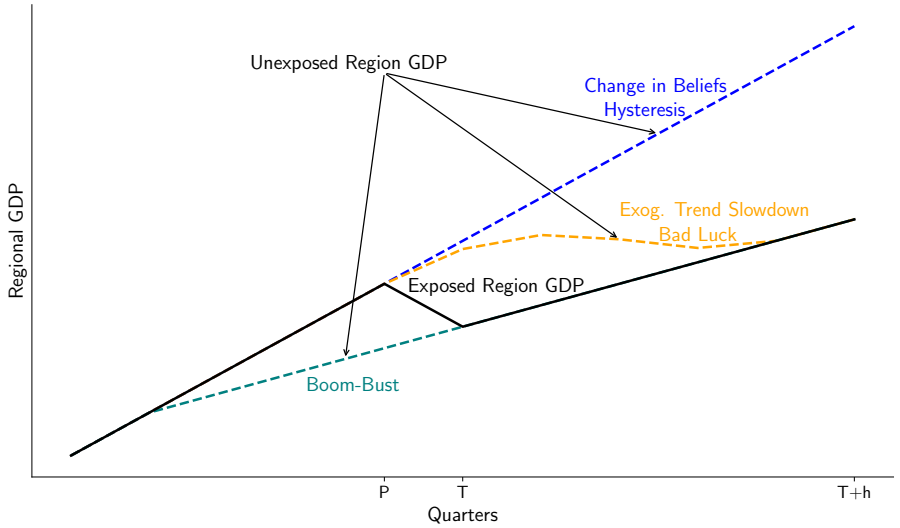
Figure 2b shows that regional data can help uncover the underlying counterfactual and thereby discriminate between theories for weak recoveries. For simplicity, suppose that there are two regions, one is completely exposed and one is completely unexposed to the recession. If the two regions are on parallel trends absent the recession, then the unexposed region provides the correct counterfactual for the exposed region.

Figure 2: Explanations for Weak Recoveries

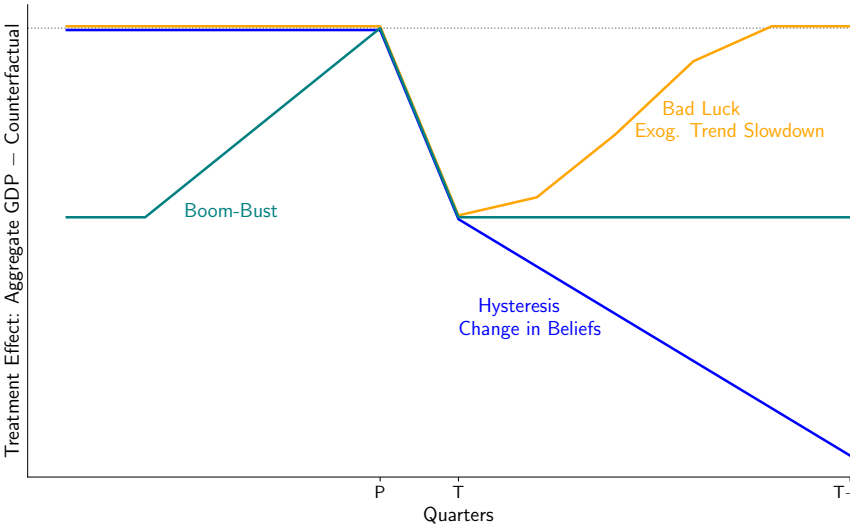
(a) Aggregate Counterfactuals



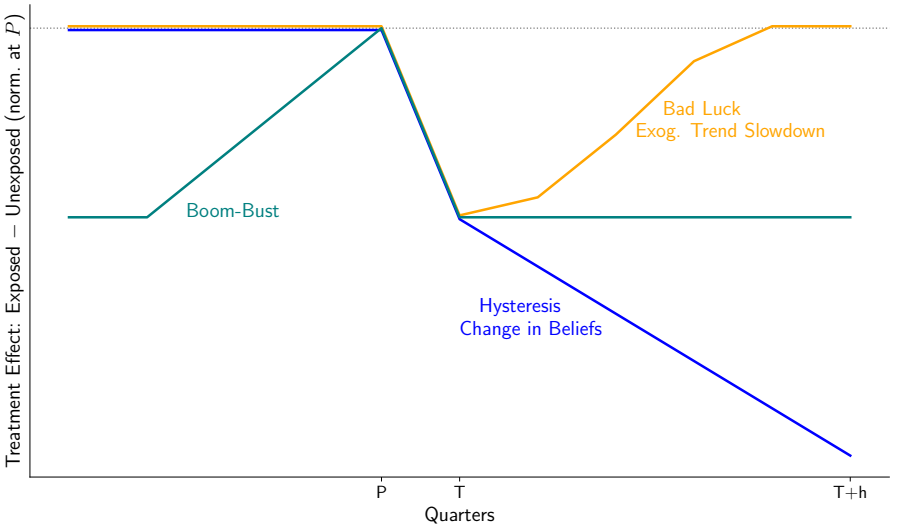
(b) Regional Identification



(c) Aggregate Recession Impulse Response Function



(d) Relative Recession Impulse Response Function



Section 2 formalizes identification. In a simple many-good, many-region new Keynesian model, we show that the aggregate effects of the recession can be recovered up to scale from the *relative recession impulse response function*: 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.

The model shows that the relative recession impulse response function can be estimated with a Bartik IV strategy. In the first stage we predict relative recession depth across regions using a cyclical Bartik that interacts regional industry shares with the cyclical component of national industry growth, and in the second stage we use this instrumented variable to predict relative regional growth before, during, and after the recession. If the cyclical Bartik is uncorrelated with local demand and supply shocks as well as region-specific trends, then this IV strategy recovers the true relative recession impulse response function, and thereby the (scaled) aggregate recession impulse response function.

Figures 2c and 2d illustrate the identification strategy in the two-region example. The aggregate recession impulse response function in Figure 2c is not directly observable and requires modeling a time-series counterfactual. However, the relative recession impulse response function represents the same object and can be feasibly estimated in the data by comparing more and less-exposed regions. In more general settings, general equilibrium forces can scale up or down the regional responses to the aggregate shock, but the aggregate recession impulse response function can still be recovered from the relative recession impulse response function up to scale. As the relative recession impulse response function remains informative about the aggregate recession impulse response, it can help discriminate among theories for weak recoveries.

The cross-region approach holds two important advantages in estimating the recession impulse response function over other methods. First, it is much more robust to trend misspecification than time-series or cross-country panel approaches. Aggregate trends are differenced out, so cross-regional comparisons do not require taking a stance on the stationarity of aggregate GDP, which is notoriously difficult to determine (Stock and Watson, 1999). The construction of the cyclical Bartik still requires a model of industry-level trends, but we show that idiosyncratic industry-level measurement error in trend growth washes out at the regional level under mild assumptions on regional specialization. This contrasts with time-series or cross-country panel estimators that are almost surely inconsistent under any trend mismeasurement.

Second, regional data holds the country and the recession fixed, which facilitates identifying the dynamics of the recession. For instance, a more persistent aggregate demand shock will induce a more persistent recession across more and less exposed regions, and thus poses

no problem for our identification strategy. In a cross-country or in a time-series setting, two recessions may correspond to shocks with different dynamics. 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 scaled "treatment effect" of -1% in the first year and -1% in the second year, and thus greater persistence than implied by either recession.

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 above 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 recession. The different recession-recovery patterns suggest that the relative recession impulse response function is different across recessions with weak and strong national recoveries.

We estimate relative recession impulse response functions 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. Our estimates are consistent with a complete relative recovery, particularly once accounting for more exposed States growing relatively slowly prior to the recession. This reflects that pre-1990 recession exposure largely correlates with the manufacturing share in State income, and manufacturing is a secularly declining sector.

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. In the national recovery phase, the more exposed States do not grow more despite their deeper local recession. However, they do grow more quickly prior to the recession. Thus, the relative recession im-

pulse response function displays an inverse V-shape consistent with boom-bust explanations for weak recoveries (Beaudry, Galizia, and Portier, 2018; Jordà, Schularick, and Taylor, 2013; Rognlie, Shleifer, and Simsek, 2018). Going back prior to the pre-recession boom also reveals that the more exposed States were originally on a slower (per capita) growth trajectory compared to the less exposed States. The weak recovery we estimate is consistent with a return to this weaker relative trend, though we cannot rule out some hysteresis effects.

Our estimates for strong national recoveries support plucking models of the business cycle (Friedman, 1993; Dupraz, Nakamura, and Steinsson, 2023), 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 for local income and GDP is consistent with an "upward" pluck, in which the recession is a drop from an unsustainable boom. However, the boom has a relatively minor effect on the pre-recession unemployment rate, which suggests that it is challenging to infer a boom-bust cycle from the unemployment rate alone.

Following Goldsmith-Pinkham, Sorkin, and Swift (2020) we decompose our relative recession impulse response functions into the relative recession impulse response functions implied by exposure to a particular sector and the associated sectoral weights. This allows us to test whether changes in the relative recession impulse response function across recessions are due to changes in recession trajectories from particular sectoral exposures or changes in sectoral exposure weights. Across different sectoral exposures, we find evidence of V-shaped relative recessions in more exposed States for the pre-1990 recessions and COVID-19, and boom-bust cycles for the 1990-1, 2001, and Great Recession. 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.

We then show that distinguishing between the two types of recessions that we identify matters for policy. We first show that a two-region New Keynesian model with a durable good can match the different recession patterns in the data. Regions that produce relatively more of the durable goods are more exposed to the aggregate business cycle, which allows us to compute relative recession impulse response functions in the same way as in the data. Boom-bust shocks predict a weak recovery in the model because the output in more exposed regions is inflated at the onset of the recession and will diverge from output in less exposed regions. Shocks that induce intertemporal trade-offs, such as monetary policy or a temporary pandemic, instead predict a strong recovery as the recession is simply a temporary deviation from the balanced growth path.

We use the model to calculate the long-run path of the natural rate of interest R^*—the real interest rate that closes the output gap. The model predicts that boom-bust recessions are followed by a prolonged slump in R^*, whereas intertemporal shocks raise R^* in the

recovery. This follows from the different dynamics of the capital stock: in boom-bust cycles the capital stock is too high and investment too low, so aggregate demand must be supported by a period of prolonged low real interest rates. However, following intertemporal recessions, the capital stock is too low as investment was delayed, which implies that the recovery is supported by pent-up demand. This shows that appropriately calibrating monetary policy in the recovery depends importantly on a correct diagnosis of the recession.

2 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. The main result is Proposition 1, which sows that regional variation in business cycle exposure can be exploited to recover the dynamics of the aggregate business cycle shock, and thereby discriminate among theories for weak aggregate recoveries. The model also illustrates key benefits of using regional data for identifying the dynamics of the business cycle shock: We impose only weak restrictions on the relative mismeasurement of cyclical trends, whereas even iid mismeasurement of aggregate trends will induce bias in other approaches. Furthermore, cross-regional variation by definition loads on the common aggregate recession, whereas approaches that leverage cross-recession variation over time or across countries are likely biased if recessions are not identical up to scale.

2.1 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 2.8 we discuss an extension of the model with hand-to-mouth consumers.

The first order conditions for the consumption goods are,

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

$$\lambda_t = \beta_t (1 + r_t) \mathbb{E}_t \lambda_{t+1} \tag{2}$$

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

$$\alpha_{g,t} = \alpha_g Z_{g,t}^{\sigma_g - 1}. \tag{3}$$

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. In our quantitative model in Section 6, durability determines cyclical sensitivity without affecting long-run growth.

2.2 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,

$$C_{i,g,t} = \alpha_{igt} \left(\frac{P_{i,g,t}}{P_{g,t}} \right)^{-\eta_g} C_{g,t}. \tag{4}$$

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

$$\alpha_{igt} = \alpha_{ig} e^{\hat{\alpha}_{it}}, \tag{5}$$

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

2.3 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,

$$Z_{g,t} = \theta_g X_{g,t} Z_{g,t-1}^{1-\xi} Y_{g,t-1}^{\xi}, \tag{6}$$

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.

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,

$$Z_t = \prod_{g=1}^{G} Z_{g,t}^{\alpha_g}, \tag{7}$$

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,

$$\frac{P_{i,g,t}}{P_t} = \frac{P_{g,t}}{P_t} = \bar{W} \frac{Z_t}{Z_{g,t}}. \tag{8}$$

In Section 2.8 we discuss an extension in which prices respond endogenously to the local business cycle.

Regional goods market clear when output equals consumption,

$$Y_{i,g,t} = C_{i,g,t} \tag{9}$$

2.4 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 (1)-(9).

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, and the outcomes in the log-linearized equilibrium by $\hat{x}_{igt} = \ln X_{igt} - \ln \bar{X}_{igt}$.

2.5 Aggregate Business Cycle We define the aggregate recession impulse response function as the sum of the deviation from the balanced growth path and any hysteresis effects to the balanced growth path,

$$\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s = \left(\sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g} \right) \left[\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right] \tag{10}$$

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

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

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

$$\Delta^h y_{t+h} = \underbrace{\Delta^h \hat{y}_{t+h} + \xi \sum_{s=0}^{h-1} \hat{y}_{t+s}}_{\text{Aggregate Recession Impulse Response}} + \underbrace{\sum_{s=0}^{h-1} x_{t+s}}_{\text{Trend Growth}}. \tag{11}$$

where Δ^h is an h-period difference.

We next show how we can recover the aggregate recession impulse response function in (11) using regional variation.

2.6 Regional Bartik Structure Combining market clearing (9), regional demand (4), and sticky relative prices (8) 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,

$$\underbrace{\Delta^h y_{i,t+h}}_{\text{Local Output Growth}} = \underbrace{\Delta^h \hat{\alpha}_{i,t+h}}_{\text{Local Demand Shifter}} + \underbrace{\sum_{g=1}^{G} \mu_{i,g,t} \Delta^h y_{g,t+h}}_{\text{Bartik Predicted Growth, } y^B_{i,t,t+h}} \tag{12}$$

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 Bartik predicted growth $y^B_{i,t,t+h}$ into a business cycle $y^{BC}_{i,t,t+h}$ and a trend component $y^{BT}_{i,t,t+h}$,

$$y^B_{i,t,t+h} = y^{BC}_{i,t,t+h} + y^{BT}_{i,t,t+h} \tag{13}$$

where

$$y^{BC}_{i,t,t+h} = \underbrace{\frac{\sum_{g=1}^{G} \mu_{i,g,t} \frac{1}{\sigma_g}}{\sum_{g=1}^{G} \mu_{g,t} \frac{1}{\sigma_g}}}_{\text{Relative Regional Business Cycle Exposure}} \times \underbrace{\left(\Delta^h \hat{y}_{t+h} + \xi \sum_{s=0}^{h-1} \hat{y}_{t+s} \right)}_{\text{Aggregate Business Cycle}} \tag{14}$$

$$y^{BT}_{i,t,t+h} = \sum_{g=1}^{G} \underbrace{\mu_{i,g,t}}_{\text{Regional Trend Exposure}} \times \underbrace{\sum_{s=0}^{h-1} x_{g,t+s}}_{\text{Sectoral Trend}}$$

The two Bartik components have different structures. The business cycle Bartik loads on aggregate business cycle shocks based on the covariance of sectoral weights and sectoral sensitivity. The trend Bartik loads on sectoral trends based on sectoral weights alone.

2.7 Relative Recession Impulse Response Function We exploit the Bartik structure to infer the dynamics of the aggregate business cycle. Let the national business cycle peak be $t = P$ and the recession trough be $t = T > P$. We estimate an IV regression, in which we predict local output growth relative to the peak $\Delta^{T-P+h} y_{i,T+h}$ using local recession depth $\Delta^{P-T} y_{i,P} = y_{i,P} - y_{i,T}$, instrumented by the Bartik predicted recession depth, $y^{BC}_{i,P,T}$,

$$\text{First Stage:} \qquad \Delta^{P-T} y_{i,P} = \gamma_0 + \gamma_1 y^{BC}_{i,P,T} + \nu_{i,T} \tag{15}$$

$$\text{Second Stage:} \qquad \Delta^{T-P+h} y_{i,T+h} = \delta_h + \beta_h \widehat{\Delta^{P-T} y_{i,P}} + \varepsilon_{i,P+h} \tag{16}$$

In what follows we call $\{\beta_h\}_h$ the *relative recession impulse response function*, because it measures the relative behavior of output of a region that experiences a 1% deeper recession compared to another region.

We next state the conditions under which the relative recession impulse response function recovers the aggregate business cycle (11) up to scale.

Proposition 1 *If Bartik predicted recession depth is uncorrelated with local demand shocks* $Cov_i[y^{BC}_{i,P,T}, \Delta^h \hat{\alpha}_{i,P+h}] = 0$ *and differential productivity trends across regions,* $Cov_i[y^{BC}_{i,P,T}, y^{BT}_{i,t,P+h}] = 0$*, then the IV estimator of the relative recession impulse response function recovers the ag-*

gregate recession impulse response function scaled by its size during the recession,

$$\beta_h^{IV} \to -\underbrace{\frac{\Delta^{T+h-P}\hat{y}_{T+h} + \xi \sum_{s=0}^{T+h-P-1} \hat{y}_{P+s}}{\Delta^{T-P}\hat{y}_T + \xi \sum_{s=0}^{T-P-1} \hat{y}_{P+s}}}_{\textit{Scaled Aggregate Recession Impulse Response at } T+h} \tag{17}$$

2.8 Extensions In Appendix A, we extend the to 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.

2.9 Mapping Relative Recession Impulse Response Functions to Theories Proposition 1 makes clear that our approach is able to discriminate between theories for slow recoveries to the extent that these differ in their scaled aggregate recession impulse response function, which we recover via the relative recession impulse response function β_h. Consider the case of an AR(1) beta-shock and an AR(1) monetary policy shock with identical persistence ρ, and assume there is no hysteresis $\xi = 0$. Then the recession trough immediately follows the peak $T = P + 1$. The relative recession impulse response function recovers the aggregate business cycle persistence:

$$\beta_h = \begin{cases} -\rho^h & \text{if } h \geq 0 \\ 0 & \text{if } h < 0 \end{cases} \tag{18}$$

However, the relative recession impulse response function is not informative about the size or identity of the shock. For example, monetary policy shocks and beta shocks with persistence ρ both predict this relative recession impulse response function regardless of their size.

