Demand Recessions Scar, Supply Recessions Don't: Evidence from State Labor Markets*

APEP Autonomous Research†(cumulative: 1h 37m). @SocialCatalystLab
@dyanag

February 26, 2026

Abstract

Not all recessions are created equal. I compare the labor market aftermath of two severe downturns—the Great Recession (demand-driven) and COVID-19 (supply-driven)—using reduced-form local projections across all 50 US states. Exploiting cross-state variation in housing price exposure and industry composition, I find that the housing-driven Great Recession produced persistent relative employment deficits: a pre-specified long-run average across horizons 48–120 months shows $\bar{\pi}_{LR} = -0.037$ (wild cluster bootstrap 95% CI: $[-0.069, -0.005]$). COVID-exposed states recovered fully within 18 months. Individual long-horizon estimates are imprecise and do not survive Romano-Wolf stepdown adjustment, reflecting power limitations inherent in 50-state cross-sections. A Diamond-Mortensen-Pissarides model with duration-dependent skill depreciation—estimated via SMM targeting the LP impulse responses—rationalizes this asymmetry. The model-implied demand-to-supply welfare ratio is approximately 7–18:1 depending on shock persistence, though this range is sensitive to calibration choices and the model is rejected by overidentification tests.

JEL Codes: E24, E32, J63, J64

Keywords: hysteresis, labor market scarring, recessions, search and matching, Great Recession, COVID-19

*This paper is revision v9 of APEP-0238. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0238 for the previous version.

†Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

The COVID-19 recession shed jobs three times faster than the Great Recession, yet it recovered them three times as quickly. In 2008, a 6 percent employment drop took six years to heal. In 2020, a 15 percent drop vanished in two. These were the two worst labor market contractions since the Great Depression, yet their aftermath could hardly look more different. The Great Recession left scars visible a decade later in depressed employment rates, elevated disability claims, and permanently lower output trajectories (Yagan, 2019; Fernald et al., 2017; Summers, 2014; Ball, 2014). The COVID recession, despite a peak contraction roughly three times as severe, left almost no detectable long-run trace. The contrast is a puzzle—and the answer lies not in the depth of the initial shock but in its nature. The Great Recession was a demand recession: a collapse in household balance sheets and aggregate spending that destroyed the incentive to hire, creating prolonged unemployment spells that eroded workers' human capital and attachment to the labor force (Mian and Sufi, 2014; Blanchard and Summers, 1986; Pissarides, 1992; Giroud and Mueller, 2017). COVID was a supply recession: a temporary shutdown of production capacity that, once lifted, allowed rapid recall of workers whose skills and employer relationships remained intact (Cajner et al., 2020; Gregory et al., 2020; Barrero et al., 2021). Think of the economy as a guitar string: a supply shock plucks it down, but the tension snaps it back to its natural position; a demand shock corrodes the string itself—weakening its structure through skill depreciation and labor force exit—so it may not fully rebound. The distinction between demand and supply origins—long central to macroeconomic theory—has first-order consequences for whether recessions scar.

I test this hypothesis using a reduced-form local projection framework applied to state-level labor market data from the Bureau of Labor Statistics, covering all 50 states at monthly frequency from 2000 to 2024. The empirical strategy exploits cross-state variation in recession exposure through two exposure measures. For the Great Recession, I use the 2003–2006 state-level housing price boom as an exposure measure for the severity of the local demand collapse, following Mian and Sufi (2014) and Charles et al. (2016). For the COVID recession, I construct a Bartik exposure measure based on pre-pandemic industry employment shares interacted with national industry-level employment changes, in the spirit of Bartik (1991) and Goldsmith-Pinkham et al. (2020).

The reduced-form results are striking. A pre-specified long-run average of the Great Recession LP coefficients across horizons 48–120 months yields $\bar{\pi}_{LR} = -0.037$ (wild cluster bootstrap 95% CI: $[-0.069, -0.005]$), confirming persistent scarring as a single test object that avoids the multiple-testing concerns inherent in horizon-by-horizon inference with 50 cross-sectional units. Individual long-horizon coefficients are imprecise—no single horizon

survives Romano-Wolf stepdown adjustment at the 5% level—but the joint evidence for persistence is robust. By contrast, the relationship between COVID Bartik exposure and subsequent employment is statistically indistinguishable from zero by 18 months. States hit hardest by COVID recovered fully and proportionally; states hit hardest by the Great Recession did not.

To understand the mechanisms behind this asymmetry, I develop a Diamond-Mortensen-Pissarides (DMP) search and matching model augmented with endogenous labor force participation and duration-dependent human capital depreciation. Four key parameters are estimated via Simulated Method of Moments (SMM), disciplining the model against the LP impulse responses. A demand shock—modeled as a highly persistent AR(1) reduction in aggregate productivity ($\rho = 0.99$, half-life ≈ 69 months)—depresses vacancy creation, lengthens unemployment durations, triggers skill depreciation, and sets off a vicious cycle of declining match quality and labor force exit. A supply shock—modeled as a temporary spike in separations—creates mass unemployment but preserves short durations because the shock itself is transient. Workers are recalled or reabsorbed quickly, before human capital depreciates.