One theory of slow recoveries is that these recessions coincided with a general slowdown in productivity growth (Fernald, Inklaar, and Ruzic, 2025). Our estimates will be unaffected by simultaneous trend shocks if regional business cycle exposure is uncorrelated with aggregate trends. In this case, if the underlying aggregate business cycle shock is the same across recessions with weak and strong aggregate recoveries, we will recover the same relative recession impulse response function β_h even when the aggregate dynamics are very different.

Another view is that slow recoveries are due to bad luck, in the sense that the economy continues to be subject to negative shocks in the recovery phase (Galí, Smets, and Wouters,

2012). One interpretation is that the bad luck manifests itself in a more persistent business cycle shock, which implies a slower recovery. In this case, we should observe that the relative recession impulse response function β_h is more persistent in recessions with slow recoveries, but eventually converges to zero as bad luck goes away.

A third view is that boom-bust cycles give rise to slow recoveries. In this framework we can capture a boom through expectations of higher permanent income relative to the balanced growth path, $\mathbb{E}_t\hat{\lambda}_\infty < 0$. For illustration we make the dynamics simple: Households form these optimistic expectations in one period, $\hat{\lambda}_t = \mathbb{E}_t\hat{\lambda}_\infty < 0$, and learn that these expectations are wrong the following period, $\hat{\lambda}_{t+1} = \mathbb{E}_{t+1}\hat{\lambda}_\infty = 0$. These periods then correspond to the recession peak and trough respectively. The relative recession impulse response function is

$$\beta_h = \begin{cases} -1 & \text{if } h \geq 0 \\ 0 & \text{if } h = -1 \\ -1 & \text{if } h < -1. \end{cases} \tag{19}$$

That the relative recession impulse response function remains at -1 after the recession implies that regions with a deeper recession do not experience a stronger recovery. Recoveries look weak in this case because the boom represents an unsustainable deviation from the balanced growth path. Importantly, this theory predicts that regions with a deeper recession should grow relatively more prior to the recession.

Now consider the hysteresis case. For simplicity we assume the underlying shock is an AR(1) shock to beta or the real interest rate at $t = 0$. The relative recession impulse response function β_h incorporates the hysteresis effect:

$$\beta_h = \begin{cases} -\rho^h - \xi\frac{1-\rho^{T-P-1+h}}{1-\rho} & \text{if } h \geq 0 \\ 0 & \text{if } h < 0 \end{cases} \tag{20}$$

The relative recession impulse response function converges to $-\frac{\xi}{1-\rho}$ because hysteresis causes a permanently lower level of output. Thus, similar to the boom-bust scenario, hysteresis predicts that relative recoveries look weak. In contrast to the boom-bust cycle, there is no pre-trend in the hysteresis relative recession impulse response function.

2.10 Identifying Assumptions The model highlights the necessity of an identification strategy for the relative recession impulse response function, as OLS estimates are almost certainly biased. In particular, local demand shocks $\hat{\alpha}_{it}$ reduce both regional recession depth $\Delta^{P-T}y_{i,P}$ and increase regional growth $\Delta^{T-P+h}y_{i,T+h}$ biasing the relative recession impulse

response function β_h downward. The identification problem is reminiscent of cross-country panel approaches, which must assume that country-specific shocks are uncorrelated with recession intensity.

The Bartik approach identifies the relative recession impulse response function under two conditions. First, States with greater exposure to the national cycle due to their industry structure cannot systematically experience either positive or negative local shocks, $Cov_i[y^{BC}_{i,P,T}, \Delta^h \hat{\alpha}_{i,P+h}] = 0$. This is the standard Bartik assumption. We remove any mechanical correlation between the Bartik instrument and local shocks by constructing leave-one-out industry growth rates.

Second, exposure to the national cycle must be orthogonal to differences in exposure to industry trends. This condition is more challenging to satisfy as manufacturing is known to be both cyclically sensitive and in secular decline. Conditioning on recession pre-trends is not valid in our setting as pre-trends may themselves be causally related to the recession, such as in a boom-bust cycle.

We therefore identify relative trends correlated with recession exposure in two ways. First, we estimate β_h sufficiently far in advance of the recession such that it is not affected by a potential boom-bust cycle. This strategy is valid under all theories of weak recovery, so long as h is sufficiently negative. Second, we use the β_h for unemployment to determine at what point the economy is at a similar point in the business cycle. As shown by Blanchard and Katz (1992) and Yagan (2019), unemployment converges across regions even when other variables do not, which makes it suitable for isolating their relative trends from the business cycle.

2.11 Comparison with Other Approaches Regional identification of the relative recession impulse response function conveys two important advantages compared to time-series or panel identification of the aggregate recession impulse response function. First, the construction of the cyclical Bartik does not require taking a stance on the aggregate trend. Thus one does not need to take a stance on whether aggregate GDP is trend-stationary or a random walk, which is notoriously difficult (Stock and Watson, 1999), nor does one need to engage in aggregate trend extrapolation, which Ball (2015) shows is very sensitive to timing decisions.

The construction of the cyclical Bartik does require taking a stance on industry trends, but we show next that it is much more robust to trend misspecification compared to time-series or cross-country panel approaches. Denote the measurement error in industry-level trend at time t by $\bar{e}_{g,t}$. In Appendix B we show that we continue to identify the relative recession impulse response function if variation in measurement error across goods g

is orthogonal to the output growth of the typical good-g producer over the recovery phase, $Cov_g[\Delta^{T-P}\bar{e}_{g,P}, \sum_i \mu_{i,g,t}\Delta^{T+h-P}y_{i,T+h-P}] = 0$.[2] Consistency follows from an appropriate law of large numbers, which imposes a restriction on the joint distribution of the covariance structure of measurement error and regional exposures. For example, this restriction is satisfied if variation in measurement error across goods is as good as randomly assigned and regional production is sufficiently diversified in the limit $E_g\left(\sum_i \mu_{i,g,t}^2\right) \to 0$.[3] Appendix B also provides alternative conditions that allow for correlated measurement error across goods. Importantly, we do not have to restrict the serial correlation in measurement error in any way. This is because these assumptions imply that good-level trend measurement error washes out asymptotically at the regional level, and thereby justify estimating relative recession impulse response functions at the regional rather than at the industry level.

These assumptions also contrast with the much stronger assumption required in a time-series or cross-country panel approach, where one must impose that measurement error in the estimated aggregate trend be orthogonal to output growth across recession episodes j, $Cov_j[\Delta^{T_j-P_j}\bar{e}_{P_j}, \Delta^{T_j+h-P_j}y_{T_j+h}] = 0$. In this case, any stochastic component in aggregate trend growth, such as a productivity slowdown that coincides with recession episodes, will result in a non-zero covariance between the measurement error term and post-recession output growth, as we show formally in Appendix B. Thus, recovering the aggregate recession impulse response function from time-series variation requires trend growth to be either deterministic or measured without error.

Second, regional data holds the country and the recession fixed, which facilitates identifying the dynamics of the recession. For instance, a more persistent aggregate demand shock will induce a more persistent recession across more and less exposed regions, and thus poses no problem for our identification strategy. In a cross-country or in a time-series setting two recessions may correspond to shocks with different dynamics. 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 scaled "treatment effect" of -1% in the first year and -1% in the second year, and thus greater persistence than implied by either recession.

[2] Any common trend measurement error across industries is differenced out in the regional regressions (15)-(16), and thus poses no threat to identification.

[3] These assumptions are equivalent to Assumption 2 in Borusyak, Hull, and Jaravel (2022). As in their work, our asymptotics implicitly assume that $G(I) \to \infty$ as $I \to \infty$.

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,[4] and, post-2005, quarterly State-level GDP from the BEA. We deflate all 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, 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 detrend both State recession depth and national industry growth using NBER-peak to NBER-peak growth rates. This gives us (1) a slightly stronger first stage and (2) puts more weight on sectors that experience a sharp change in growth rates in a recession relative to sectors that grow fast or slow due to an underlying trend. 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.[5]

4 Illustration of the Empirical Approach

Goldsmith-Pinkham, Sorkin, and Swift (2020) show that estimation using the Bartik instrument (14) is equivalent to weighted GMM using the shares s_{i,k,P_j} as instruments.[6]

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

[5] The Personal Income files switch to a more detailed SIC classification in 1957. We aggregate these to the coarser historical classification for the purpose of constructing peak-to-peak growth rates for the 1957-8 recession.

[6] Adao, Kolesár, and Morales (2019) and Borusyak, Hull, and Jaravel (2022) also show that identification

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.[7] The average recession begins five quarters before the 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.[8]

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 areas 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. Our analysis does not require taking a stance on the aggregate pre-recession 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

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.

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

[8] The relative recovery adjusting for the trend is 1.8% + 15*0.11% = 3.4%, which is equal to the 3.4% relative recession depth.

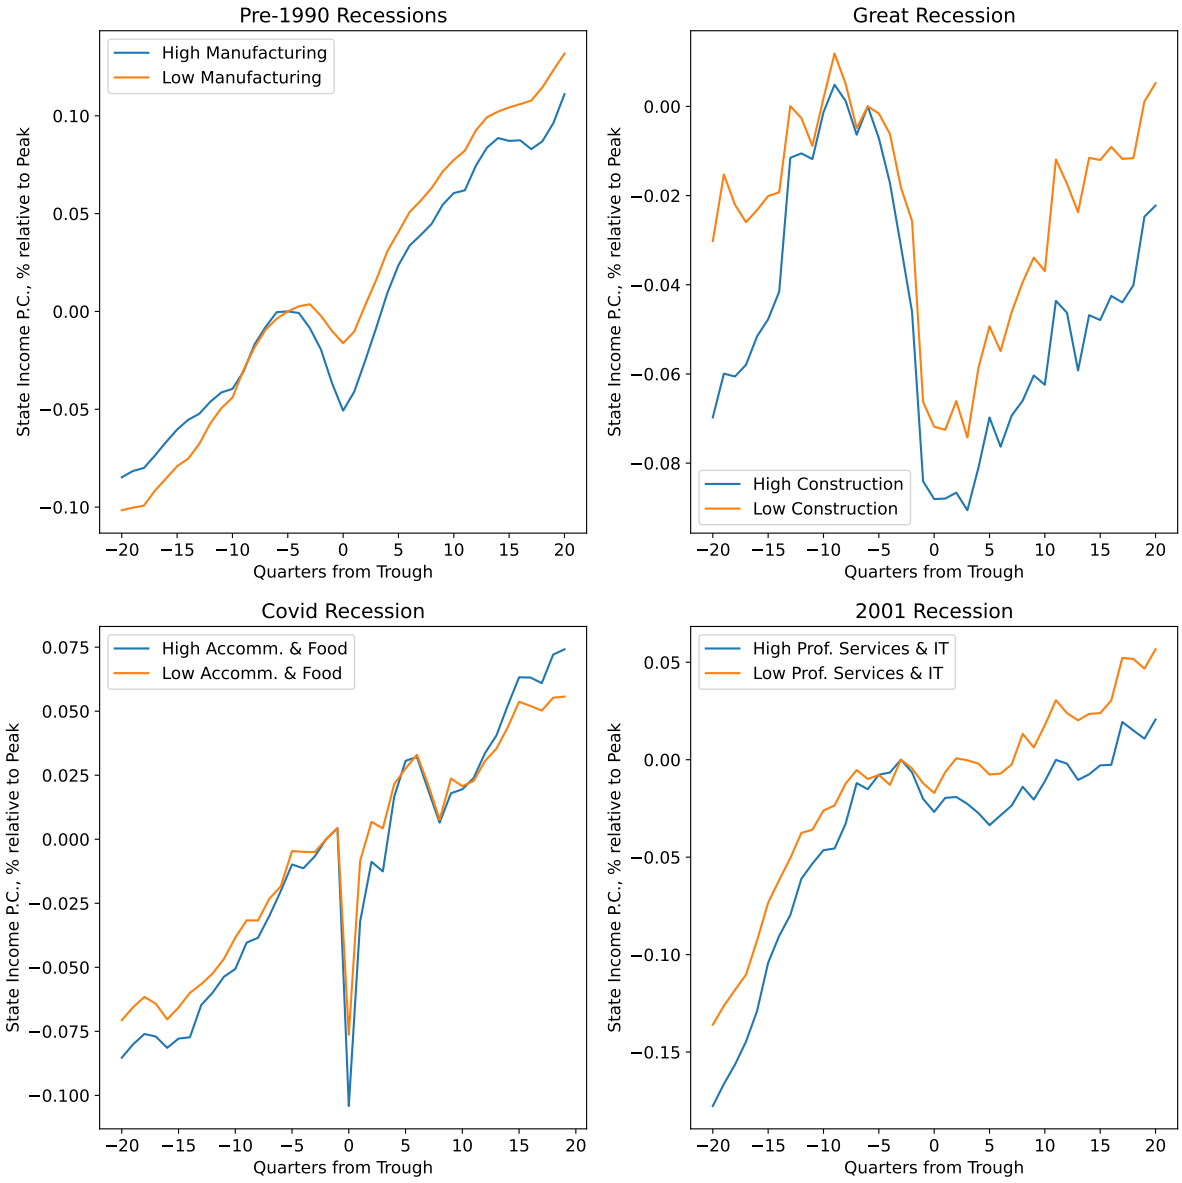


Figure 3: Outcomes by Key Sources of Variation of the Bartik Instrument

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 steeper 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 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 relative recession impulse response functions, 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 relative recession impulse response function: 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 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 and is supported by the similar qualitative recession pattern we document above.

We start with the pooled first stage,

$$\begin{aligned}\hat{y}_{i,P_j} - \hat{y}_{i,T_j} = \gamma_1(\hat{y}^b_{i,P_j} - \hat{y}^b_{i,T_j})I(P_j \in \{1990Q3, 2001Q1, 2007Q4\}) \\ + \gamma_2(\hat{y}^b_{i,P_j} - \hat{y}^b_{i,T_j})I(P_j \in \{\text{Other Recessions}\}) + \theta_j + \varepsilon_{i,j,h}\end{aligned}$$

in which we regress recession depth relative to trend $\hat{y}_{i,P_j} - \hat{y}_{i,T_j}$ on Bartik-predicted recession depth $\hat{y}^b_{i,P_j} - \hat{y}^b_{i,T_j}$. We allow the relationship to vary across recession groups. γ_1 is the elasticity of actual recession depth 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. We control for

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

Local Recession Depth:	Inc. Earnings	Earnings	Unemploy.	Employ.	GDP
	(1)	(2)	(3)	(4)	(5)
Coeff. on Predicted Recession Depth:					
90-1, 01, 07-9 Recessions	3.11***	2.49***	−54.7**	1.89***	5.66***
	(0.71)	(0.67)	(25.1)	(0.37)	(0.85)
Pre-1990 and 2020 Recessions	1.42***	0.67***	−63.3***	0.86***	0.85***
	(0.28)	(0.22)	(9.76)	(0.17)	(0.28)
Recession FE	Yes	Yes	Yes	Yes	Yes
F-Statistic 90-1, 2001, 07-9	19.0	13.9	4.8	25.6	44.3
F-Statistic Other Rec.	25.8	8.8	42.1	25.1	9.3
State Clusters	100	100	100	100	
R^2	0.36	0.28	0.62	0.62	0.42
Observations	544	544	541	541	100

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

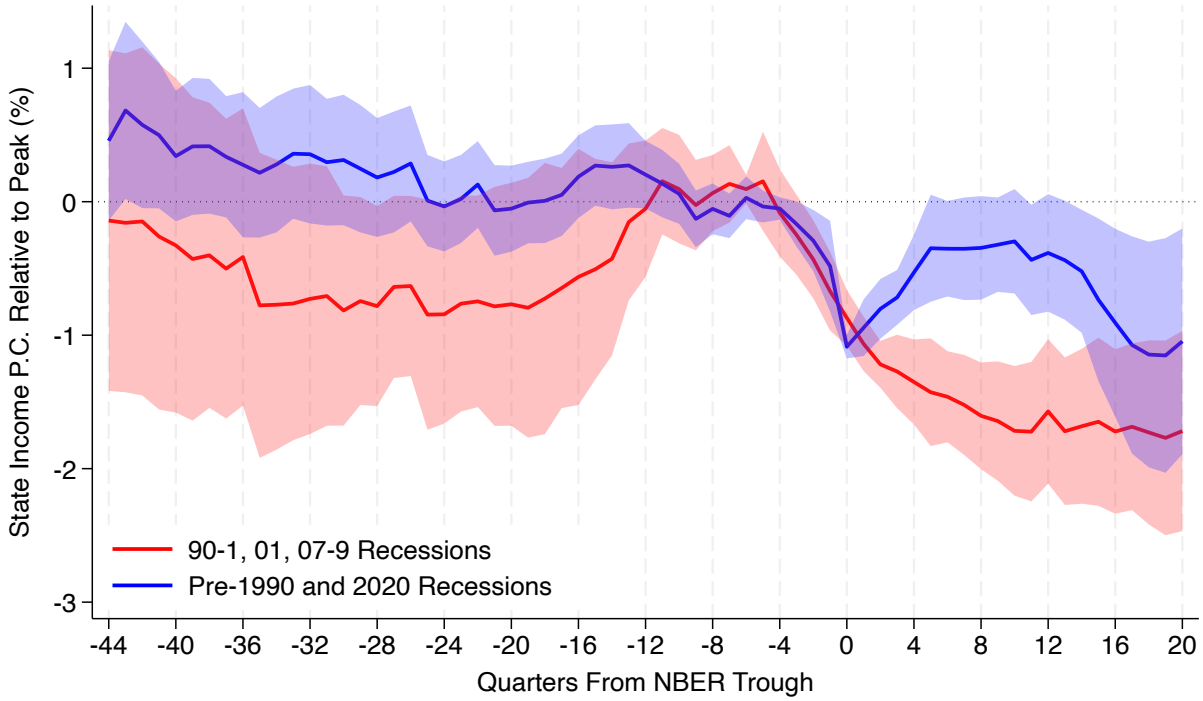
recession fixed effects θ_j to isolate within-recession variation. Within each Recession Group we cluster by State.[9]

The pooled first stage estimates are shown in Table 1. In columns (1) and (2) we show that predicted recession depth based on State income shares and national trends strongly predicts State recession depth in income. For the 1990-1, 2001, and 2007-9 recessions the first stage coefficient for income included in the instrument is $\gamma_1 = 3$ (t-statistic 4.3). This means that a 1% deeper predicted recession translates into a 3% deeper local recession. For the other recessions the first stage coefficient is $\gamma_2 = 1.4$ (t-statistic 5.1). In columns (3) through (5) we measure State recession depth relative to trend with the unemployment rate, employment, and GDP. In each case, our Bartik instrument is a powerful predictor of the State recession depth.

5.2 Relative Recession Impulse Response Functions We next estimate relative recession impulse response functions across recession groups. Within each group, these capture the differential movement of income across more and less exposed States both before and after the national recession.

[9]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 10. This suggests that our Bartik approach largely captures cross-state correlations, consistent with Adao, Kolesár, and Morales (2019).

Figure 4: Impulse Response Function of State Personal Income P.C. by Recession Group



Specifically, we pool equation (16) across recession groups,

$$y_{i,T_j+h} - y_{i,P_j} = \beta_{1h}(\hat{y}_{i,P_j} - \hat{y}_{i,T_j})I(P_j \in \{1990Q3, 2001Q1, 2007Q4\}) + \beta_{2h}(\hat{y}_{i,P_j} - \hat{y}_{i,T_j})I(P_j \in \{\text{Other Recessions}\}) + \alpha_{j,h} + \varepsilon_{i,j,h} \quad (21)$$

where $y_{i,T_j+h} - y_{i,P_j}$ is h-quarter growth in the outcome variable since the recession peak P_j. We instrument recession depth relative to trend $(\hat{y}_{i,P_j} - \hat{y}_{i,T_j})$ with predicted recession depth relative to trend $(\hat{y}^b_{i,P_j} - \hat{y}^b_{i,T_j})$ (both interacted with the recession group indicators). The parameters β_{1h} and β_{2h} measure the differential trajectory for a State with a 1% deeper recession compared to another State h quarters after the national recession trough. Since h can be negative we use the same regression to estimate differences in the pre-recession trend. β_{1h} captures the differential trajectories for the 1990-1, 2001, and 2007-9 recessions and β_{2h} for the other recessions.

Figure 4 plots our estimates β_{1h} and β_{2h} for $h = -44$ quarters before the recession trough to $h = 20$ quarters after the recession trough. For $h = 0$, the estimate is close to -1 for both recession groups. If there were no trends in our measure of local recession depth, these

estimates would be equal to -1 exactly. They are close to -1 because the relative trends are not very important over the short recession duration.

Other than the immediate period prior to the recession trough $h = 0$, the two recession groups look different. For the 1990-1, 2001, and 2007-9 recessions, a deeper relative recession predicts a weaker relative recovery. 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 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 $h - 44$, the more exposed States are growing at either the same or a slightly slower pace than the less exposed States. Our estimates for the recovery are consistent with a return to the slower relative trend five years before the recession trough.

The impulse response functions we estimate for the 1990-1, 2001, and 2007-9 recessions are consistent with the boom-bust cycle based explanations for weak recoveries (Beaudry, Galizia, and Portier, 2018; Jordà, Schularick, and Taylor, 2013; Rognlie, Shleifer, and Simsek, 2018). Comparing our estimates with Figure 2d shows that they are not consistent with bad luck or an exogenous trend slowdown, because we do not see regional convergence after the recession. They are also not consistent with an important role for hysteresis, because we estimate a pre-recession boom and the return to the pre-boom trend.

In contrast, the more exposed States in the pre-1990 and 2020 recessions grow relatively more quickly in the recovery, which is captured by the increasing estimates β_{2h} for $h > 0$. Five quarters after the national recession trough, β_{2h} is close to, and statistically indistinguishable from, zero. A zero effect implies that any deeper local recession was subsequently made up for by a stronger recovery.[10] The argument for a complete recovery is stronger once we account for the downward relative trend of more exposed States prior to the recession. For the pre-1990 and 2020 recessions the average relative trend is continuously downward and does not display a break.

In summary, the more exposed States to national recessions pre-1990 and 2020 experience a more pronounced V-shape recession followed by a recovery relative to trend. This pattern is consistent with plucking models of the business cycle (Friedman, 1993; Dupraz, Nakamura, and Steinsson, 2023). The fact that we find no evidence for boom-bust cycles for these recessions suggests that the pre-1990 and 2020 recessions and the 1990-1, 2001, and 2007-9

[10]The subsequent decline for $h > 12$ quarters reflects a subsequent recession: the peak of the 1960-1 recession follows only eleven quarters after the 1957-8 recession.

recessions are caused by 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 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 appears to be more complete based on all income, but the relative downward trend in more exposed States prior to the recession is also less apparent.

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.

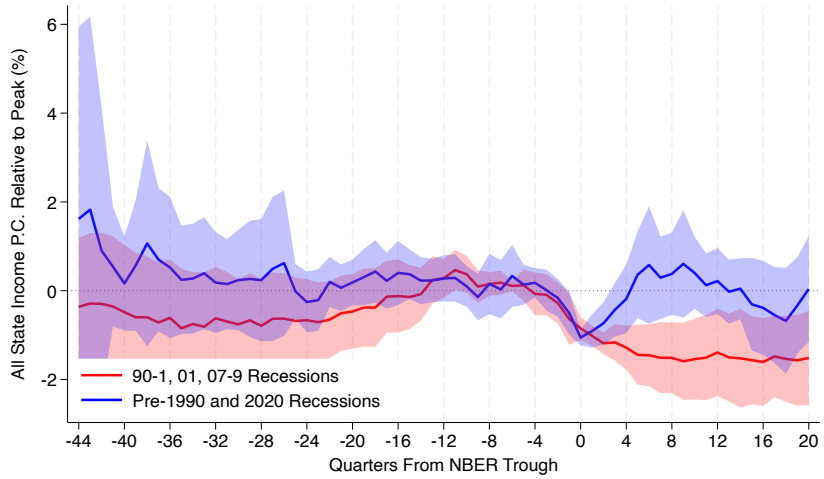
There are no noticeable pre-trends in unemployment for the pre-1990 and 2020 recessions. This is consistent with our interpretation that the relative downward trend in State income for more exposed States (Figure 4) reflects structural rather than cyclical forces.

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 small. The boom in State income in Figure 4 in the five years prior to the recession is almost as large as the bust. The decline in unemployment in the five years prior to the recession is only one-fifth as large as the subsequent increase in unemployment. This suggests that boom-bust cycles may have asymmetric effects on labor markets and that unemployment rates are less informative about the boom-bust nature of a recession.

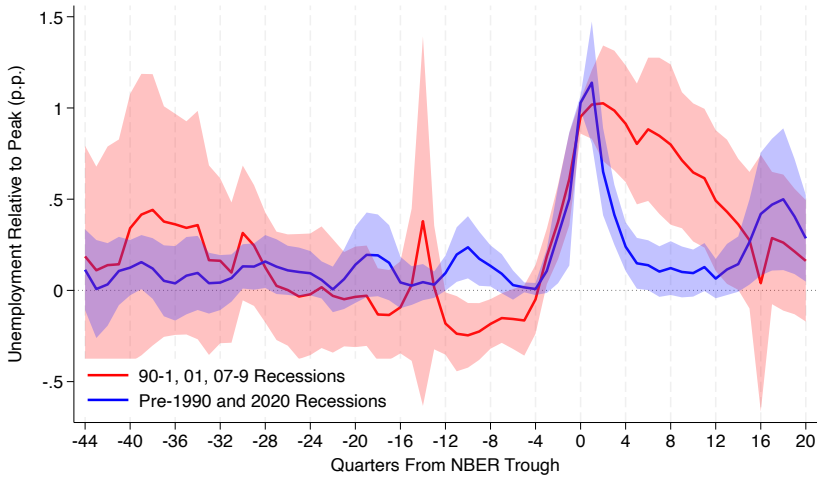
From $h = -44$ to $h = -20$ relative unemployment is relatively flat or slightly decreasing, suggesting that the length of this period is informative about the relative structural trend for the more exposed States in the 1990-1, 2001, and 2007-9 recessions. This points to a relative downward trend in the more exposed States based on both State income and (see below) employment.

Figure 5: Recession Impulse Response Function for Other Outcomes

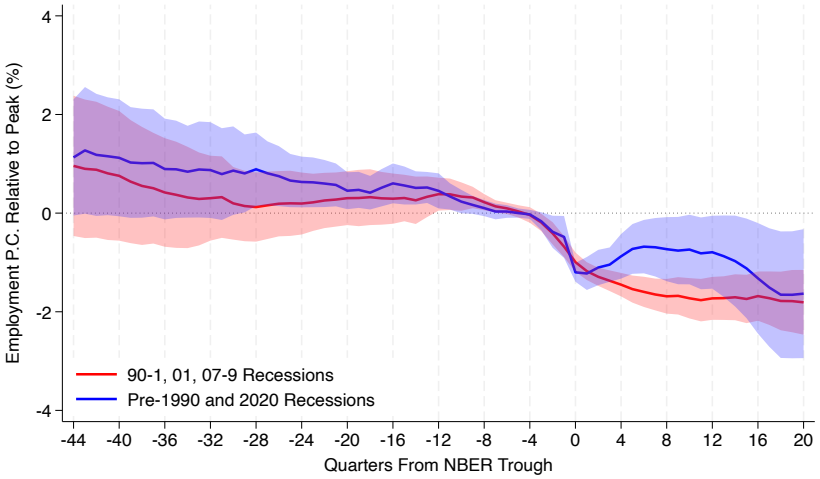
(a) All State Income P.C.



(b) Unemployment Rate



(c) Employment P.C.



(d) GDP P.C.

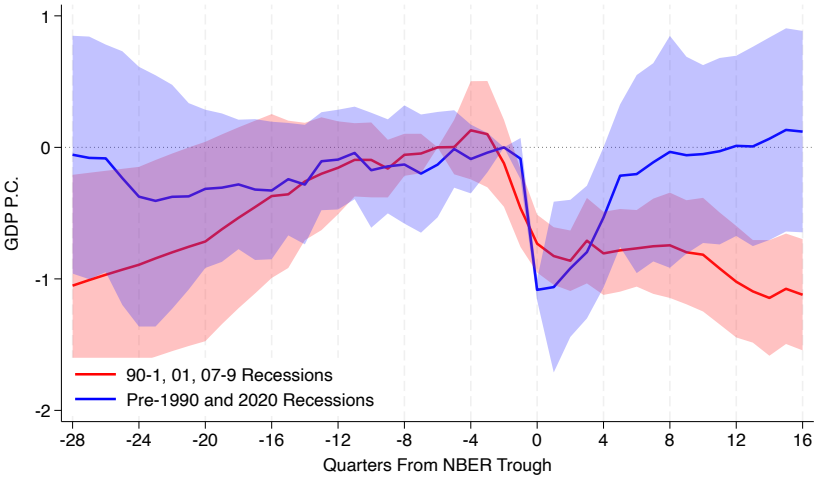


Figure 5c shows that more exposed States experience a much stronger relative decline in employment relative to population. The boom prior to the 1990-1, 2001, and 2007-9 recessions barely turns this trend positive. What looks like a level downward-shift in employment after the recession is consistent with a return to the pre-boom trend.[11] For the pre-1990 and 2020 recessions, the downward trend in employment is particularly noticable. Nevertheless, the recovery is V-shaped and employment quickly returns to its pre-recession trend.

The GDP data is more limited since it is only available for the last two business cycles. Yet the patterns 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 group in Appendix Figure 11. 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) 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 of the Impulse Response Function We decompose our Bartik estimates to (1) show what variation is most important and (2) determine whether the changes across recessions come from different behavior within or across industries (or both). 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} in (14) are the instruments and industry growth rates $\hat{g}_{-i,k,P_j,T_j}$ are the weights. When the industry growth rates are common across locations, then we can decompose the Bartik estimate for β_h in (16) 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 (16) 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,

[11] Yagan (2019) performs a different exercise by projecting the local employment to population ratio on the local unemployment increase over the recession. This regression in our data also shows no pre-trend over 2003-7, but a pre-trend prior to 2003. Furthermore, projecting the local employment to population ratio on the local decline in the employment to population ratio shows a boom period prior to the recession. This suggests that an analysis of pre-trends for the Great Recession depends importantly on the horizon and the variation used.

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 remains 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 12 includes all sectors.

The three major sectors that inform recession depth in the 1990-1, 2001, and 2007-9 recession are construction ($\alpha = 0.28$), transportation and utilities ($\alpha = 0.06$), and other services ($\alpha = 0.59$). 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. Thus, the boom-bust nature of these recessions is broad-based rather than confined to a particular industry. It also suggests 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 the same 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 ($\alpha = 0.54$), wholesale and retail trade ($\alpha = 0.09$), 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. There is no evidence of boom-bust cycles. 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.

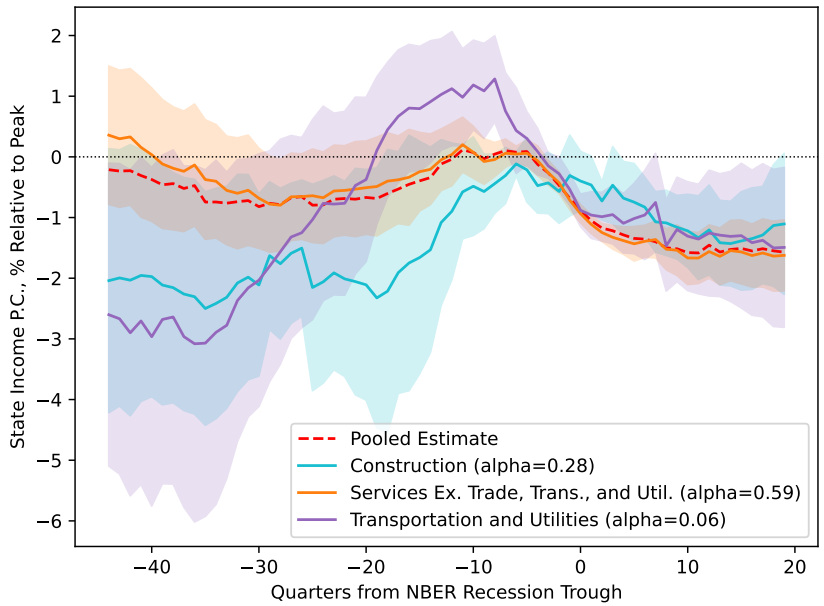
Figures 6c and 6d show the impulse response functions when the variation is restricted to particular recessions. For both the 2001 and 2007-9 recession the boom-bust pattern is evident and each contains the pooled estimate within their 95% confidence bands. This is despite these recessions loading on 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 $\alpha = 0.02$ and wide confidence band (Appendix Figure 12c).

There is no evidence of a boom-bust cycle in the pre-1990 or the Covid recessions in Figure 6d. Instead, 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. However, States exposed to the pre-1990 recessions and Covid recession differ in their pre-trends and thus in the extent to which their relative recovery looks complete in levels, though not in trends. For the pre-1990 recessions, the more exposed States are on a relative downward trend due to their higher manufacturing exposure. The recovery, while incomplete in levels, converges to this downward trend. In the Covid recession, the more exposed States are on an upward trend prior to the recession and recover completely both in levels and to their prior relative trend.

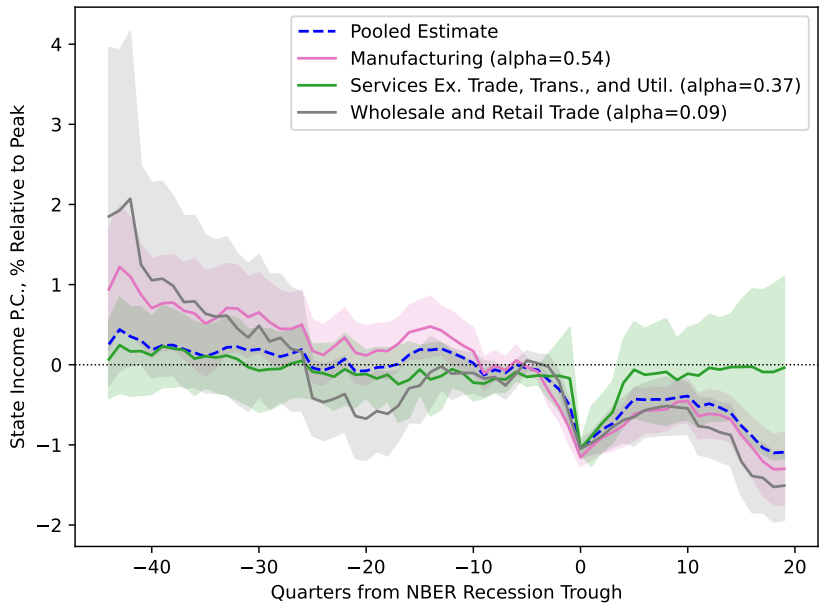
In short, weaker recoveries following a boom-bust cycle are a broad-based phenomenon across sectors and recessions for the 1990-1, 2001, and 2007-9 recessions. In contrast, the strong recoveries from the pre-1990 and Covid recessions are also broad-based, and they do not display any boom-bust dynamics.

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

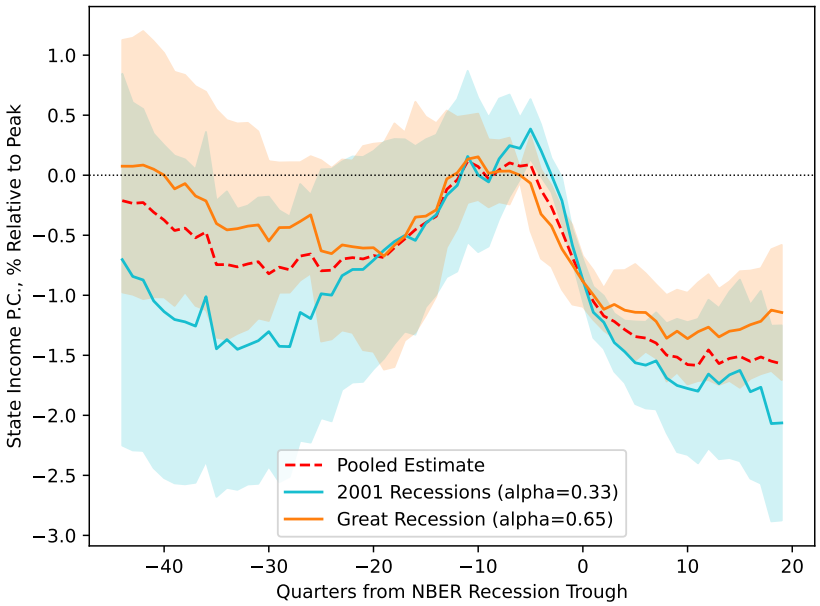
(a) By Industry: 1990-1, 2001, 2007-9



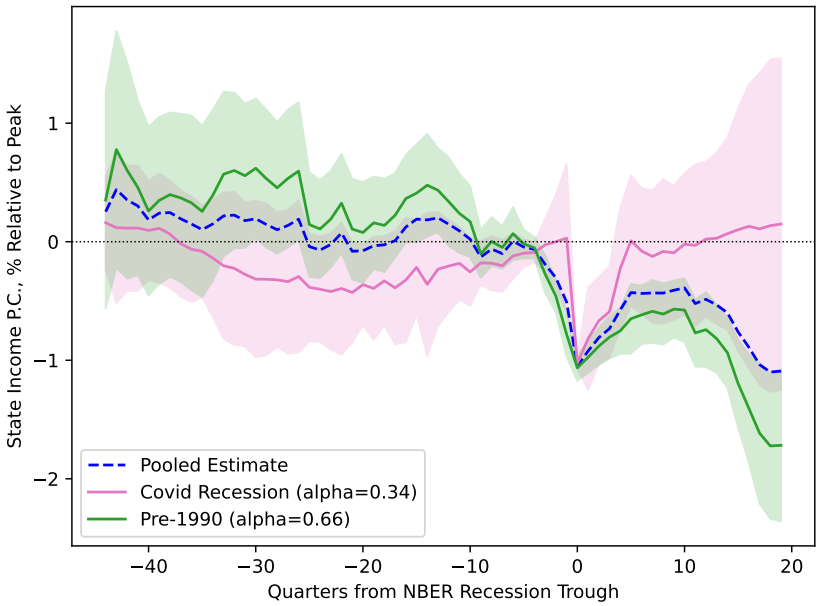
(b) By Industry: Other Recessions



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



(d) By Recessions: Other Recessions



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 7a to 7f. 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 7a and 7b), the middle panel varies the fixed effects (Figures 7c and 7d), and the bottom panel changes the sample (Figures 7e and 7f). In each case, we re-estimate the first stage and second stage for each variation.