The estimated model reproduces the qualitative asymmetry. A 5% demand shock with AR(1) persistence generates a peak employment decline of 1.2% that dissipates slowly—employment remains 0.5% below steady state at 120 months through the interaction of persistent productivity loss and the scarring channel. A temporary supply shock generates a sharper initial decline of 6.8% but recovers to within 0.3% of steady state by 12 months. The model-implied demand-to-supply welfare ratio ranges from approximately 7:1 to 18:1 depending on shock persistence (ρ_a), with a baseline of 12:1 at $\rho_a = 0.99$. This range should be interpreted cautiously: the model is rejected by overidentification tests (the simple AR(1) shock process cannot match the long-run plateau in the data), and the welfare ratio is sensitive to calibration choices. The scarring channel accounts for 9% of the demand welfare loss.

This is fundamentally a comparative case study of two episodes, not a general statement about all demand and supply recessions. Within this comparison, demand-driven exposure produces persistent relative employment deficits while supply-driven exposure does not—but the COVID-19 recovery cannot be separated from the unprecedented fiscal response, and the Great Recession's scarring reflects the full episode-specific treatment package of shock type, policy response, and sectoral dynamics.

This paper contributes to and advances the macroeconomic hysteresis literature pioneered by Blanchard and Summers (1986) and extended by Cerra and Saxena (2008), Jordà et al. (2016), and Cerra et al. (2023). While this literature has established that recessions *can* leave permanent scars, it has not systematically examined *which types* scar most. I provide the

first direct comparison of demand-driven versus supply-driven recession dynamics using the same identification framework and the same labor markets observed across two episodes, contributing to the literature on local labor market adjustment (Blanchard and Katz, 1992; Autor et al., 2013; Notowidigdo, 2020; Amior and Manning, 2021; Beraja et al., 2019) by showing that the speed of adjustment depends critically on the nature of the shock, not just its magnitude—complementing Hershbein and Stuart (2020), Dao et al. (2017), and Foote et al. (2021). I also connect the COVID recession literature (Chetty et al., 2020; Cajner et al., 2020; Forsythe et al., 2021) to the broader question of recession persistence, arguing that the rapid COVID recovery reflects a simpler explanation than effective fiscal policy (Autor et al., 2022) or inherent labor market resilience (Hall and Kudlyak, 2020): supply shocks do not create the prolonged unemployment that generates hysteresis. A unified DMP framework that nests both demand and supply recessions, building on Shimer (2005), Hall (2005), and the skill depreciation mechanism of Pissarides (1992), shows that the interaction of search frictions with human capital depreciation creates a powerful amplification mechanism for demand shocks that is absent for supply shocks.

2. Institutional Background

2.1 The Great Recession: Anatomy of a Demand Collapse

The Great Recession originated in the US housing market. Between 2000 and 2006, national home prices rose by approximately 90%, but the boom was geographically concentrated: supply-constrained markets with speculative demand (Nevada, Arizona, Florida, California) saw price increases exceeding 100%, while elastic-supply states (Texas, the Midwest) experienced modest appreciation (Mian et al., 2013; Glaeser and Nathanson, 2017; Saiz, 2010). This cross-state variation in the boom became the cross-state variation in the bust.

The collapse began in mid-2006. By 2008, falling home values had wiped out trillions in household wealth, triggering mortgage defaults, bank losses, and a full-blown financial crisis. Mian and Sufi (2014) demonstrate that the employment decline between 2007 and 2009 was closely tied to the decline in local housing net worth, operating through the household balance sheet channel. Giroud and Mueller (2017) provide complementary evidence that firm leverage amplified these losses through the credit channel.

The critical feature for understanding hysteresis is the nature of the employment decline. The demand collapse was not a temporary disruption—it was a sustained reduction in the willingness and ability to spend. Firms permanently closed establishments, laid off workers, and did not rehire even as conditions slowly improved. The share of unemployed workers out of work for 27 weeks or more rose from 17.4% in December 2007 to 45.5% in April 2010,

unprecedented in postwar data (Kroft et al., 2016). Mean unemployment duration peaked at 39.4 weeks. These extended spells are the mechanism through which demand recessions scar: workers lose skills, employer networks atrophy, and stigma reduces callback rates (Kroft et al., 2016; Elsby et al., 2010).

The policy response arrived with a lag. The Federal Reserve reached the zero lower bound by December 2008; the ARRA provided $787 billion beginning February 2009. By the time stimulus reached full force, millions had already crossed the duration thresholds associated with skill depreciation and discouragement.

2.2 The COVID Recession: Anatomy of a Supply Disruption

The COVID-19 recession was fundamentally different. In March 2020, stay-at-home orders and business closures drove nonfarm payrolls down by 22.4 million—a 14.7% decline that dwarfed any postwar contraction. Yet by July 2022, just 29 months later, all lost jobs had been recovered.