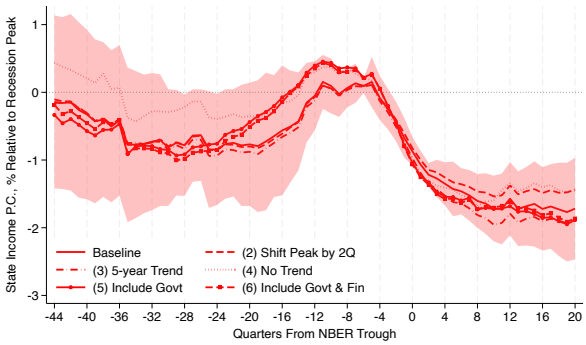
Figures 7a and 7b show how the IRFs for the 1990-1, 2001, 2007-9 recessions and the other recessions change when we construct the instrument differently. In (2) we construct the peak-to-peak national industry growth rates two quarters prior to the NBER peak and in (3) we construct it using five-year pretrend from the NBER peak. In (4) we omit detrending local recession depth. Neither of these changes materially affects the impulse response functions. In (5) and (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 and is a bit more pronounced.

In Figures 7c and 7d 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.

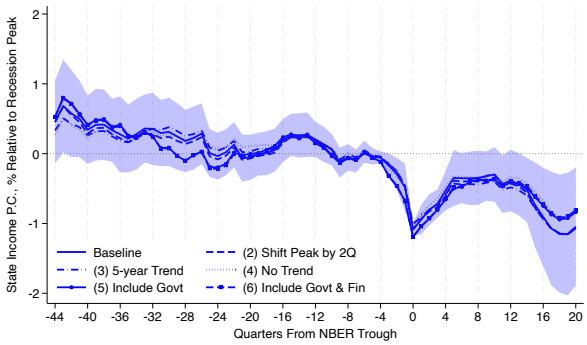
In Figures 7e and 7f we show the sensitivity to the sample. In (2) we drop the three most mining-intensive States and in (3) the three most finance-intensive States, the two categories in which we are least confident that the location of income corresponds to production. The differences from our baseline 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 7: IRFs Robustness

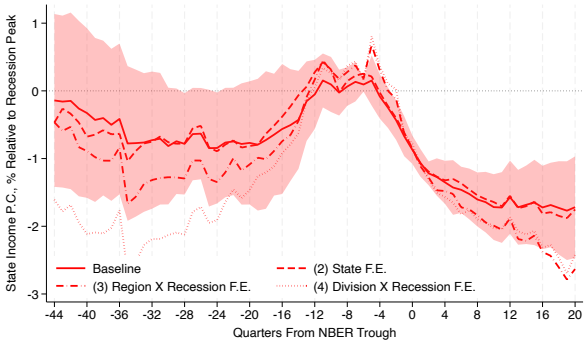
(a) Instrument Variation: 90-1, 01, 07-9



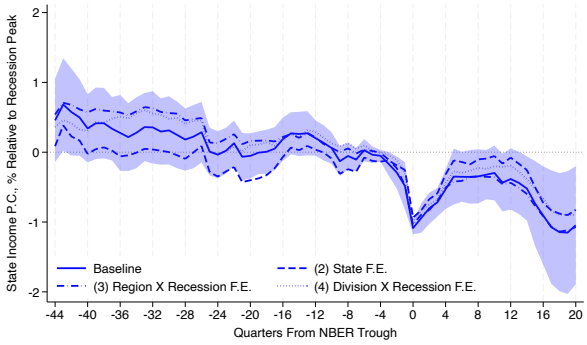
(b) Instrument Variation: Other Recessions



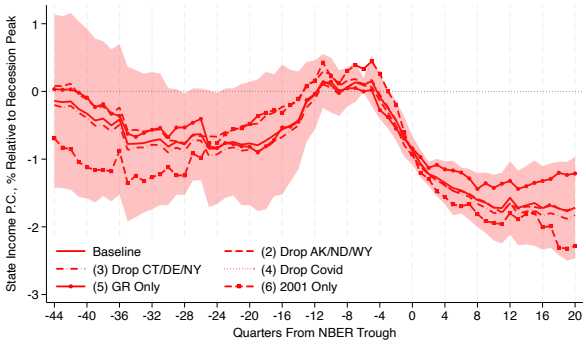
(c) Fixed Effects: 90-1, 01, 07-9



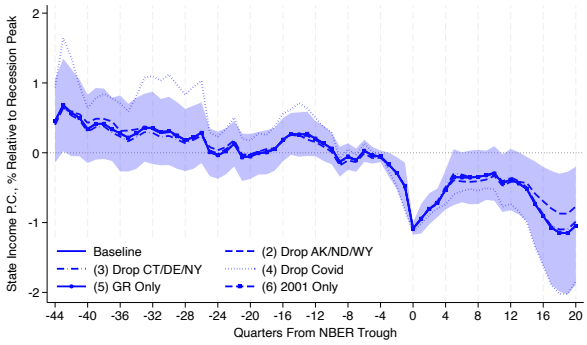
(d) Fixed Effects: Other Recessions



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



(f) Sample Variation: Other Recessions



6 Quantitative Model

The objective of this section is to build a quantitative model that is consistent with the two different types of recessions we document above and then to use that model for policy analysis. The model has two region and two goods, with goods subject to standard New-Keynesian price stickiness. The two-good structure allows for the regions to be differentially exposed to a national recession so we can run the same regressions in the model as we do in the data.

6.1 Households Households maximize expected utility subject to their budget constraint. The expected utility function of household i is given by:

$$\mathrm{E}_{i,0} \sum_{t=0}^{\infty} \beta^t \left[\frac{(C_t)^{1-\sigma}}{1-\sigma} + \psi_t \varepsilon_{\psi t} \frac{{D_{i,t}}^{1-\sigma_d}}{1-\sigma_d} - \varphi_t \frac{{L_t}^{1+\nu}}{1+\nu} \right]$$

where C_t is consumption of a composite nondurable good, $D_{i,t}$ is the durable stock of household i and L_t is total labor supply. ψ_t and φ_t are time-varying preference shifters for durable goods consumption and labor supply that keep hours worked and the ratio of nondurable to durable consumption constant along the balanced growth path.[12] $\varepsilon_{\psi t}$ is an exogenous stationary shock to durable preferences that we use to model a boom-bust cycle.

The durable stock D_{it} is specific to the household. We assume that each period only a fraction $(1-\theta^d)$ of households is allowed to re-optimize its durable stock. This formulation, which follows Evans and Ramey (1992) and Orchard, Ramey, and Wieland (2023), reduces the counterfactually high sensitivity of durable consumption to interest rate changes present in standard models, while preserving tractability. Nondurable consumption C_t is identical between households due to consumption insurance and labor supply L_t is identical because it is set by labor unions.

Expectations are specific to household i due to information stickiness as in Auclert, Rognlie, and Straub (2020) and Carroll, Crawley, Slacalek, Tokuoka, and White (2020). The current state at t is common knowledge, so households never violate their budget constraint.

[12] The process for φ_t is such that it implies constant hours along a balanced growth path:

$$\varphi_t = (C_t)^{-\sigma} Z_t$$

where Z_t is the level of productivity of the aggregate economy. The taste for durables, ψ_t, contains a productivity component as well as an exogenous stationary process $\varepsilon_{\psi t}$:

$$\psi_t = \bar{\psi} Z_t^{\sigma - \sigma^d}.$$

There is uncertainty about the future path of aggregate shocks and households update their expectations with i.i.d. probability $1-\theta$ in each period.

The aggregate budget constraint and the law of motion for the aggregate durable stock are given by:

$$\begin{aligned}
A_t &= (1+i_{t-1})A_{t-1} + W_t L_t - P_t C_t - \iota P_t D_t - P_t^x X_t \\
D_t &= (1-\delta)D_{t-1} + X_t \\
D_t &= \int_0^1 D_{i,t} di
\end{aligned}$$

where P_t is the price index for nondurable goods, A_t are nominal bond holdings, i_t is the nominal interest rate in terms of nondurable goods, $W_t L_t$ is nominal labor income, D_t is the aggregate durable stock, P_t^x is the price index of durable goods , X_t is the composite durable good and ιD_t are operating costs from using the durable goods (e.g. gas, insurance), which we assume are denominated in units of nondurable goods.

Durable purchases, X_t, include maintenance costs, given by $\chi\delta D_{t-1}$, and new durable purchases, I_t. Maintenance costs require households to pay a portion χ of depreciation as maintenance expenditures in each period. This yields:

$$\begin{aligned}
X_t &= \chi\delta D_{t-1} + I_t \\
I_t &= D_t - (1-(1-\chi)\delta)D_{t-1}
\end{aligned}$$

Households optimally choose the schedule for $\{C_{it}, D_{i,t}, A_{it}, X_{it}\}_{t=0}^{\infty}$ that maximizes their utility given their information set. Labor supply is not chosen directly by the household, but by regional labor unions which we describe later on. The separability of preferences allows us to break up the household problem in two steps. We first solve for the optimal choices of C_t and A_t and, given these, we solve the durable choice problem. The first order conditions for C_t and A_t yield a standard Euler equation:

$$\begin{aligned}
\lambda_t &= \beta \frac{(1+i_t)}{\Pi_{t+1}} \lambda_{t+1} \\
\lambda_t &= (C_t)^{-\sigma}
\end{aligned}$$

where λ_t is the Lagrange multiplier on the aggregate budget constraint and $\Pi_t = P_t/P_{t-1}$ is the gross inflation rate for nondurable goods. The optimal choice for $D_{i,t}$ conditional on

household i making an adjustment at time t solves:

$$\max_{D_{i,t}} E_t \sum_{s=0}^{\infty} (\beta\theta^d)^s \left[\psi_{t+s}\varepsilon_{\psi t+s} \frac{\left((1-(1-\chi)\delta)^s D_{i,t}\right)^{1-\sigma_d}}{1-\sigma_d} - \lambda_{t+s}\iota(1-(1-\chi)\delta)^s D_{i,t} \right]$$
$$- \lambda_t p_t^x D_{i,t} + \sum_{s=0}^{\infty} \beta^s \theta^{d s-1}(1-\theta^d)(1-(1-\chi)\delta)^s \lambda_{t+s} p_{t+s}^x D_{i,t}$$

where θ^{ds} is the probability that the durable stock chosen at t survives s periods into the future and $p_t^x = P_t^x/P_t$ is the relative price of the durable good. The first term in brackets, $\psi_t\varepsilon_{\psi t}\frac{\left((1-(1-\chi)\delta)^s D_{i,t}\right)^{1-\sigma_d}}{1-\sigma_d}$, captures the utility flow of the durable stock in $t+s$. The term $\lambda_{t+s}\iota(1-(1-\chi)\delta)^s D_{i,t}$ measures the cost of operating the durable stock in utils of time $t+s$ while the term $p_t^x D_{i,t}$ measures the cost of purchasing the durable stock in period t. The last term captures the resale value of the durable stock if an adjustment opportunity arrives in period $t+s$. All households that re-optimize at t solve the same problem, therefore, we can denote their optimal durable stock choice by D_t^*. It can be shown that D_t^* satisfies the following Calvo-type specification[13]:

$$D_t^* = \left(\frac{\Omega_{1,t}}{\Omega_{2,t}}\right)^{\frac{1}{\sigma_d}}$$
$$\Omega_{1,t} = \psi_t\varepsilon_{\psi t+s} + \beta\theta^d(1-(1-\chi)\delta)^{1-\sigma_d}\Omega_{1,t+1}$$
$$\Omega_{2,t} = \lambda_t\left(p_t^x + \iota - \frac{(1-(1-\chi)\delta)}{(1+i_t)\Pi_{t+1}^{-1}} p_t^x\left(cg_{t+1}-1\right)\right) + \beta\theta^d(1-(1-\chi)\delta)\Omega_{2,t+1}$$

where $cg_{t+1} = 1 + \frac{p_{t+1}^x}{p_t^x}$ is the capital gain from selling the stock of durable goods in period $t+1$. The optimal aggregate purchases of durable goods are given by:

$$X_t = (1-\theta^d)\Big[D_t^* - (1-(1-\chi)\delta)D_{t-1}\Big] + \chi\delta D_{t-1}$$

where the first term captures new durables purchases by households that are allowed to reset their durable stock and the second term captures maintenance expenditures which are incurred by all households, regardless of their re-set status. Aggregate nondurable expenditure, ND_t, is defined as the sum of consumption of the nondurable good plus operating costs:

$$ND_t = C_t + \iota D_t$$

[13]This formulation is analogous to the optimal re-set price in a standard model with Calvo-type price rigidities.

6.2 Regional Demand The composite nondurable good is given by:

$$ND_t = \left[\gamma_c {ND_{1t}}^{\frac{\eta_c - 1}{\eta_c}} + (1-\gamma_c){ND_{2t}}^{\frac{\eta_c - 1}{\eta_c}}\right]^{\frac{\eta_c}{\eta_c - 1}}$$

where ND_{rt} denotes consumption of nondurable goods produced in region r.

The composite X_t aggregates durable goods produced in each region according to:

$$X_t = \left[\gamma_x {X_{1t}}^{\frac{\eta_x - 1}{\eta_x}} + (1-\gamma_x){X_{2t}}^{\frac{\eta_x - 1}{\eta_x}}\right]^{\frac{\eta_x}{\eta_x - 1}}$$

where X_{rt} denotes purchases of durable goods produced in region r.

Similarly, total labor supply is a CES aggregate over labor supplied to each region:

$$L_t = \left[\gamma_\ell {L_{1t}}^{\frac{\eta_\ell - 1}{\eta_\ell}} + (1-\gamma_\ell){L_{2t}}^{\frac{\eta_\ell - 1}{\eta_\ell}}\right]^{\frac{\eta_\ell}{\eta_\ell - 1}}$$

where L_{rt} denotes hours worked in region r. η_ℓ governs the elasticity of substitution of labor supply across regions and γ_ℓ captures the relative preference of households to work in region 1.

Given these preferences, we have the following regional demand schedules:

$$ND_{1t} = \gamma_c ND_t \left(\frac{P_{1t}}{P_t}\right)^{-\eta_c} \qquad X_{1t} = \gamma_x X_t \left(\frac{P^x_{1t}}{P^x_t}\right)^{-\eta_x}$$

$$ND_{2t} = (1-\gamma_c) ND_t \left(\frac{P_{2t}}{P_t}\right)^{-\eta_c} \qquad X_{2t} = (1-\gamma_x) X_t \left(\frac{P^x_{2t}}{P^x_t}\right)^{-\eta_x}$$

With the corresponding CES price indices given by:

$$P_t = \left[\gamma_c {P_{1t}}^{1-\eta_c} + (1-\gamma_c){P_{2t}}^{1-\eta_c}\right]^{1-\eta_c}$$

$$P^x_t = \left[\gamma_x {P^x_{1t}}^{1-\eta_x} + (1-\gamma_x){P^x_{2t}}^{1-\eta_x}\right]^{1-\eta_x}$$

6.3 Nondurable Goods Production Nondurable goods are produced by perfectly competitive firms using a technology linear in labor. The problem of a representative firm located in region r is:

$$\max \quad \pi_t = P_{rt} Y_{rt} - W_{rt} L_{rt} \qquad \text{s.t.} \quad Y_{rt} = Z_t Z_r L_{rt}$$

where Z_t is aggregate TFP, Z_r is a region-specific productivity shifter, and W_{rt} is the nominal wage prevailing in region r. Under flexible prices, profit-maximizing firms set prices equal to nominal marginal costs:

$$P_{rt} = \frac{W_{rt}}{Z_t}$$

6.4 Durable Goods In each region, a continuum of perfectly competitive firms operate with a technology that transforms local nondurable goods into durable goods subject to a congestion externality (McKay and Wieland, 2022). This externality introduces decreasing returns to scale in durable goods production. The technology of a representative firm i in region r is:

$$X_{irt} = Z_r^x M_{irt} \left(\frac{X_{rt}}{\bar{X}_r} \right)^{-\zeta}$$

where Z_r^x is region-specific productivity in durable production, M_{irt} are inputs of locally-sourced nondurable goods, X_{rt} is total durable good production in region r and $\bar{X}_r$ is durable production of r on the balanced growth path. The term $\left(\frac{X_r}{\bar{X}_r}\right)^{-\zeta}$ captures the congestion externality.

The optimal price of the perfectly competitive firm is:

$$P_{rt}^x = \frac{P_{rt}}{Z_r^x} \left(\frac{X_{rt}}{\bar{X}_r} \right)^{\zeta}$$

Firms set prices equal to region-specific marginal costs. P_r^x is the price of non-durable goods produced in region r. Parameter ζ, the inverse price-elasticity of durable supply, governs the strength of the congestion externality.

6.5 Aggregate TFP growth Total factor productivity Z_t grows at a constant rate $g_{zt} = \bar{g}$.

6.6 Labor Unions There is a continuum of labor unions in each region. Each union j supplies a differentiated labor type to a local competitive firm that aggregates all labor types and sells them at price W_{rt} to local nondurable producers. The aggregator is,

$$L_{rt}^d = \left[\int_0^1 L_{rt}^d(j)^{\frac{\varepsilon^w - 1}{\varepsilon^w}} dj \right]^{\frac{\varepsilon^w}{\varepsilon^w - 1}}$$

where ε^w measures the elasticity of substitution across labor types within a region and L^d_{rt} is labor demand from nondurable producers located in region r.

The labor demand faced by union j is given by:

$$L^d_{rt}(j) = L^d_{rt}\left(\frac{W_{rt}(j)}{W_{rt}}\right)^{-\varepsilon^w}$$

The union chooses wages to maximize the average expected utility across households subject to a Calvo-type friction. The problem of a union j in region r that gets to re-optimize its wage at time t is as follows:

$$\max_{W^*_{rt}(j)} \sum_{s=0}^{\infty} (\beta\theta^w)^s \left[U_{c_{t+s}} \frac{W^*_{rt+s|t}(j)}{P_{t+s}} L^d_{rt+s|t}(j) - U_{L_{r,t+s}} L^d_{rt+s|t}(j) \right]$$

where:

$$\begin{aligned} L^d_{rt+s|t}(j) &= L^d_{rt+s} W^{\varepsilon^w}_{rt+s} W^*_{rt}(j)^{-\varepsilon^w} \\ U_{L_{rt+s}} &= \varphi_{t+s} L^{\nu}_{t+s} \gamma_\ell^{\frac{1}{\eta_\ell}} L_{rt+s}^{-\frac{1}{\eta_\ell}} \\ U_{c_{t+s}} &= \lambda_{t+s} \end{aligned}$$

where θ^{ws} is the probability that the wage contract survives s periods into the future, U_{ct} is the marginal utility of consumption and $U_{L_{rt}}$ is the marginal disutility of labor supplied to region r. In absence of an adjustment opportunity, the prevailing wage in $t+s$ stays fixed at its period t value. The labor union chooses $W^*_{rt}(j)$ such that it maximizes household utility over the expected duration of the wage contract. Since the problem of all unions adjusting in period t in region r is symmetric, we omit indexing the optimal reset wage by j. The optimal wage contract solves:

$$\begin{aligned} w^*_{rt} &= \frac{\varepsilon^w}{\varepsilon^w - 1} \frac{F_{1rt}}{F_{2rt}} \\ F_{1rt} &= \gamma_\ell^{-\frac{1}{\eta_\ell}} \varphi_t L_t^{\nu - \frac{1}{\eta_\ell}} L_{rt}^{\frac{1}{\eta_\ell}} L^d_{rt} w^{\varepsilon^w}_{rt} + \beta\theta^w \Pi^{\varepsilon^w}_{t+1} F_{1r,t+1} \\ F_{2rt} &= \lambda_t L^d_{rt} w^{\varepsilon^w}_{rt} + \beta\theta^w \Pi^{\varepsilon^w - 1}_{t+1} F_{2r,t+1} \end{aligned}$$

where lowercase letters denote real wages and Π_t denotes the gross price inflation rate. With perfectly flexible wages, optimality implies that wages equal a constant mark-up over the marginal rate of substitution. Away from this scenario, the wage mark-up will fluctuate across time and unions. This wage dispersion induces a wedge (in terms of *utils*) between

hours worked, L_{rt}, and effective hours worked, L^d_{rt}:

$$L_{rt} = s^w_{rt} L^d_{rt}$$

where s^w_{rt} measures the degree of wage dispersion across unions within a region. This wedge evolves according to:

$$\begin{aligned} s^w_{rt} &= \int_0^1 \left(\frac{W_{rt}(j)}{W_{rt}} \right)^{-\varepsilon^w} dj \\ &= (1-\theta^w) \left(\frac{W^*_{rt}}{W_{rt}} \right)^{-\varepsilon^w} + \theta^w \left(\frac{W_{rt-1}}{W_{rt}} \right)^{-\varepsilon^w} s^w_{rt-1} \\ &= (1-\theta^w) \left(\frac{w^*_{rt}}{w_{rt}} \right)^{-\varepsilon^w} + \theta^w \left(\frac{w_{rt-1}}{w_{rt}} \Pi_t \right)^{-\varepsilon^w} s^w_{rt-1} \end{aligned}$$

The real wage index for jobs located in region r satisfies:

$$w_{rt} = \left[(1-\theta^w) {w^*_{rt}}^{1-\varepsilon^w} + \theta^w \left(w_{rt-1} \Pi_t^{-1} \right)^{1-\varepsilon^w} \right]^{\frac{1}{1-\varepsilon^w}}$$

Lastly, the aggregate real wage, w_t, is characterized by the following CES index:

$$w_t = \left[\gamma_\ell w_{1t}^{1-\eta_\ell} + (1-\gamma_\ell) w_{2t}^{1-\eta_\ell} \right]^{\frac{1}{1-\eta_\ell}}$$

6.7 Monetary Policy The central bank sets the real interest rate according to the following monetary policy rule:

$$1 + r_t = (1 + r_{ss}) \left(\frac{L_t}{L} \right)^{\phi_L} \varepsilon_{rt}$$

where $\bar{r}$ is the real rate along the balanced growth path, L is aggregate labor along the balanced growth path and ε_{rt} is an exogenous shock to the real interest rate.

6.8 Market Clearing Regional goods market clearing implies:

$$Y_{rt} = C_{rt} + \frac{P^x_{rt}}{P_{rt}} X_{rt} \qquad \forall \quad r = 1, 2$$

Aggregate output, Y_t, is given by:

$$Y_t = \frac{P_{1t}}{P_t} Y_{1t} + \frac{P_{2t}}{P_t} Y_{2t}$$

Lastly, asset market clearing implies:

$$A_t = 0$$

6.9 Stationary Equilibrium We normalize all prices by the nondurable price index, P_t, and make the economy stationary by expressing all the growing variables relative to aggregate productivity. We denote with lowercase letters these relative variables and solve the model around a zero-inflation balanced growth path with $y_{ss} = l_{ss} = 1$, $w_{ss} = w_{i,ss} = 1$, and equal-sized regions, $p_{1,ss}y_{1,ss} = p_{2,ss}y_{2,ss} = .5y_{ss}$. The stationary equilibrium is defined by a sequence of productivity-adjusted quantities and wages, a sequence of goods prices $\{p_{it}, p^x_{it}\}$, plus a sequence for the real rate $\{r_t\}$ such that, given the realizations for $\{\varepsilon_\psi, \varepsilon_r\}$: (1) household utility and firm profits are maximized, (2) durable and nondurable goods markets clear, and (3) the asset market clears.

6.10 Calibration Table 2 shows the calibrated parameters. We calibrate the model to a quarterly frequency. We set the parameters for the household sector as follows. For the degree of household inattention we choose $\theta = .93$, in line with the estimates of Auclert, Rognlie, and Straub (2020). We calibrate the effective discount factor to $\frac{\beta}{1+g_z} = .99$ to match the average real rate during our sample. We calibrate the depreciation rate to match the average depreciation of the stock of private fixed assets plus durable goods in the BEA fixed asset table between 1947-2019. Operating and maintenance costs are computed as in McKay and Wieland (2022). Operating costs include expenditures in households utilities (PCE), expenditures in motor vehicle fuels and fluids (PCE), and taxes on the housing sector. We calibrate $\iota = .005$ to match the average operating costs as a share of the durable goods stock. Maintenance costs are calculated by aggregating household and motor vehicle maintenance expenditures from the Personal Consumption Expenditures (PCE) data, as well as intermediate goods and services from the housing output table. We set $\chi = .17$ to match the average ratio of maintenance costs to depreciation. The taste for durable goods, $\bar{\psi}$, is calibrated to match the average ratio of nondurable consumption to the durable stock during our sample. We set θ^d,the Calvo friction in the durable choice problem, at $\theta^d = .8$. This implies an elasticity of investment to the real rate of -4 (Gilchrist and Himmelberg (1995), Auclert and Rognlie (2018)). We set the elasticity of substitution for durables equal to $\eta_x = 7$ and the elasticity of substitution for nondurables to $\eta_c = 5$. Higher values for these elasticities

Table 2: Calibrated Parameters

Parameters		Value
$1/\sigma$	Nondurable EIS	.25
$1/\sigma^d$	Durable EIS	.25
$\bar{\psi}$	Taste for Durables	6.3
ζ	Inverse Price Elasticity of Durable Supply	.1
ι	Operating cost	0.005
χ	Maintenance cost	0.17
δ	Depreciation rate	0.017
$\frac{\beta}{(1+g_z)}$	Effective Discount factor	0.99
$\bar{g}$	Productivity Growth	.005
Y/Z	Output to TFP Ratio	1
ν	Inverse Frisch elasticity	1
ϕ_L	Taylor Rule Coefficient - Output	0
η_c	Elasticity of Substitution - Nondurables	5
η_x	Elasticity of Substitution - Durables	7
η_ℓ	Elasticity of Substitution - Labor Supply	5
ε_w	Elasticity of Substitution - Labor Demand	7
θ	Degree of Information Stickiness	.93
θ^d	Calvo Parameter - Durable Consumption	.8
θ^w	Degree of Wage Stickiness	.9
γ_x	Regional Bias in Durable Demand	.5
γ_ℓ	Regional Bias in Labor Supply	.5
γ_c	Regional Bias in Nondurable Demand	.5
Z_1^x	Productivity of Investment Sector in Region 1	1.04
Z_2^x	Productivity of Investment Sector in Region 2	.95
Z_1	Productivity of Nondurable Sector in Region 1	.98
Z_2	Productivity of Nondurable Sector in Region 2	1.01

imply stronger expenditure switching effects in response to changes in prices across regions. We choose these values to reflect the relatively lower tradability of nondurable goods which we characterize as mostly being labor services. Lastly, we set the elasticity of substitution of labor supply between regions at $\eta_\ell = 5$.

We calibrate the production side as follows. We normalize the output-to-TFP ratio to one in the stationary equilibrium. The deterministic component of TFP growth, g_Z, is calibrated to match the average quarterly growth rate of real GDP per capita over the period 1947-2019. We calibrate the inverse supply elasticity of durable goods to $\zeta = .1$ based on House and Shapiro (2008). We set the elasticity of substitution of labor demand across unions to $\varepsilon_w = 7$. This is in the middle range of values used in the literature and implies a wage markup of 15% in the steady state. We set the degree of wage stickiness to $\theta_w = .9$, which implies that unions re-optimize wages on average every 10 quarters.

We set the productivity parameters in the durables sector, Z_1^x and Z_2^x, and nondurables sector, Z_1 and Z_2, to generate heterogeneity in the industrial composition between regions.

Specifically, we choose Z_1^x such that the participation of region 1 in total durable production is 55%. Then Z_2^x is set to guarantee that the relative aggregate price of durable goods is equal to one in steady state (that is, $p_{ss}^x = 1$). Lastly, we choose productivity values in the nondurable sector to ensure that both regions have the same level of steady-state output. As a result, Region 1 is relatively specialized in the production of durable goods, while Region 2 is relatively specialized in nondurable goods.

6.11 Matching the Data We model boom-bust recessions as the result of a fake news shock to households' taste for durables ψ_t. We assume a two-wave information structure to construct the fake news shock which we briefly describe next. Appendix E provides further details. In period $t = 0$, news arrive about an increase in ψ_t beginning in period $s > 0$. This constitutes the first wave of information. Starting at $t = 0$, agents gradually learn of this future path at a rate $1 - \theta$, where $\theta \in [0, 1)$ captures the degree of information stickiness. Households who update their information set expect a higher return on durable consumption beginning in period s. In anticipation of this, they build up their durable stock, which boosts output and the investment share. Because households update their information set infrequently, this results in a gradual build-up of the aggregate durable stock. In period $z < s$, a second wave of information begins. In each period $t \geq z$, a share $1 - \theta$ of agents who expect ψ_t to increase at time s learn that the taste for durables will not change. Households who receive this new piece of information find themselves with an excess stock of durable goods they wish to liquidate. Initially, this results in a deceleration of the process of accumulation of durable goods at the aggregate level. Accumulation decelerates but does not come to an immediate halt. This is because the stock of households that still expect ψ_t to increase is large relative to those who learn of the second wave. As the second wave of information disseminates, the liquidation of excess stocks generates a bust in output, a sustained decrease in the investment share, and a slow recovery of the economy relative to the elevated pretrend.

We generate a V-shaped recession pattern for the pre-1990 and Covid recessions using a temporary monetary policy shock. In practice, any transitory shock in the model will generate a V-shaped recession in the model. We choose monetary policy shocks because they are considered important drivers of pre-1990 recessions (Coibion, 2012; Romer and Romer, 2023). One can interpret the Covid recession as a period in which the price of consumption today relative to the future—the real interest rate—was high due to health risk associated with certain types of consumption. An unexpected increase in the real interest rate induces intertemporal substitution away from consumption and investment that cause output to temporarily drop below trend before bouncing back once the monetary shock recedes. Due

to the importance of intertemporal motives in such shocks, we call them intertemporal shocks.

We model the relatively weaker trend of the more exposed region as the result of a change in γ_x, the regional bias in durable demand. A change in γ_x triggers the reallocation of durable household purchases across regions, holding total purchases fixed. This allows us to compute the path of γ_x that rationalizes a given time path for the differential in growth rates between regions, absent any national recessions. Given the linearity of the model, we can then additively feed in the aggregate shocks that generate national recessions, as discussed above.

We solve the model in the sequence space using the methods developed by Auclert, Bardóczy, Rognlie, and Straub (2021). We then simulate the path of the economy for different shocks that generate heterogeneous recessions and recovery dynamics. Importantly, we keep all the parameters that affect the steady state fixed across recessions, which means that all differential dynamics can be traced back to the type of shock that hits the economy. Next, using the simulated data, we compute the model analog of the relative recession impulse response function in Figure 4. Specifically, we estimate the same regression (21) using a Bartik instrument in the model as we do in the data.

We estimate the path of γ_x[14] and the parameters that govern the shock processes to match the impulse responses in Figure 4. We assume AR(1) processes for both the monetary policy and taste for durables shocks:

$$\varepsilon_t^r = \rho_r e_{t-1}^r + e_t^r, \qquad e_t^r \overset{\text{iid}}{\sim} N(0, \sigma_r^2)$$
$$\varepsilon_t^\psi = \rho_\psi \varepsilon_{t-1}^\psi + e_t^\psi, \qquad e_t^\psi \overset{\text{iid}}{\sim} N(0, \sigma_\psi^2)$$

where ρ_x and σ_x are the persistence and standard deviation of shock x. For each alternative explanation, we estimate the persistence of the shock processes. We also estimate the timing of each information wave, s and z. Let $f(\mathbf{\Gamma})$ denote the vector of model-implied cross-sectional impulse responses evaluated at parameter values $\mathbf{\Gamma}$. Let $\boldsymbol{\beta}$ denote the vector of empirical impulse responses estimated in Section 5 .

We estimate the parameters by solving:

$$\hat{\mathbf{\Gamma}} = argmin_{\mathbf{\Gamma}} \quad (\boldsymbol{\beta} - f(\mathbf{\Gamma}))' \mathcal{W}^{-1} (\boldsymbol{\beta} - f(\mathbf{\Gamma}))$$

where $\mathcal{W}$ is a diagonal weighting matrix containing the estimated variances of the elements of β and $\hat{\mathbf{\Gamma}}$ is the value of the parameter vector that minimizes the weighted distance between the model-implied and empirical impulse responses. Table 3 presents the estimated parameters.

[14] In particular, we estimate a single path of γ_x, which means that both types of recessions occur around the same relatively weaker pretrends.

Table 3: Estimated Patameters

		Value
Intertemporal *vs.* Boom-bust		
Persistence of Monetary Policy Shock	ρ_r	.946
Persistence of Taste Shock	ρ_ψ	1
Fake News Date	s	17
Second Wave of News Date	z	14

Throughout the exercise, we calibrate the standard deviation of each aggregate shock to match the average depth of the recession in our sample, which is 5.3%.

In Figure 8 we show that when recessions are driven by different aggregate shocks, the model can generate impulse response functions that closely match the cross-sectional estimates for the 1990-1, 2001, and 2007-9 recessions. Regions that are more exposed to the boom-bust recession have a relatively higher above-trend growth rate due to the prior boom. Recoveries in these regions will look weaker because they can never catch up to this inflated pre-recession trend. The estimation favors a slight downward trend in the more exposed regions due to the negative pretrend prior to the boom.

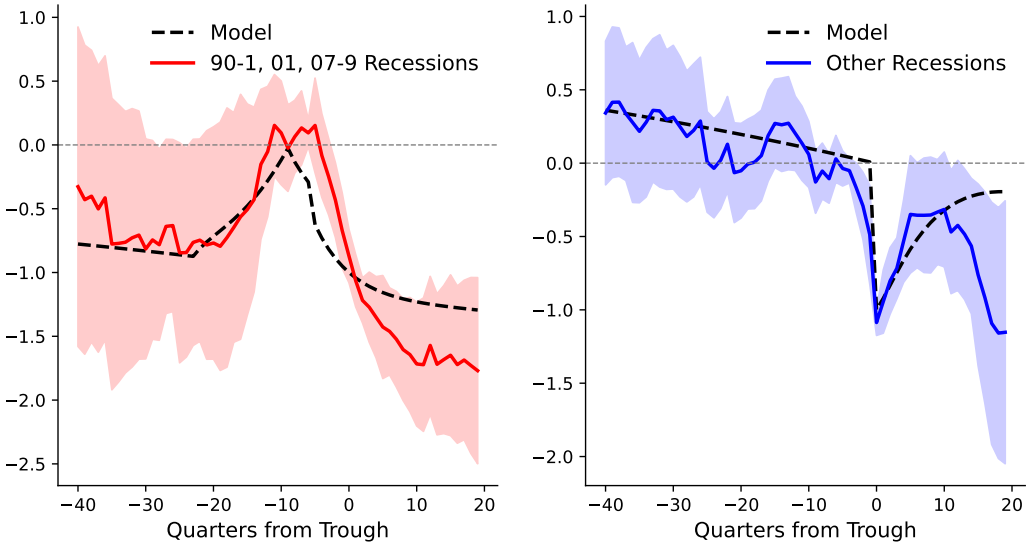
Intertemporal shocks, such as a monetary policy shock, do not affect the pretrend and represent only a temporary departure from it. Thus, these shocks predict a V-shaped relative recovery in the absence of pretrends. However, our estimates point to a significantly lower trend growth rate in the more exposed States which helps the model match the observed recovery in the data.