COVID was a supply shock: an exogenous disruption to production, not a reduction in the desire to consume. Consumer spending on durable goods actually *increased* during 2020 (Guerrieri et al., 2022). The shock was sectoral—leisure and hospitality lost 8.2 million jobs (49% of the sector)—and geographic, hitting tourism-dependent states (Hawaii, Nevada, New York) hardest (Cajner et al., 2020).

Crucially, the COVID shock preserved employer-employee matches. Many workers were furloughed rather than permanently separated. The PPP provided $800 billion in forgivable payroll loans, directly subsidizing match preservation (Autor et al., 2022). Enhanced unemployment insurance maintained household income (Ganong et al., 2020). The result: unemployment durations remained short even though the peak unemployment rate (14.7%) exceeded the Great Recession peak (10.0%). Median unemployment duration peaked at only 10.4 weeks during COVID versus 25.0 weeks after the Great Recession. Short durations meant minimal skill depreciation, minimal discouragement, and rapid reabsorption.

2.3 Why the Comparison Is Informative

Both recessions were severe, affected all 50 states, featured substantial cross-state variation in severity driven by well-understood sources, and were followed by aggressive policy responses. The fundamental difference—demand versus supply—allows me to test whether hysteresis operates specifically through prolonged unemployment generated by demand deficiency, rather than through the simple fact of job loss. If hysteresis were driven purely by the experience of unemployment regardless of cause, COVID would have left similar scars. That it did not

isolates the demand channel.

3. Conceptual Framework

I develop a search and matching model in the tradition of Diamond (1982), Mortensen and Pissarides (1994), and Pissarides (1985), augmented with endogenous labor force participation and human capital depreciation during unemployment, building on the duration-dependence mechanisms in Pissarides (1992) and Gertler and Trigari (2009). The model generates three testable predictions about the differential persistence of demand versus supply shocks.

3.1 Environment

Time is discrete, with periods indexed by $t = 0, 1, 2, \ldots$. The economy is populated by a unit measure of infinitely-lived, risk-neutral workers who discount the future at rate $\beta \in (0, 1)$. Workers can be in one of three states: employed (E), unemployed and actively searching (U), or out of the labor force (O). There is a continuum of firms, each operating a single job.

Each worker is characterized by a human capital level $h \in (0, 1]$ that determines the productivity of a filled job. A match between a firm and a worker with human capital h produces output $a \cdot h$ per period, where $a > 0$ is aggregate productivity—the key object that demand shocks affect.

3.2 Matching Technology

Unemployed workers search for jobs by entering a labor market where firms post vacancies. The number of new matches formed per period is given by a constant-returns-to-scale matching function:

$$m(u_t, v_t) = A \cdot u_t^{\alpha} \cdot v_t^{1-\alpha}, \tag{1}$$

where u_t is the measure of unemployed workers, v_t is the measure of vacancies, $A > 0$ is matching efficiency, and $\alpha \in (0, 1)$ is the elasticity of matches with respect to unemployment. Define labor market tightness as $\theta_t \equiv v_t/u_t$. The job finding rate for an unemployed worker is:

$$f(\theta_t) = A \cdot \theta_t^{1-\alpha}, \tag{2}$$

and the vacancy filling rate for a firm is:

$$q(\theta_t) = A \cdot \theta_t^{-\alpha}. \tag{3}$$

Tighter labor markets (higher θ) make it easier for workers to find jobs but harder for firms to fill vacancies.

3.3 Timing and Transitions

Within each period, the following events occur. First, aggregate shocks are realized (changes to a or the separation rate δ). Second, employed workers separate from their jobs with probability δ_t. Newly separated workers enter the unemployment pool in the following period and are first exposed to job-finding and OLF exit hazards at $t+1$. In steady state, δ reflects the average monthly separation rate. A supply shock operates by temporarily elevating δ_t above its steady-state value. Third, three transitions operate as *competing hazards* on the unemployment pool: workers match with vacancies (probability $f(\theta_t)$), exit to non-participation (probability χ_t), or remain unemployed (probability $1 - f(\theta_t) - \chi_t$). Separately, a fraction ψ of non-participants re-enter unemployment. Fourth, human capital is determined by unemployment duration: workers who have been unemployed for at least d^* periods have reduced human capital $h(d) = 1 - \lambda$, where $\lambda \in (0, 1)$ parameterizes the severity of scarring, while workers with $d < d^*$ have full human capital $h(d) = 1$.