6.12 Policy Analysis We next show that distinguishing between boom-bust cycles and intertemporal shocks is important for policy in the recovery. Specifically, for each recession we solve for the path of the real interest rate R^* that closes the aggregate output gap starting at the recession trough. We scale each recession to a 5.2% peak-to-trough decline in aggregate output, which is the empirical average for both types of recessions. Figure 9 shows the 10-year R^* zero coupon rate computed using the expectations hypothesis.

The model implies that long-run R^* persistently declines following a boom-bust cycle, reaching a trough of 20 basis points. In contrast, the V-shaped monetary policy recession predicts an increase of long-run R^* of several basis points.

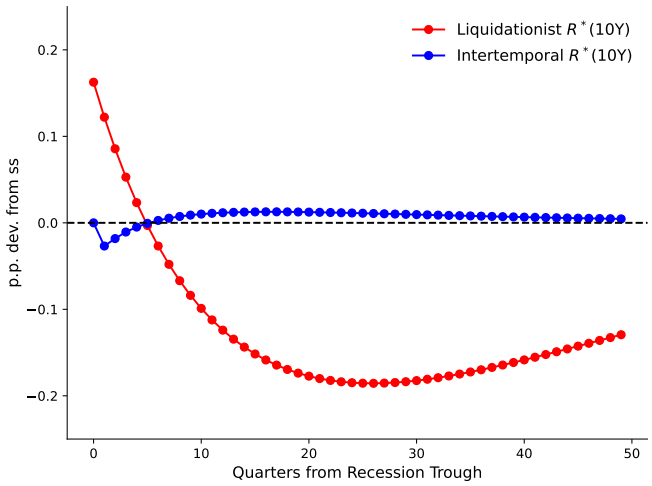
These distinct patterns for R^* are due to the dynamics of the capital stock. Following a boom-bust cycle, the economy finds itself with excess capital. Investment is persistently low and aggregate demand needs to be supported using a low real interest rate. However, an

Figure 8: Matching the Recession IRF - Liquidationist *vs.* Intertemporal



intertemporal shock that delayed investment and consumption implies that the capital stock is below steady-state in the recovery. Thus, there is pent-up demand for investment which pushes up R^* in the recovery path. This implies that the degree of monetary accommodation required in the recovery depends importantly on the type of shock.

Figure 9: R_t^* - Liquidationist *vs.* Intertemporal



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 suggests the recovery patterns are due to different shocks: Boom-bust cycles in the 1990-1, 2001, and 2007-9 recessions give rise to weak recoveries and intertemporal shocks in the pre-1990 and 2020 recessions predict strong recoveries. Our model that quantitatively matches these patterns shows that recessions caused by boom-bust cycles require extended periods of accommodative policy relative to recessions caused by intertemporal shocks.

References

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

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89(5), 2375–2408.

Auclert, A., and M. Rognlie (2018): "Inequality and aggregate demand," Discussion paper, National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Discussion paper, National Bureau of Economic Research.

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

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.

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., and C. K. Wolf (2022): "Demand Composition and the Strength of Recoveries," .

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.

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

Carroll, C. D., E. Crawley, J. Slacalek, K. Tokuoka, and M. N. White (2020): "Sticky expectations and consumption dynamics," *American economic journal: macroeconomics*, 12(3), 40–76.

Coibion, O. (2012): "Are the effects of monetary policy shocks big or small?," *American Economic Journal: Macroeconomics*, 4(2), 1–32.

Dupraz, S., E. Nakamura, and J. Steinsson (2023): "A Plucking Model of Business Cycles," .

Evans, G. W., and G. Ramey (1992): "Expectation calculation and macroeconomic dynamics," *The American Economic Review*, pp. 207–224.

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.

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*, 2004(1), 265–319.

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

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.

Gilchrist, S., and C. P. Himmelberg (1995): "Evidence on the role of cash flow for investment," *Journal of monetary Economics*, 36(3), 541–572.

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

House, C. L., and M. D. Shapiro (2008): "Temporary investment tax incentives: Theory with evidence from bonus depreciation," *American Economic Review*, 98(3), 737–768.

Howard, G., R. Martin, and B. A. Wilson (2011): "Are recoveries from banking and financial crises really so different?," .

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.

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 J. F. Wieland (2022): "Forward guidance and durable goods demand," *American Economic Review: Insights*, 4(1), 106–122.

Orchard, J., V. A. Ramey, and J. F. Wieland (2023): "Micro MPCs and macro counterfactuals: the case of the 2008 rebates," Discussion paper, National Bureau of Economic Research.

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.

Romer, C. D., and D. H. Romer (2023): "Presidential address: Does monetary policy matter? The narrative approach after 35 years," *American Economic Review*, 113(6), 1395–1423.

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 Simple Model Extention

A.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,

$$Y_{i,g,t} = (1-\mu)C^o_{i,g,t} + \mu C^m_{i,g,t} \tag{22}$$

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,

$$C^m_{i,g,t} = \alpha_{gt}\left(\frac{P_{i,g,t}}{P_{it}}\right)^{-\sigma_g} C^m_{it}. \tag{23}$$

Given optimal prices (8), we get

$$C^m_{g,t} = \alpha_g Y_t \frac{Z_{t,g}}{Z_t} \tag{24}$$

$$C^o_{g,t} = \alpha_g \left(\frac{\lambda_t}{Z_t}\right)^{-\frac{1}{\sigma_g}} Z_{t,g} \tag{25}$$

Market clearing implies,

$$Y_{g,t} = (1-\mu)\alpha_g \left(\frac{\lambda_t}{Z_t}\right)^{-\frac{1}{\sigma_g}} Z_{t,g} + \alpha_g \mu Y_t \frac{Z_{t,g}}{Z_t} \tag{26}$$

Total local output is,

$$Y_{i,t} = \sum_{g=1}^{G}(1-\mu)\alpha_{ig}e^{\hat{\alpha_{it}}}\left(\alpha_g\left(\frac{\lambda_t}{Z_t}\right)^{-\sigma_g} Z_{t,g}\right) + \mu C^m_{i,t} \tag{27}$$

$$= e^{\hat{\alpha_{it}}}\sum_{g=1}^{G}\alpha_{ig}Y_{g,t} - e^{\hat{\alpha_{it}}}\mu Y_t\frac{Z_{it}}{Z_t} + \mu Y_{i,t} \tag{28}$$

$$= e^{\hat{\alpha_{it}}}\left[\frac{1}{1-\mu}\sum_{g=1}^{G}\alpha_{ig}Y_{g,t} - \frac{\mu}{1-\mu}Y_t\frac{Z_{it}}{Z_t}\right] \tag{29}$$

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

$$\hat{y}_{g,t} = -\frac{1}{\sigma_g}(1-\mu)\hat{\lambda}_t + \mu\hat{y}_t \tag{30}$$

$$\hat{y}_t = -\sum_{g=1}^{G}\left(\alpha_g \frac{1}{\sigma_g}\right)\hat{\lambda}_t \tag{31}$$

Local output growth is

$$\Delta^h y_{i,t,t+h} \approx \hat{\alpha_{it}} + \frac{1}{1-\mu}\sum_{g=1}^{G}\mu_{i,g,t}\Delta^h y_{g,t+h} + \frac{\mu}{1-\mu}\Delta^h y_{t+h} \tag{32}$$

The aggregate business cycle at t is unchanged,

$$\hat{y}_t + \xi\sum_{s=0}^{t-1}\hat{y}_s. \tag{33}$$

Therefore, the cyclical Bartik is,

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

Under conditions from Proposition 1, we get a first stage and reduced form coefficient equal to,

$$\gamma_1 = \frac{1}{1-\mu} \tag{35}$$

$$\delta_1 = \frac{1}{1-\mu}\frac{\Delta^{T+h-P}\hat{y}_{T+h} + \xi\sum_{s=0}^{T+h-P-1}\hat{y}_{P+s}}{\Delta^{T-P}\hat{y}_T + \xi\sum_{s=0}^{T-P-1}\hat{y}_{P+s}} \tag{36}$$

The IV coefficient is the ratio δ_1/γ_1 and unchanged from Proposition 1.

Intuitively, the Bartik underpredicts the recession in more exposed regions because it captures only the hand-to-mouth response in the average region. The Bartik exposure therefore gets scaled up in the first stage (and reduced form), and the common hand-to-mouth response in the average region gets absorbed by the recession fixed effect.

A.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,

$$\frac{P_{i,g,t}}{P_{g,t}} = \frac{\bar{W}_{ig}}{\bar{W}_g} \left(\frac{Y_{i,t}/Z_{i,t}}{Y_t/Z_t} \right)^{\kappa_I} \tag{37}$$

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

$$\frac{P_{g,t}}{P_t} = \frac{\bar{W}_g}{\bar{W}} \left(\frac{Y_{g,t}/Z_{g,t}}{Y_t/Z_t} \right)^{\kappa_G} \frac{Z_t}{Z_{g,t}} \tag{38}$$

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

Then local output of good g is equal to

$$Y_{i,g,t} = \alpha_{igt} \left(\frac{\bar{W}_g}{\bar{W}} \right)^{-\eta} \left(\frac{Y_{i,t}/Z_{i,t}}{Y_t/Z_t} \right)^{-\eta\kappa_I} Y_{g,t} \tag{39}$$

and aggregate output of good g is equal to,

$$Y_{g,t} = Z_{g,t} \left[\alpha_g \left(\frac{\bar{W}_g}{\bar{W}} \right)^{-1} \left(\frac{Y_t}{Z_t} \right)^{\kappa_G} \lambda_t^{-1} Z_t \right]^{\frac{1}{\sigma_g - \kappa_G}} \tag{40}$$

Local output growth has a Bartik structure

$$\underbrace{\Delta^h y_{i,t+h}}_{\text{Local Output Growth}} = \frac{1}{1+\eta\kappa_I} \underbrace{\Delta^h \hat{\alpha}_{i,t+h}}_{\text{Local Demand Shifter}} + \frac{1}{1+\eta\kappa_I} \underbrace{\sum_{g=1}^{G} \mu_{i,g,t} \Delta^h y_{g,t+h}}_{\text{Bartik Predicted Growth, } y^B_{i,t,t+h}} + \frac{\eta\kappa_I}{1+\eta\kappa_I} \underbrace{\Delta^h y_{t+h}}_{\text{Aggregate Growth}} \tag{41}$$

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

$$\hat{y}_{g,t} + \xi \sum_{s=0}^{t-1} \hat{y}_{g,s} = \frac{\kappa_G}{\tilde{\sigma}_g} \left(\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s \right) - \frac{1}{\tilde{\sigma}_g} \left(\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right) \tag{42}$$

The cyclical component of output is,

$$\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s = - \left(\frac{\sum_{g=1}^{G} \mu_{gt} \frac{1}{\tilde{\sigma}_g}}{1 - \kappa_G \sum_{g=1}^{G} \mu_{gt} \frac{1}{\tilde{\sigma}_g}} \right) \left(\hat{\lambda}_t + \xi \sum_{s=0}^{t-1} \hat{\lambda}_s \right) \tag{43}$$

which implies,

$$\hat{y}_{g,t} + \xi \sum_{s=0}^{t-1} \hat{y}_{g,s} = \frac{1}{\tilde{\sigma}_g} \left[\kappa_G + \frac{1 - \kappa_G \sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}}{\sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}} \right] \left(\hat{y}_t + \xi \sum_{s=0}^{t-1} \hat{y}_s \right) \tag{44}$$

The cyclical component of Bartik predicted growth is,

$$y^{BC}_{i,t,t+h} = \underbrace{\left(\frac{\sum_{g=1}^{G} \mu_{i,g,t} \frac{1}{\tilde{\sigma}_g}}{\sum_{k=1}^{G} \mu_{kt} \frac{1}{\tilde{\sigma}_k}}\right)}_{\text{Relative Regional Business Cycle Exposure}} \times \underbrace{\left(\Delta^h \hat{y}_{t+h} + \xi \sum_{s=0}^{h-1} \hat{y}_{t+s}\right)}_{\text{Aggregate Business Cycle}} \tag{45}$$

Under conditions from Proposition 1, we get a first and reduced form coefficient equal to,

$$\gamma_1 = \frac{1}{1+\eta\kappa_I} \tag{46}$$

$$\delta_1 = \frac{1}{1+\eta\kappa_I} \frac{\Delta^{T+h-P} \hat{y}_{T+h} + \xi \sum_{s=0}^{T+h-P-1} \hat{y}_{P+s}}{\Delta^{T-P} \hat{y}_T + \xi \sum_{s=0}^{T-P-1} \hat{y}_{P+s}} \tag{47}$$

The IV coefficient is the ratio δ_1/γ_1 and unchanged from Proposition 1.

The Bartik now overpredicts the recession in more exposed regions because it captures only the price response in the average region. The Bartik exposure therefore gets scaled down in the first stage (and reduced form). Good-specific price responses do not pose an issue as these only affect how the aggregate shock affects the Bartik, but not how the Bartik differentially affects regions.

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

$$\hat{r}_t = \phi \hat{y}_t + \hat{\nu}_t \tag{48}$$

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

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

$$\hat{\lambda}_t = \frac{1}{1+(\phi-\xi)\sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}} \sum_{s=0}^{\infty} \left(\frac{1}{1+(\phi-\xi)\sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}}\right)^s (\hat{\beta}_{t+s} + \hat{\nu}_{t+s}) + \lim_{T\to\infty} \left(\frac{1}{1+(\phi-\xi)\sum_{g=1}^{G} \frac{\alpha_g}{\sigma_g}}\right)^T \mathbb{E}_t \lambda_T \tag{49}$$

Since no other equation changes, Proposition 1 goes through unchanged.

B Trend Mismeasurement

B.1 Trend Mismeasurement in the Cross-Region Approach The IV strategy discussed in Section 2 requires an estimate of the Bartik predicted recession depth which we denote by $\tilde{y}^{BC}_{i,P,T}$:

$$\tilde{y}^{BC}_{i,P,T} = \sum_{g} \mu_{i,g,P} \Delta^{P-T} \tilde{\hat{y}}_{g,P} \tag{50}$$

We construct the estimated cyclical component of sectoral output as

$$\tilde{\hat{y}}_{g,t} = y_{g,t} - \tilde{\bar{y}}_{g,t} \tag{51}$$

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

$$\bar{y}_{g,t} = \gamma_g t + u_{g,t} + \eta_g u_t + u_t \tag{52}$$

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, $E(u_{g,t}) = E(u_t) = 0$. The parameter $\eta_g \sim (0, \sigma^2_\eta)$ governs the differential exposure of sector g to the aggregate trend shock.

Throughout, we assume that sector-specific shocks are (i) uncorrelated across goods, (ii) uncorrelated with the aggregate shock, and (iii) orthogonal to sectoral trend parameters and specialization patterns, $\{\gamma_g, \eta_g\}_g$ and $\{\mu_{i,g}\}_i$.

Assumption 1

$$\begin{aligned}
&u_{g,t} \perp u_{g',s} \quad \forall\, g, g' \quad \textit{and} \quad \forall\, s \\
&u_{g,t} \perp u_s \quad \forall g \in G \quad \textit{and} \quad s \in T \\
&u_{g,t} \perp \{\{\gamma_g, \eta_g\}_g, \{\mu_{i,g}\}_i\} \\
&E(u_{g,t}) = E(u_t) = 0 \quad \forall\, g, s
\end{aligned}$$

Trend mismeasurement $\bar{e}_{g,t}$ after fitting a linear trend is

$$\bar{e}_{g,t} = (\hat{\gamma}_g - \gamma_g)t + u_{g,t} + \eta_g u_t + u_t \tag{53}$$

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

$$\bar{e}_{g,t} \to u_{g,t} + \eta_g u_t + u_t \tag{54}$$

Then, the estimated Bartik instrument $\tilde{y}^{BC}_{i,P,T}$ can be expressed as

$$\tilde{y}^{BC}_{i,P,T} = \sum_g \mu_{i,g,P}[\Delta^{P-T}\hat{y}_{g,P} - \Delta^{P-T}\bar{e}_{g,P}] \tag{55}$$

and the IV estimator of the *relative recession impulse response function* as

$$\hat{\beta}^{IV}_h = \frac{Cov_i\Big(\Delta^{T+h-P}y_{i,T+h}, \tilde{y}^{BC}_{i,P,T}\Big)}{Cov_i\Big(\Delta^{P-T}y_{i,P}, \tilde{y}^{BC}_{i,P,T}\Big)}. \tag{56}$$

Under the assumptions stated in Proposition 1 and further assuming that trend mismeasurement $\Delta^{P-T}\bar{e}_{g,P}$ over the recession is uncorrelated with the cyclical component of sectoral output, it follows that:

$$\hat{\beta}^{IV}_h = \frac{Cov_i\Big(\Delta^{T+h-P}\hat{y}_{i,T+h}, y^{BC}_{i,P,T}\Big) - Cov_i\Big(\Delta^{T+h-P}\bar{y}_{i,T+h}, \sum_g \mu_{i,g,t}\Delta^{P-T}\bar{e}_{g,P}\Big)}{Cov_i\Big(\Delta^{P-T}\hat{y}_{i,P}, y^{BC}_{i,P,T}\Big) - Cov_i\Big(\Delta^{P-T}\bar{y}_{i,P}, \sum_g \mu_{i,g,t}\Delta^{P-T}\bar{e}_{g,P}\Big)} \tag{57}$$

Assumption 2

$$Cov_i\Big(\sum_g \mu_{i,g,t}\Delta^{T+s-P}\hat{y}_{g,T+s}, \sum_g \mu_{i,g,t}\Delta^{P-T}\bar{e}_{g,P}) = 0 \quad \forall \quad s$$

Then, it follows that $\hat{\beta}^{IV}_h \to \beta_h$ if, and only if

$$Cov_i\Big(\Delta^{T+s-P}\bar{y}_{i,T+s}, \sum_g \mu_{i,g,t}\Delta^{T-P}\bar{e}_{g,T}\Big) \to 0 \quad \text{for } s = 0, h \tag{58}$$

where the last expression uses $\Delta^{P-T}\bar{e}_{g,P} = -\Delta^{T-P}\bar{e}_{g,T}$. This is the identifying assumption for trend mismeasurement to wash-out in the regional design. The covariance in (58) can be re-written as:

$$Cov_g\Big(\Delta^{D}\bar{e}_{g,T}, \sum_i \mu_{i,g}\Delta^{D+s}\bar{y}_{i,T+s}\Big) \to 0 \tag{59}$$

where $D = T - P$ denotes recession duration. For $T \to \infty$, replacing for the definitions of $\bar{y}_{i,T+s}$ and $\bar{e}_{g,t}$ yields:

$$Cov_g\Big(\Delta^{D}u_{g,T} + \eta_g\Delta^{D}u_T + \Delta^{D}u_T, \\ \sum_{g'}\Theta_{g,g'}\Big\{\gamma_{g'} \times (D+s) + \Delta^{D+s}u_{g',T+s} + \eta_{g'}\Delta^{D+s}u_{T+h} + \Delta^{D+s}u_{T+h}\Big\}\Big) \to 0$$

where $\Theta_{g,g'} = \sum_i \mu_{i,g}\mu_{i,g'}$. This covariance can be further broken up into the following six

terms:

$$C_1 = (D+s)\sum_{g'} \gamma_{g'} E_g(\Theta_{g,g'} \Delta^D u_{g,T}) \tag{60}$$

$$C_2 = \Delta^D u_T (D+s) \times \sum_{g'} \gamma_{g'} \Big[E_g(\Theta_{g,g'} \eta_g) - \bar{\eta} E_g(\Theta_{g,g'}) \Big] \tag{61}$$

$$C_3 = E_g(\Delta^D u_{g,T} \Delta^{D+s} u_{g,T+s}) \times E_g\Big(\sum_i \mu_{i,g}^2 \Big) \tag{62}$$

$$C_4 = \Delta^{D+s} u_{T+s} \times E_g\Big(\Delta^D u_{g,T} \sum_{g'} \Theta_{g,g'} \eta_{g'} \Big) \tag{63}$$

$$C_5 = \Delta^D u_T \times \sum_{g'} \Delta^{D+s} u_{g',T+s} \Big[E_g\Big(\Theta_{g,g'} \eta_g \Big) - \bar{\eta} E_g\Big(\Theta_{g,g'} \Big) \Big] \tag{64}$$

$$C_6 = \Delta^D u_T \Delta^{D+s} u_{T+s} \times \underbrace{\sum_{g'} \eta_{g'} \Big[E_g\Big(\Theta_{g,g'} \eta_g \Big) - \bar{\eta} E_g\Big(\Theta_{g,g'} \Big) \Big]}_{Var_i(\eta_i)} \tag{65}$$

where $\bar{\eta} = \frac{1}{G}\sum_g \eta_g$ and $\eta_i = \sum_g \mu_{i,g,t} \eta_g$. We omit the covariance terms that involve the average effect of the aggregate trend shock as those are differenced-out. First, C_1 and C_4 converge to zero by Assumption 1—trend shocks have mean zero and are orthogonal to sectoral parameters. Second, for C_3 to converge to zero we impose the following assumption on regional output composition.

Assumption 3

$$E_g\Big(\sum_i \mu_{i,g,t}^2 \Big) \to 0 \quad as \quad G \to \infty$$

This asymptotic condition guarantees that sectoral-level trend shocks *wash-out* in the cross-section and can be thought of as a bound on regional specialization. In other words, we require that regions are sufficiently diversified across sectors g. This is a necessary condition.

Lastly, for C_2, C_5 and C_6 to converge to zero it suffices to establish conditions under which the term in brackets—$E_g\Big(\Theta_{g,g'} \eta_g \Big) - \bar{\eta} E_g\Big(\Theta_{g,g'} \Big)$—cancels out, although there exist weaker conditions specific to each covariance term that could make orthogonality hold. Below we establish a condition on sectoral loadings η_g that guarantees that these covariances collapse to zero:

Assumption 4 *The sectoral-level loadings η_g are independently and identically distributed across goods with mean $\bar{\eta}$ and variance σ_η^2:*

$$\eta_g \overset{iid}{\sim} (\bar{\eta}, \sigma_\eta^2).$$

A corollary of Assumption 4 is that regions are in expectation homogeneously exposed to the aggregate trend shock—$E(\eta_i) = \sum_g \mu_{i,g,t} E(\eta_g) = \bar{\eta}$.

B.2 Trend Mismeaurement 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

$$\bar{y}_t^e = \bar{y}_t + \bar{e}_t \tag{66}$$

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

$$\hat{y}_t^e = y_t - \bar{y}_t^e \tag{67}$$

$$= \hat{y}_t - \bar{e}_t \tag{68}$$

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

$$\Delta^h \hat{y}_{t+h}^e = \hat{y}_{t+h}^e - \hat{y}_t^e = \Delta^h \hat{y}_{t+h} - \Delta^h \bar{e}_{t+h} \tag{69}$$

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

$$\hat{\beta}_h^{ts} = \frac{Cov_j(y_{T_j+h} - y_{P_j}, \hat{y}_{P_j}^e - \hat{y}_{T_j}^e)}{Var_j(\hat{y}_{P_j}^e - \hat{y}_{T_j}^e)} \tag{70}$$

$$= -\frac{Cov_j(\Delta^{T_j+h-P_j} \hat{y}_{T_j+h} + \Delta^{T_j+h-P_j} \bar{y}_{T_j+h}, \Delta^{T_j-P_j} \hat{y}_{T_j} - \Delta^{T_j-P_j} \bar{e}_{T_j})}{Var_j(\Delta^{T_j-P_j} \hat{y}_{T_j}^e)} \tag{71}$$

Further assuming that the trend and cyclical component of aggregate output are orthogonal to each other we get:

$$\hat{\beta}_h^{ts} = -\frac{Cov_j(\Delta^{T_j+h-P_j}\hat{y}_{T_j+h}, \Delta^{T_j-P_j}\hat{y}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} + \frac{Cov_j(\Delta^{T_j+h-P_j}y_{T_j+h}, \Delta^{T_j-P_j}\bar{e}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} \tag{72}$$

Proposition 2

$$\beta_h^{ts} \to -\frac{Cov_j(\Delta^{T_j+h-P_j}\hat{y}_{T_j+h}, \Delta^{T_j-P_j}\hat{y}_{T_j})}{Var_j(\Delta^{T_j-P_j}\hat{y}^e_{T_j})} = \beta_h^{ts} \times \kappa \tag{73}$$

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

The assumption that $Cov_j(\Delta^{T_j+h-P_j}y_{T_j+h}, \Delta^{T_j-P_j}\bar{e}_{T_j}) = 0$ is strong. Even if the estimated trend $\bar{y}_t^e$ 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

$$\bar{y}_t = \gamma t + \upsilon_t \tag{74}$$

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

$$\bar{y}_t^e \to \hat{\gamma} t, \tag{75}$$

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 $\bar{y}_t^e \perp \{\hat{y}_s\}_s$, the covariance between output growth over the recovery and measurement error reduces to

$$Cov_j(\Delta^{D_j+h}\bar{y}_{T_j+h}, \Delta^{D_j}\bar{e}_{T_j}) = Cov_j(\gamma(D_j + h) + \Delta^{D_j+h}\upsilon_{T_j+h}, (\hat{\gamma} - \gamma)(D_j) + \Delta^{D_j}\upsilon_{T_j}) \tag{76}$$

$$= Cov_j(\gamma(D_j + h) + \Delta^{D_j+h}\upsilon_{T_j+h}, (\hat{\gamma} - \gamma)D_j + \Delta^{D_j}\upsilon_{T_j}) \tag{77}$$

$$= \gamma(\hat{\gamma} - \gamma)Var_j(D_j) + Cov_j(\Delta^{D_j+h}\upsilon_{T_j+h}, \Delta^{D_j}\upsilon_{T_j}) \tag{78}$$

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

$$\Delta^{D_j+h} \upsilon_{T_j+h} = (\rho_\upsilon^{D_j+h} - 1)\upsilon_{P_j} + \sum_{s=1}^{D_j+h} \rho_\upsilon^{D_j+h-s} \eta_{P_j+s}, \tag{79}$$

$$\Delta^{D_j} \upsilon_{T_j} = (\rho_\upsilon^{D_j} - 1)\upsilon_{P_j} + \sum_{s=1}^{D_j} \rho_\upsilon^{D_j-s} \eta_{P_j+s} \tag{80}$$

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

C Figures

Figure 10: Impulse Response Function of State Personal Income by Recession Group Double-Clustering by State and Recession

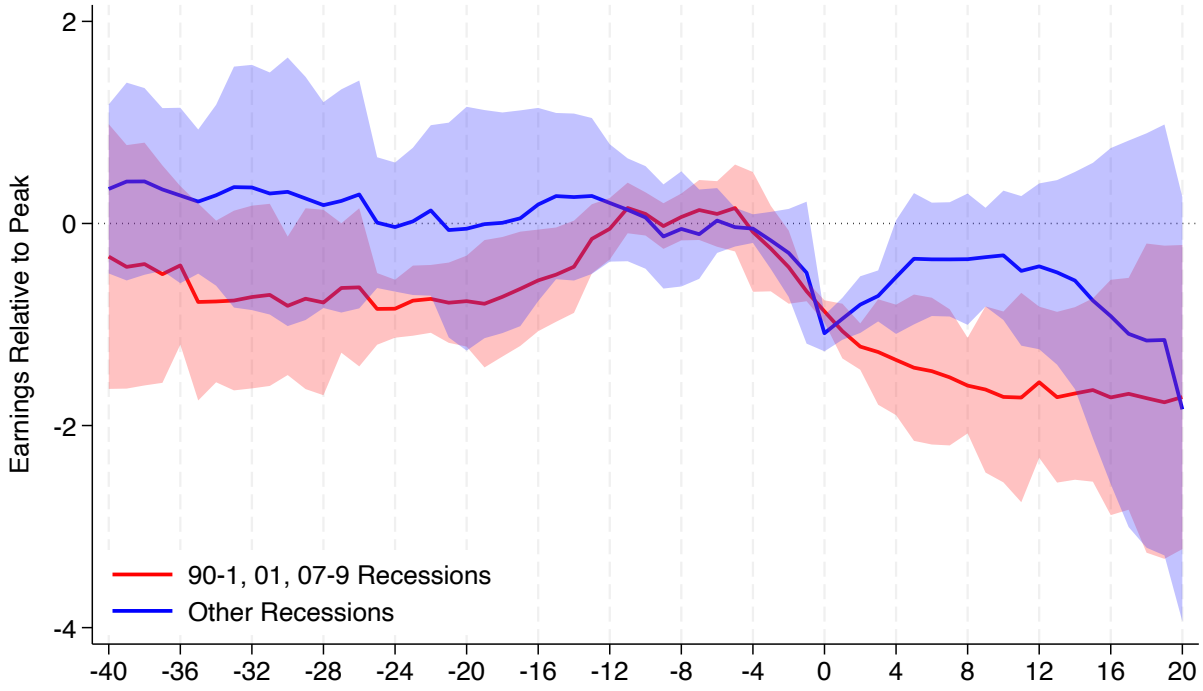
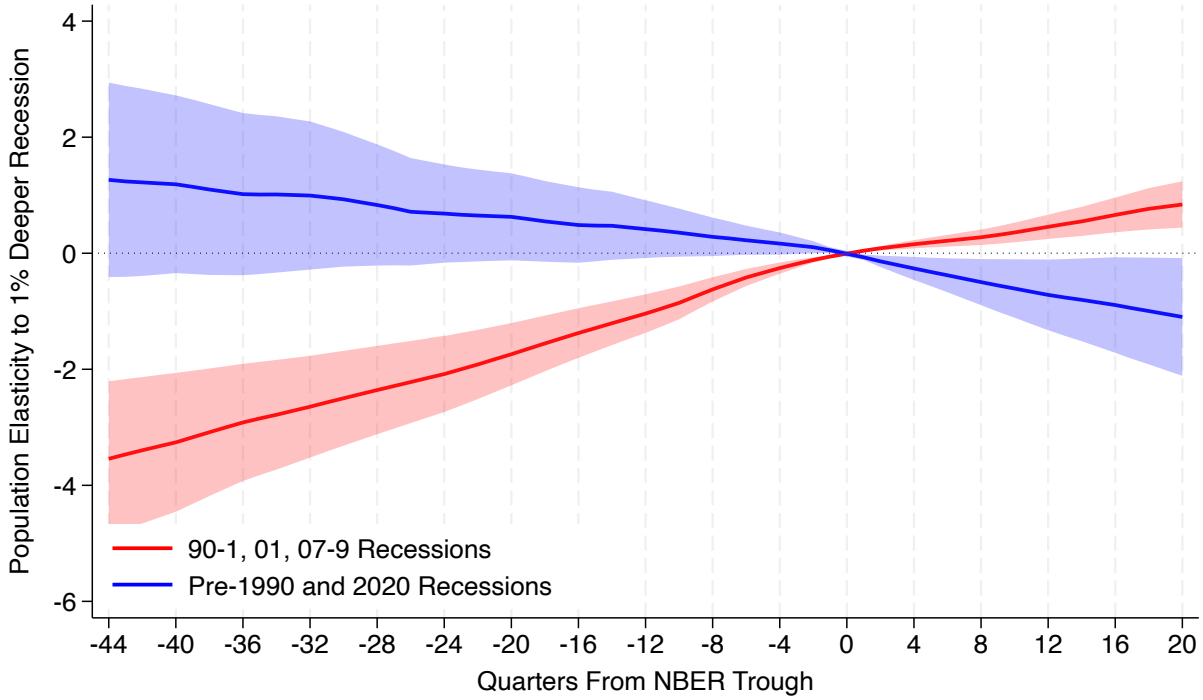


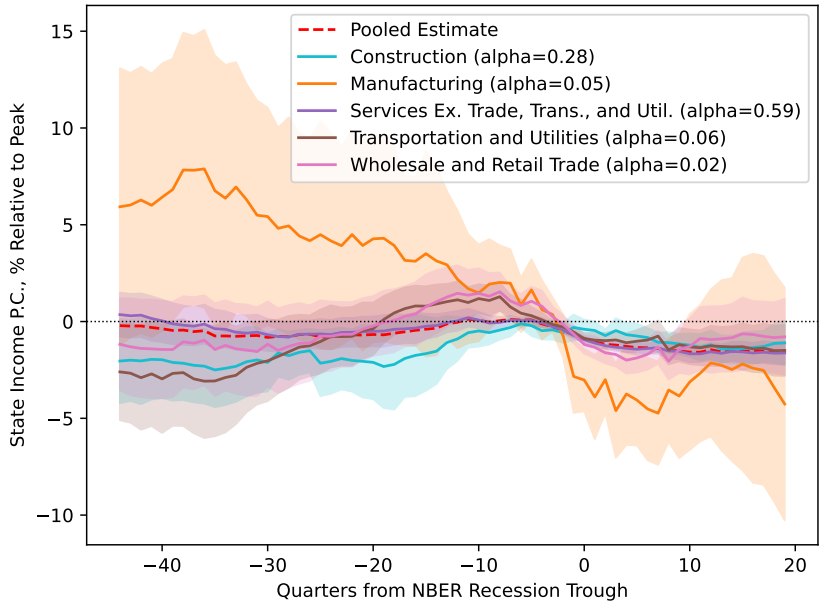
Figure 11: Recession Impulse Response Function for Population



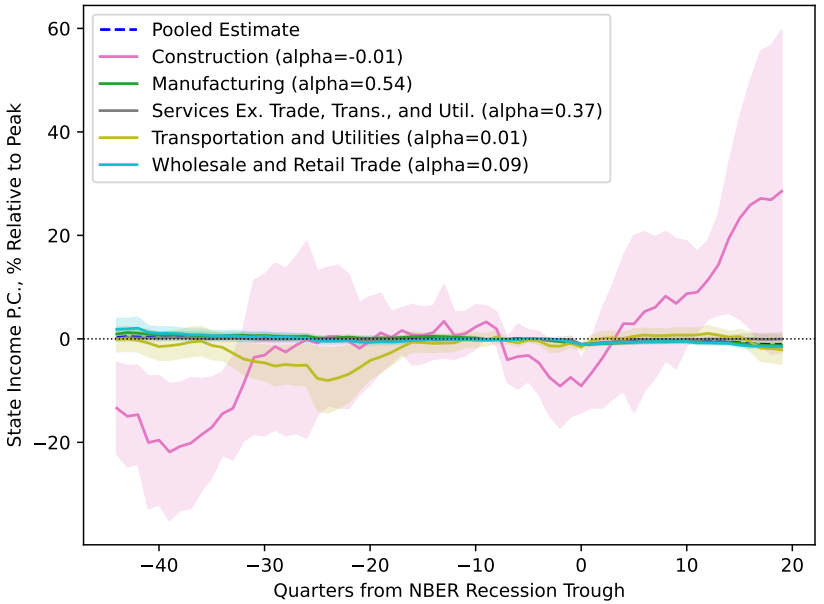
D Rotemberg Weights

Figure 12: Impulse Response Functions based on Major Industry and Recession Variation