The exit rate from unemployment to non-participation has two components:

$$\chi_t = \chi_0 + \chi_1 \cdot s_t, \tag{4}$$

where χ_0 is a baseline exit rate and χ_1 governs the additional discouragement effect when the fraction s_t of unemployed workers who have experienced skill depreciation ("scarred" workers) is high. Here s_t denotes the scarred fraction at the *beginning* of period t (pre-transition), computed from the duration distribution inherited from $t-1$. This beginning-of-period convention—standard in discrete-time search models—means that χ_t and the effective human capital $\bar{h}_t = 1 - \lambda s_t$ are determined by the state entering the period, before matching and transitions occur. Equivalently, in the simulation algorithm (Appendix A), the scarred fraction entering period t is denoted s_{t-1} under end-of-period indexing: the algorithm computes the distribution at the end of period $t-1$ and uses it to set χ_t and $\bar{h}_t$ at the start of period t. Both conventions refer to the same object. This captures the empirical regularity that labor force participation declines during prolonged recessions as discouraged workers abandon job search (Elsby et al., 2010; Kroft et al., 2016). The scarred fraction s_t is determined by the full unemployment duration distribution, whose law of motion is specified in the simulation approach (Section 3.8).

3.4 Value Functions

The value of employment for a worker with human capital h determines the worker's willingness to accept job offers and remain employed. The model environment is risk-neutral for tractability; welfare is evaluated under both risk-neutral and CRRA preferences as a robustness check in Section 8.5.

$$W_t(h) = w_t(h) + \beta \left[(1 - \delta_t) W_{t+1}(h) + \delta_t \cdot U_{t+1}(0) \right], \tag{5}$$

where $w_t(h)$ is the wage. Employed workers produce at their match-formation human capital h throughout the employment spell; match output does not change during the match. Upon separation, workers re-enter the unemployment pool at duration $d = 0$. Since $h(d)$ captures the erosion of job-search effectiveness and labor market attachment during extended unemployment—rather than general human capital—recent employment restores search capital to its baseline level $h(0) = 1$. This is consistent with evidence that earnings penalties from displacement diminish with subsequent employment tenure (Jacobson et al., 1993). The value of unemployment for a worker at duration d is:

$$U_t(d) = b + \beta \left[f(\theta_t) W_{t+1}(h(d)) + \chi_t V_{t+1}^{OLF} + (1 - f(\theta_t) - \chi_t) U_{t+1}(d+1) \right], \tag{6}$$

[1]where b is the flow value of unemployment and $h(d)$ is the duration-dependent human capital defined by:

$$h(d) = \begin{cases} 1 & \text{if } d < d^* \\ 1 - \lambda & \text{if } d \geq d^* \end{cases}. \tag{7}$$

Since human capital is fully determined by unemployment duration, the value function requires only d as an individual state variable. Workers at $d < d^*$ have full human capital $h = 1$; workers at $d \geq d^*$ have reduced human capital $h = 1 - \lambda$. This is exactly a two-point distribution, consistent with the aggregate relation $\bar{h}_t = 1 - \lambda s_t$. Upon reemployment, workers rebuild human capital through on-the-job experience: separated workers re-enter unemployment at $d = 0$ with $h(0) = 1$ (Eq. 5), so the duration distribution fully characterizes the scarred fraction $s_t = \sum_{d \geq d^*} u_d / U_t$ without requiring a separate "type" state variable. This interpretation is consistent with evidence that earnings losses from displacement, while initially large, diminish over subsequent employment spells (Jacobson et al., 1993; Jarosch, 2023).

[1]The Cobb-Douglas matching function can generate $f(\theta) > 1$ for sufficiently high θ. In simulations, we cap all transition probabilities at 1 and ensure $f(\theta_t) + \chi_t \leq 1$. Similarly, $q(\theta_t)$ is capped at 1. At calibrated values, the maximum realized $q(\theta_t) = 0.85$ and $f(\theta_t) = 0.42$, so the bounds are never approached.

The value of non-participation is:

$$V_t^{OLF} = b_{OLF} + \beta \left[\psi \cdot U_{t+1}(0) + (1 - \psi) \cdot V_{t+1}^{OLF} \right], \tag{8}$$

where b_{OLF} is the flow value of home production, assumed to be less than the unemployment benefit ($b_{OLF} < b$), and re-entrants start with unemployment duration $d = 0$, which implies human capital $h(0) = 1$ by Eq. 7. Under the perceived-permanent assumption (Section 3.6), the aggregate state is treated as stationary within each period, so V_t^{OLF} is effectively determined by the current-period equilibrium objects. We suppress the time subscript in steady-state analysis.

The reduced-form hazard χ_t in Equation (4) captures the participation margin: when unemployment durations are long and human capital has depreciated, more workers exit the labor force. This is the channel through which demand recessions generate persistent employment losses. The competing-hazard formulation—where job finding and OLF exit are parallel transitions from unemployment—is consistent with the aggregate laws of motion in Equations (12) and (13).

These value functions characterize the welfare implications of each labor market state and provide the foundation for the Nash bargaining solution. However, under the adopted simplified wage rule (Eq. 9) and the perceived-permanent free-entry condition (Eq. 10), equilibrium tightness θ_t is determined directly from aggregate state variables without explicitly solving the individual value functions. The value functions thus serve primarily an interpretive and welfare-evaluation role.