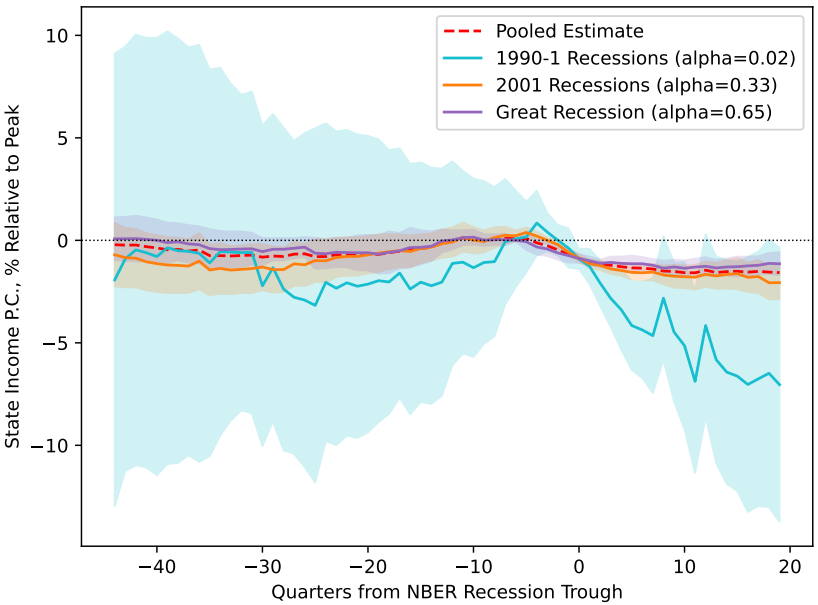
(a) Industry: 1990-1, 2001, 2007-9



(b) Industry: Other Recessions



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



(d) Recessions: Other Recessions

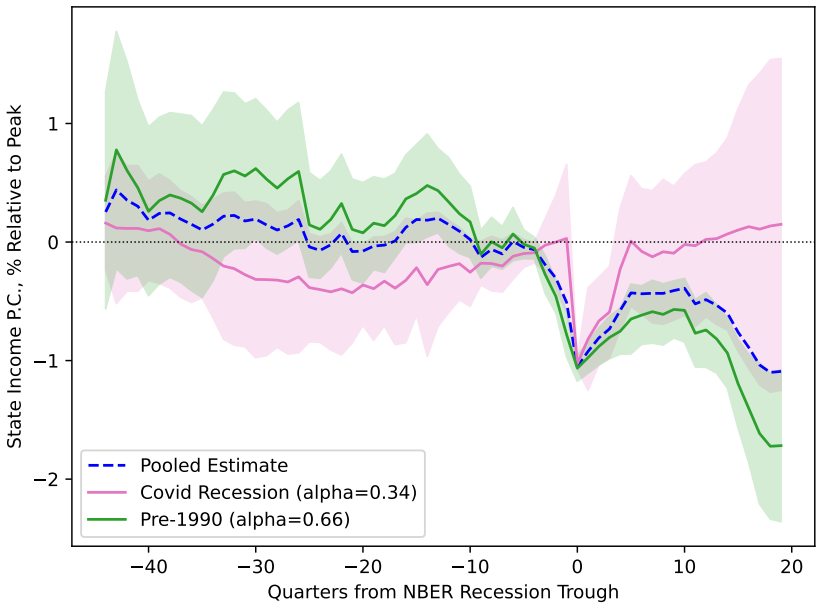
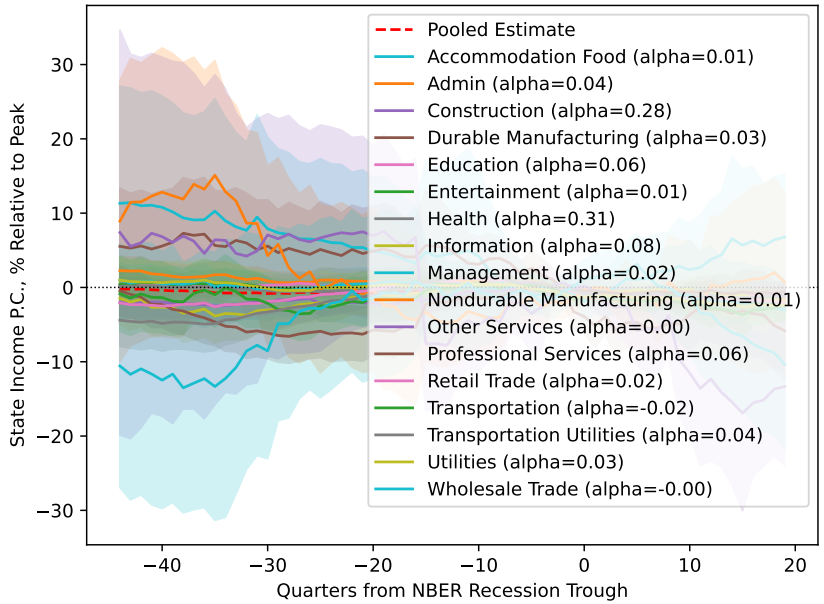
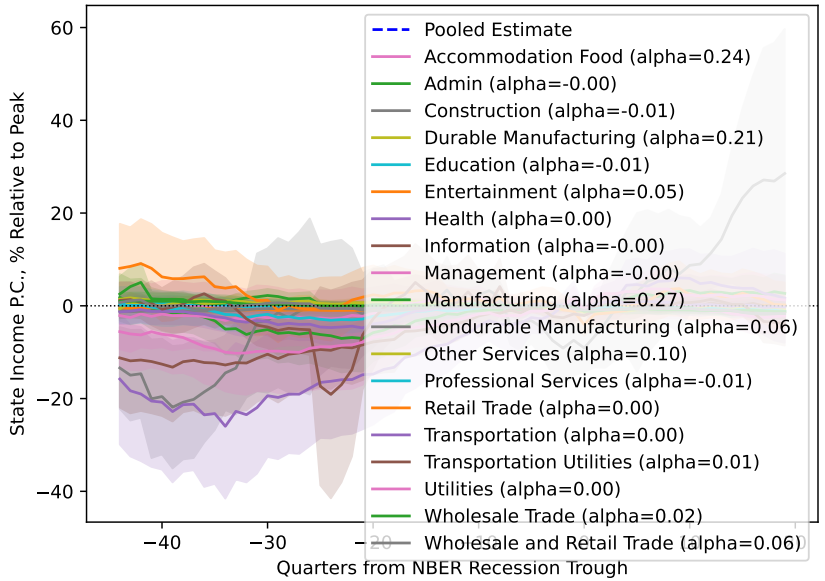


Figure 13: Impulse Response Functions based on Industry and Recession Variation

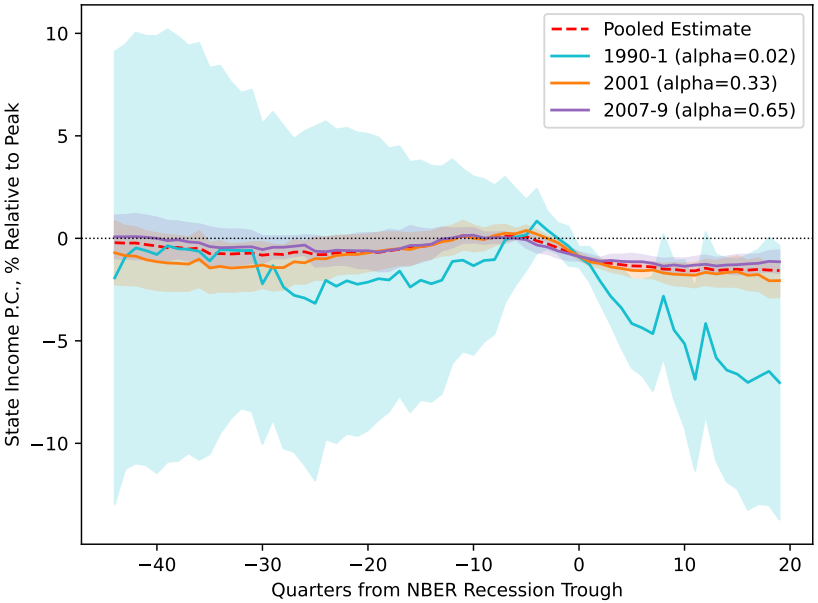
(a) Industry: 1990-1, 2001, 2007-9



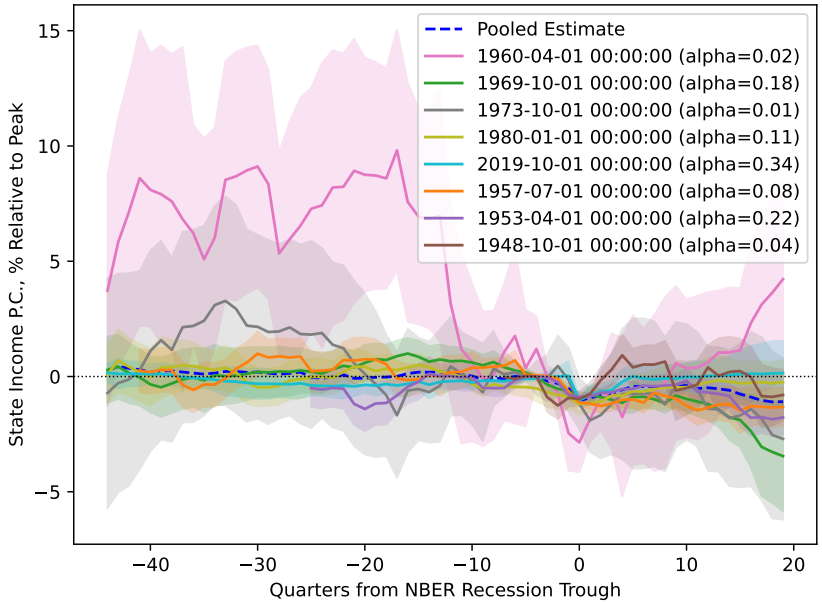
(b) Industry: Other Recessions



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



(d) Recessions: Other Recessions



E Information Structure

This appendix details how we model boom-bust or liquidationist recessions using a fake news shock to households's taste for durables ψ_t.

Notation Let $J^{i,o}$ be the full-information Jacobian for an input i and output o. Each column s of this matrix tracks the response of output o to a one-time innovation to input i in period s. Concretely, we have

$$J^{o,i} = \begin{pmatrix} \frac{\partial o_0}{\partial i_0} & \cdots & \frac{\partial o_{T-1}}{\partial i_0} \\ & \ddots & \\ \frac{\partial o_{T-1}}{\partial i_0} & \cdots & \frac{\partial o_{T-1}}{\partial i_{T-1}} \end{pmatrix} \qquad J^{o,i}_{t,s} = \frac{\partial o_t}{\partial i_s}$$

Throughout, we fix $i = \psi$ and $o = y_t$ so we omit superscripts. We use the notation $J_{[a\bullet b,c\bullet d]}$ to denote the submatrix of J from row a to b and column c to d, and is zero otherwise. We label the first column and row as 0, so the final column and row are labeled $T-1$. We also define 0_m as a row of m zeros and $0_{m,n}$ as an $m \times n$ matrix of zeros. The taste shock is a $T \times 1$ vector $\psi_{0,T-1} = (\psi_t)_{t=0}^{T-1}$. We denote the subsequence from t to $t+k$ by $\psi_{t,t+k}$. Under full information the output sequence, $y_{0,T-1}$, in response to the shock path $\psi_{0,T-1}$ is:

$$y_{0,T-1} = J\psi_{0,T-1}$$

Importantly, we can construct impulse responses that incorporate different information structures by combining different columns of the full information Jacobian J. Below we detail our approach.

Fake News Shocks We assume the following two-wave information structure to model the fake news shock. In period $t = 0$, news arrive about an increase in ψ_t starting at time $s > 0$. This is the first wave of information. We denote the shock sequence associated to this wave by $\psi^0_{0,T-1}$:

$$\psi^0_{0,T-1} = \begin{pmatrix} 0_{0,s-1} \\ \bar{\psi}_{s,T-1} \end{pmatrix}$$

where $\bar{\psi}_k > 0$ for some $k \geq s$. Starting at $t = 0$, agents learn of the sequence $\psi^0_{0,T-1}$ at rate $1-\theta$, where θ governs the degree of information stickiness. Households that update their information set expect a higher return from durable consumption beginning in period s. In anticipation of this, they build-up their durable stock which boosts output and the investment

share. Because households update their information set infrequently this results in a gradual build-up of the aggregate durable stock. At time $z < s$ a new wave of information starts. In each period $t \geq z$, a share $1-\theta$ of the agents that expect the shock sequence to be $\bar{\psi}_{t,T-1}$ learn that the sequence will instead be $\psi_{t,T-1} = \mathbf{0}$. This is equivalent to learning of a shock with path given by $\tilde{\psi}_{t,T-1} = -\psi^0_{t,T-1}$. Households that receive this new piece of information, find themselves with an excess stock of durable goods that they wish to liquidate. Initially, this results in a deceleration of the accumulation process of durable goods at the aggregate level. Accumulation decelerates but does not come to an immediate halt. This is because the stock of households that still expect the shock sequence $\psi^0_{0,T-1}$ is large relative to the stock that learns of $\tilde{\psi}_{0,T-1}$. As the second wave of information disseminates, the liquidation of excess stocks generates a bust in output, a sustained decrease in the investment share and a slow recovery of the economy relative to the preceding boom period. Next, we provide further details on how we implement and put together each of these information waves.

Implementation of the first wave We can construct the outcome of the first information wave as the following series of learning steps:

$$
\begin{aligned}
y_{0,T-1} = {} & (1-\theta) J \psi^0_{0,T-1} && \text{[learn at 0]} \\
& + \theta(1-\theta) \begin{pmatrix} 0_{1,1} & 0_{1,T-1} \\ 0_{T-1,1} & J_{[0\bullet T-2, 0\bullet T-2]} \end{pmatrix} \psi^0_{0,T-1} && \text{[learn at 1]} \\
& + \theta^2(1-\theta) \begin{pmatrix} 0_{2,2} & 0_{2,T-2} \\ 0_{T-2,2} & J_{[0\bullet T-2, 0\bullet T-2]} \end{pmatrix} \psi^0_{0,T-1} && \text{[learn at 2]} \\
& \vdots \\
& + \theta^{T-1}(1-\theta) \begin{pmatrix} 0_{T-1,T-1} & 0_{T-1,1} \\ 0_{1,T-1} & J_{[0\bullet 0, 0\bullet 0]} \end{pmatrix} \psi^0_{0,T-1} && \text{[learn at } T-1\text{]}
\end{aligned}
$$

The summation representation is given by:

$$
y_{0,T-1} = \left[\sum_{k=0}^{T-1} \theta^k (1-\theta) \begin{pmatrix} 0_{k,k} & 0_{k,T-k} \\ 0_{T-k,k} & J_{[0\bullet T-1-k, 0\bullet T-1-k]} \end{pmatrix} \right] \psi^0_{0,T-1}
$$

where each term k in the summation captures the response of households that learn about the shock sequence $\psi^0_{0,T-1}$ in period $t = k$. We can define the sticky-information Jacobian S such that the outcome effect of the first wave is then

$$
y_{0,T-1} = S \psi^0_{0,T-1}
$$

S can be constructed recursively using

$$S_{[k\bullet T-1,k\bullet T-1]} = S_{[k\bullet T-1,k\bullet T-1]} + \theta^k(1-\theta)J_{[0\bullet T-1-k,0\bullet T-1-k]}, \qquad k = 0, ..., T-1$$

Implementation of the second wave The second wave of information, which starts at $z < s$, gradually updates the information set of households that expect the shock sequence $\psi^0_{0,T-1}$ by feeding in the sequence $\tilde{\psi}_{0,T-1} = -\psi^0_{0,T-1}$. We refer to the stock of households that believe the shock sequence to be $\psi^0_{0,T-1}$ as believers. In each period $t \geq z$, a portion of believers learns that the sequence will be $\psi_{0,T-1} = 0$ and the remaining believers only learn of the current realization $\tilde{\psi}_t$. This means the latter know that the taste for durables did not change at t, $\psi_t = 0$, but still expect ψ^0_{t+k} for $k > 0$.

For implementation, we need to keep track of the stock of households that expect $\psi^0_{0,T-1}$ at each point of time. Every period the flow of new believers is given by:

$$\text{new believers}_t = (1-\theta)\theta^t, \qquad t = 0, ..., T-1$$

These are the households that continue to learn about the first wave of information. Next, we assume that believe updating from the second wave of information happens immediately after the first wave and, as mentioned before, it can only affect the current stock of believers. Current believers are the sum of past believers and new believers. Starting from $t \geq z$, each period a fraction $(1-\theta)$ of them gets converted to non-believers (i.e. learns about the second wave):

$$\text{new non-believers}_t = \begin{cases} 0, & \text{if} \quad t < s \\ (1-\theta)(\text{believers}_{t-1} + \text{new believers}_t), & \text{if} \quad t \geq z \end{cases}$$

This means that the stock of believers is simply past believers plus the net change in believers:

$$\text{believers}_t = \text{believers}_{t-1} + \text{new believers}_t - \text{new non-believers}_t$$

We can compute the sticky information Jacobian associated to the second wave as follows:

$$
\begin{aligned}
\tilde{y}_{0,T-1} = {} & \text{new non-believers}_z \begin{pmatrix} 0_{z,z} & 0_{z,T-z} \\ 0_{T-z,z} & J_{[0\bullet T-z-1,0\bullet T-z-1]} \end{pmatrix} \tilde{\psi}_{0,T-1} && \text{[unlearn at } z\text{]} \\
& + \text{believers}_z \begin{pmatrix} 0_{z,1} \\ J_{[0\bullet T-1-z,0]} \end{pmatrix} \tilde{\psi}_z \\
& + \text{new non-believers}_{z+1} \begin{pmatrix} 0_{z+1,z+1} & 0_{z+1,T-z-1} \\ 0_{T-z-1,z+1} & J_{[0\bullet T-z-2,0\bullet T-z-2]} \end{pmatrix} \tilde{\psi}_{0,T-1} && \text{[unlearn at } z+1\text{]} \\
& + \text{believers}_{z+1} \begin{pmatrix} 0_{z+1,1} \\ J_{[0\bullet T-2-z,0]} \end{pmatrix} \tilde{\psi}_{z+1} \\
& \vdots \\
& \left. + \text{new non-believers}_{T-1} \begin{pmatrix} 0_{T-1,T-1} & 0_{T-1,1} \\ 0_{1,T-1} & J_{[0\bullet 0,0\bullet 0]} \end{pmatrix} \tilde{\psi}_{0,T-1} \right\} \\
& + \text{believers}_{T-1} \begin{pmatrix} 0_{T-1,1} \\ J_{[0\bullet 0,0]} \end{pmatrix} \tilde{\psi}_{T-1} && \text{[unlearn at } T-1\text{]}
\end{aligned}
$$

We can define the sticky-information Jacobian W such that the outcome effect of the second wave is:

$$\tilde{y}_{0,T-1} = W\tilde{\psi}_{0,T-1}$$

W can be constructed recursively using:

$$
\begin{aligned}
W^1_{[k\bullet T-1,k\bullet T-1]} &= W^1_{[k\bullet T-1,k\bullet T-1]} + \text{new non-believers}_k J_{[0\bullet T-1-k,0\bullet T-1-k]}, \qquad k = s, ..., T-1 \\
W^2_{[k\bullet T-1,k]} &= \text{believers}_k J_{[0\bullet T-1-k]}, \qquad k = s, ..., T-1 \\
W &= W^1 + W^2
\end{aligned}
$$

Combining both waves The total effect of the two waves is

$$
\begin{aligned}
y_{0,T-1} &= S\psi^0_{0,T-1} + W\tilde{\psi}_{0,T-1} \\
&= (S-W)\psi^0_{0,T-1}
\end{aligned}
$$

The matrix $K = S - W$ governs the transmission of a fake news shock with path given by $\psi^0_{0,T-1}$.