3.5 Wage Determination

In the continuous-time DMP limit without the OLF margin, standard Nash bargaining yields $w = (1-\gamma)b + \gamma(ah + \kappa\theta)$. In the discrete-time environment with discounting and separation, the exact Nash wage is modified by discount factors; however, the key qualitative properties are preserved. With the additional OLF exit option, the modified Nash wage would replace b with a function of b, χ_t, and V^{OLF}. For tractability in the transition dynamics, I adopt a simplified proportional-surplus wage rule that abstracts from both tightness and OLF effects:

$$w_t(h) = \gamma \cdot a_t \cdot h + (1 - \gamma)b, \tag{9}$$

where γ is the worker's bargaining power. This wage rule captures the key comparative statics—wages rise with productivity a_t and the outside option b—while permitting a closed-form solution for equilibrium tightness. The omission of the $\kappa\theta$ term means wages do not respond

to tightness, which allows θ to fluctuate more freely than under the full Nash wage (where rising wages in tight markets dampen vacancy creation). The qualitative demand/supply asymmetry is robust to this simplification because it arises from the interaction of duration dependence with the permanent/temporary nature of shocks, not from wage dynamics.

3.6 Free Entry

Firms post vacancies as long as the expected return covers the posting cost κ per period. The free entry condition is:

$$\frac{\kappa}{q(\theta_t)} = \beta \cdot (1-\gamma) \cdot \frac{a_t \cdot \bar{h}_t - b}{1-\beta(1-\delta_t)}, \tag{10}$$

This closed-form expression is derived under the assumption that firms treat the current aggregate state (a_t, δ_t) as permanent when computing match values—a standard computational simplification in the DMP literature (Shimer, 2005; Hall, 2005). Under this *perceived-permanent* (myopic) beliefs assumption, the match value $J_t(h) = (1-\gamma)(a_t h - b)/[1-\beta(1-\delta_t)]$ admits a closed-form solution and pins down θ_t without backward induction. For the AR(1) demand shock (productivity follows $a_t = a_{ss}(1 - \Delta a \cdot \rho_a^t)$ with $\rho_a = 0.99$), the perceived-permanent assumption means firms treat the current a_t as if it will persist forever. Because the shock is highly persistent (half-life ≈ 69 months), this overstates the shock at each t only modestly. Importantly, this creates a conservative bias: under rational expectations with known mean reversion, firms would post more vacancies and recovery would be faster, implying a smaller demand-supply welfare gap than we report. For the temporary supply shock (δ_t elevated for T_s months), firms slightly overreact to the separation spike, depressing θ_t more than under rational expectations. This overstates the supply-shock trough depth but again biases against our main result.

Here $\bar{h}_t$ is the expected human capital of a matched worker (approximated by $1 - \lambda s_{t-1}$ in simulations). This pins down equilibrium tightness θ_t as a function of productivity a_t, expected human capital $\bar{h}_t$, and the separation rate δ_t. A persistent reduction in a (demand shock) lowers the right-hand side, reducing θ_t and the job finding rate $f(\theta_t)$. The lower finding rate extends unemployment durations, triggering skill depreciation and labor force exit. Even as a_t mean-reverts, scarred workers depress effective human capital $\bar{h}_t$, creating endogenous persistence beyond the exogenous shock. A temporary increase in δ_t (supply shock) reduces the present value of a match (via $1-\beta(1-\delta_t)$ in the denominator), lowering θ_t while the shock lasts. However, because the shock is transient, θ_t recovers quickly once δ_t returns to its steady-state value.

The equilibrium requires three regularity conditions: (i) positive match surplus $a_t \bar{h}_t > b$,

ensuring interior vacancy creation; (ii) $f(\theta_t) \leq 1$, ensuring well-defined transition probabilities; and (iii) $f(\theta_t) + \chi_t \leq 1$, ensuring the competing hazards are feasible. At calibrated parameter values, all three conditions hold with ample margin throughout the shock simulations (the maximum realized $f(\theta_t) = 0.42$, and $\chi_t \leq 0.031$).

3.7 Laws of Motion

The aggregate state of the economy is described by the distribution of workers across employment states.[2] The laws of motion are:

$$E_{t+1} = E_t + f(\theta_t) \cdot U_t - \delta_t \cdot E_t, \tag{11}$$

$$U_{t+1} = U_t + \delta_t \cdot E_t + \psi \cdot O_t - f(\theta_t) \cdot U_t - \chi_t \cdot U_t, \tag{12}$$

$$O_{t+1} = O_t + \chi_t \cdot U_t - \psi \cdot O_t, \tag{13}$$

with the adding-up constraint $E_t + U_t + O_t = 1$. In steady state:

$$\delta \cdot E = f(\theta) \cdot U, \quad \chi \cdot U = \psi \cdot O, \quad E + U + O = 1. \tag{14}$$

Solving yields $U^{ss} = \left(\frac{f}{\delta} + 1 + \frac{\chi}{\psi}\right)^{-1}$, $E^{ss} = \frac{f}{\delta} \cdot U^{ss}$, and $O^{ss} = \frac{\chi}{\psi} \cdot U^{ss}$. In steady state with $f \approx 0.40$ and monthly OLF exit $\chi_0 = 0.03$, the survival probability to $d^* = 12$ months is approximately $(1 - f - \chi_0)^{12} \approx 0.57^{12} \approx 0.001$, making s^{ss} negligible. The algorithm initializes the steady-state duration distribution as a geometric decay with this property.

3.8 Demand Versus Supply Shocks

The model nests two types of recession:

Demand shock. Aggregate productivity follows a mean-reverting AR(1) process: $a_t = a_{ss}(1 - \Delta a \cdot \rho_a^t)$ for $t \geq 1$, where $\rho_a \in (0, 1)$ governs the persistence of the shock. With $\rho_a = 0.99$ (externally calibrated from the GR housing price recovery timeline), the half-life is approximately 69 months. This reduces the match surplus, lowering θ and $f(\theta)$. Unemployment durations lengthen, the scarred fraction s_t rises, human capital depreciates, and the participation exit rate χ_t increases. The key mechanism: even as productivity mean-reverts, workers who accumulated long unemployment spells during the trough suffer permanent skill depreciation, creating endogenous persistence beyond the exogenous shock process. The scarring channel thus transforms a transient (though highly persistent) productivity shock

[2]To avoid confusion: U_t in the laws of motion (Eqs. 11–13) denotes the aggregate unemployment *stock*, while $U_t(d)$ in the value functions (Eq. 6) denotes the *value* of unemployment at duration d. Context and the presence of arguments distinguish the two.

into a long-lasting employment decline.

Supply shock. The separation rate spikes temporarily: $\delta_t = \delta(1 + \sigma_s)$ for $t \in \{1, \ldots, T_s\}$ and $\delta_t = \delta$ thereafter. This creates a mass of newly unemployed workers, but because the shock dissipates quickly, θ and $f(\theta)$ recover rapidly. Unemployment durations remain short, few workers cross the scarring threshold d^*, and labor force participation is largely unaffected.

Simulation approach. The quantitative simulations track the full discrete unemployment duration distribution. Each period, new separations enter at duration $d = 0$; surviving unemployed workers at duration d advance to $d + 1$; workers with $d \geq d^*$ experience skill depreciation. The scarred fraction $s_t = \sum_{d \geq d^*} u_d / \sum_d u_d$ is computed exactly from the distribution, and the participation exit rate responds to s_t through Equation (4). Four model parameters (λ, χ_1, A, κ) are estimated via Simulated Method of Moments targeting six moments from the LP impulse responses and steady-state labor market targets (see Section 8.2). Formally, let $u_{d,t}$ denote the mass of unemployed workers at duration d in period t. The duration distribution evolves as:

$$u_{0,t+1} = \delta_t E_t + \psi O_t, \tag{15}$$

$$u_{d+1,t+1} = (1 - f(\theta_t) - \chi_t) \cdot u_{d,t} \quad \text{for } d \geq 0. \tag{16}$$

The scarred fraction and effective human capital are then $s_t = \sum_{d \geq d^*} u_{d,t}/U_t$ and $\bar{h}_t = 1 - \lambda s_t$, closing the model. This approach follows Shimer (2005) in working with aggregate flows, but improves on reduced-form proxies by maintaining exact consistency between the Bellman equations that take d^* as given and the simulation dynamics.

The model's asymmetric prediction connects to the "plucking" model of Dupraz et al. (2024): supply shocks displace the economy temporarily, while demand shocks erode its capacity to recover.

An important caveat: the clean demand/supply dichotomy can break down. Guerrieri et al. (2022) show that supply shocks can generate "Keynesian supply shocks"—demand deficiencies triggered by supply disruptions. That COVID recovery was rapid suggests this channel was empirically muted, consistent with the massive fiscal transfers that sustained household income and prevented balance-sheet deterioration.

3.9 Testable Predictions

The model generates three predictions:

Prediction 1 (Persistent Demand Effects). *Cross-state employment responses to demand-shock exposure exhibit persistent negative effects at horizons of 48 months or more.* This

follows from the persistent reduction in θ and the amplification through skill depreciation, which creates endogenous persistence beyond the exogenous shock process.

Prediction 2 (Transient Supply Effects). *Cross-state employment responses to supply-shock exposure converge to zero within 18–24 months.* The temporary nature of the separation shock means the economy returns to its original steady state.

Prediction 3 (Mechanism: Duration and Participation). *The persistence gap operates through unemployment duration, skill depreciation, and participation exit rather than through the initial depth of the shock.*

4. Data

4.1 Data Sources

I use state-level labor market data from the Bureau of Labor Statistics, covering all 50 states at monthly frequency from 2000 to 2024, combined with recession exposure measures from public federal data sources accessed through the Federal Reserve Economic Data (FRED) API.

Employment. State-level total nonfarm payroll employment comes from the BLS Current Employment Statistics (CES) survey, available monthly for all 50 states from January 2000 through June 2024. I use seasonally adjusted series (FRED mnemonic: `[ST]NA` (e.g., `CANA` for California)).

Unemployment and Labor Force Participation. State-level unemployment rates come from the BLS Local Area Unemployment Statistics (LAUS) program, monthly seasonally adjusted for all 50 states. State-level labor force participation rates come from the same program for all 50 states; LFPR results are reported in Section C.

Housing Prices. State-level house price indices come from the Federal Housing Finance Agency (FHFA). I construct the housing price boom measure as the log change in the FHFA index between 2003Q1 and 2006Q4.

Industry Employment. State-level employment by major industry sector from the BLS Current Employment Statistics (CES) program, accessed through FRED. I use 10 major industry supersectors. Pre-recession industry shares are computed from CES data in the year preceding each recession (2006 for the Great Recession, 2019 for COVID).

JOLTS. National-level data on job openings, hires, quits, and layoffs from the Job Openings and Labor Turnover Survey, available monthly from December 2000.

4.2 Sample Construction

The analysis sample is a balanced panel of all 50 U.S. states observed monthly from January 2000 through June 2024 ($50 \times 294 = 14{,}700$ state-month observations). The cross-sectional local projections use all 50 states for both recession analyses.

For the Great Recession, the NBER peak is December 2007; I track outcomes through December 2017 (120 months). For COVID, the peak is February 2020; I track outcomes through February 2024 (48 months), deliberately truncating at 48 months even though the sample extends to June 2024 (52 months post-peak) to use round horizons consistent with the Great Recession analysis. All employment variables are expressed in logs relative to the pre-recession peak.

4.3 Variable Definitions

Employment change. For each state s and horizon h: $\Delta y_{s,h} = \ln(E_{s,t_0+h}) - \ln(E_{s,t_0})$.

Housing price boom (Great Recession instrument). $HPI_s = \ln(P_{s,2006Q4}) - \ln(P_{s,2003Q1})$, capturing the state-level intensity of the housing bubble.

Bartik shock (COVID instrument). $B_s = \sum_j \omega_{s,j,2019} \cdot \Delta E^{nat}_{-s,j}$, where $\omega_{s,j,2019}$ is the share of industry j in state s's total employment in 2019, and $\Delta E^{nat}_{-s,j}$ is the leave-one-out national employment change in industry j between February and April 2020, following Goldsmith-Pinkham et al. (2020). Intuitively, the Bartik measure asks how much a state's employment would have declined if each of its industries followed national trends.

Peak-to-trough employment. $\text{Trough}_s = \min_t \{\ln(E_{s,t}) - \ln(E_{s,t_0})\}$ for $t \in [t_0, t_0 + 24]$.

4.4 Summary Statistics

Table 1 presents summary statistics. Panel A reports the state-month panel. Average state employment is 2,773 thousand, ranging from 237 thousand (Wyoming) to 18,010 thousand (California). The average unemployment rate is 5.3%, with substantial variation (with a peak of 30.5% in Nevada during April 2020). Mean labor force participation is 65.0%.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max	N
Panel A: State-level outcomes (monthly, 2000–2024)					
Nonfarm payrolls (thousands)	2,773.1	2,963.4	237.3	18,010.0	14,700
Unemployment rate (%)	5.3	2.2	1.7	30.5	14,700
Labor force participation rate (%)	65.0	4.3	52.7	75.8	14,700
Panel B: Recession exposure measures					
Bartik shock: Great Recession	-0.0557	0.0097	-0.0966	-0.0353	50
Bartik shock: COVID	-0.1760	0.0236	-0.2779	-0.1018	50
Housing price boom (log, 2003–2006)	0.3033	0.1466	0.0653	0.6146	50
Peak-to-trough employment (GR)	-0.0593	0.0272	-0.1393	0.0000	50
Peak-to-trough employment (COVID)	-0.1547	0.0447	-0.2708	-0.0963	50

Notes: Panel A reports summary statistics for the balanced state-month panel. Nonfarm payrolls are from BLS Current Employment Statistics via FRED. Panel B reports cross-state recession exposure measures. Bartik shocks are constructed using pre-recession industry employment shares weighted by national industry employment changes. Housing price boom is the log change in the FHFA state-level house price index from 2003Q1 to 2006Q4.

Panel B reports the cross-state recession exposure measures. The housing price boom averages 0.30 log points (standard deviation 0.14), ranging from 0.065 (Texas) to 0.615 (Arizona/Nevada). The COVID Bartik shock averages -0.176 with a standard deviation of 0.023. Peak-to-trough employment declines averaged 5.9% during the Great Recession and 15.5% during COVID—the initial COVID shock was roughly 2.6 times more severe, yet the long-run effects are reversed. Table 12 in Section B reports the five most and least affected states by each recession. The question is whether this cross-state variation in exposure translates into lasting differences in recovery—and the answer turns on identification.

5. Empirical Strategy

5.1 Local Projections Framework

I estimate the dynamic effects of recession exposure on state employment using the local projections (LP) method of Jordà (2005). For each horizon $h = 0, 3, 6, 12, 24, \ldots, 120$ months, I estimate:

$$\Delta y_{s,h} = \alpha_h + \pi_h \cdot Z_s + \phi_h' X_s + \varepsilon_{s,h}, \tag{17}$$

where $\Delta y_{s,h}$ is the log employment change in state s at horizon h, Z_s is the recession exposure measure (housing price boom for the Great Recession, Bartik shock for COVID), and X_s is a vector of pre-recession state characteristics (log nonfarm employment at the recession peak, pre-recession employment growth, census region indicators).

The coefficient π_h traces an impulse response function: it measures how cross-state differences in recession exposure map into cross-state differences in employment outcomes at each horizon. Persistent π_h—remaining significantly negative at long horizons—indicates scarring; convergence toward zero indicates recovery.

Sign convention. For the Great Recession, negative π_h indicates that more-exposed states experienced larger employment declines. For COVID, the Bartik shock is constructed so that negative values indicate larger predicted employment declines. I multiply the COVID LP coefficients by -1 so that negative $\hat{\pi}_h$ indicates employment loss in both panels, matching the Great Recession convention. In both cases, convergence toward zero indicates recovery.

I adopt local projections rather than TWFE or staggered difference-in-differences because the cross-state variation in exposure is continuous, not a discrete treatment adoption. The staggered DiD concern (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021) does not apply here because I exploit cross-sectional variation at a single event date. LPs are also robust to lag structure misspecification (Ramey, 2016; Plagborg-Møller and Wolf, 2021) and naturally accommodate different sample lengths across recessions.

Estimation approach. The estimates are reduced-form local projections: I regress employment outcomes directly on Z_s, not on an endogenous treatment instrumented by Z_s. The coefficient π_h captures the total effect of exogenous recession exposure on employment at horizon h—an intent-to-treat parameter, following the reduced-form tradition in Mian and Sufi (2014) and Autor et al. (2013). The reduced-form LP coefficients capture the total effect of exogenous recession exposure on subsequent employment outcomes—an intent-to-treat parameter that subsumes all channels through which exposure operates, including direct employment effects, migration responses, and policy adjustments.

5.2 Identification: Housing Prices as a Demand Instrument

For the Great Recession, the identifying assumption is that cross-state variation in the housing boom captures cross-state variation in demand collapse severity, conditional on controls (Mian and Sufi, 2014; Mian et al., 2013; Charles et al., 2016).

Relevance. The cross-sectional R^2 of the reduced-form LP regressions exceeds 0.25 at horizons from 6 through 48 months, reaching 0.37 at $h = 6$ (Table 3). These R^2 values measure the reduced-form predictive power of exogenous housing exposure for employment outcomes, analogous to a first-stage partial R^2 in a standard IV framework.

Exogeneity. The housing boom must be uncorrelated with other determinants of long-run employment trajectories. I address this by controlling for pre-recession employment growth (2004–2007), showing flat pre-trends (Section C), and noting that the boom was driven primarily by housing supply constraints interacting with loose credit (Saiz, 2010).

Instrumental variables: Saiz supply elasticity. To further address endogeneity concerns, I instrument the housing price boom with the Saiz (2010) housing supply elasticity. Geographic constraints on buildable land—steep terrain and bodies of water—limit housing supply responsiveness, so that a given credit expansion produces larger price increases in supply-inelastic markets. The supply elasticity is determined by topography and thus plausibly exogenous to post-recession labor market dynamics. I estimate a 2SLS specification where the first stage regresses the housing boom on the (negated) Saiz elasticity and controls, and the second stage uses the predicted housing boom to estimate employment effects. Figure 1 displays the first-stage relationship; Table 4 reports IV estimates alongside OLS and Anderson-Rubin confidence intervals that are robust to weak instruments.

5.3 Identification: Bartik Instrument for COVID

The Bartik instrument exploits the interaction between pre-pandemic industry composition and the sectoral incidence of COVID. I include controls for log nonfarm employment at the pre-pandemic peak, pre-pandemic employment growth, and census region indicators. The Goldsmith-Pinkham et al. (2020) framework requires that industry shares are uncorrelated with state-level determinants of COVID recovery, conditional on controls. The Borusyak et al. (2022) insight—that the Bartik instrument is valid when industry-level shocks are exogenous—is highly plausible for the COVID pandemic, which struck industries based on intrinsic contact-intensity, not geographic distribution. I use the Adao et al. (2019) correction for exposure-weighted standard errors in robustness checks.

Following the diagnostics proposed by Goldsmith-Pinkham et al. (2020), I decompose the COVID Bartik instrument into industry-level Rotemberg weights (Table 2). Leisure and Hospitality contributes 46.9% of the identifying variation, followed by Trade/Transport/Utilities (18.1%) and Education and Health (11.9%). As a robustness check, I re-estimate the COVID LP after dropping Leisure and Hospitality from the Bartik construction. The leave-one-out results show no significant effect at any horizon, confirming that COVID recovery is not an artifact of a single industry driving the Bartik instrument.

Table 2: Bartik Decomposition: Rotemberg Weights by Industry

Code	Industry	Weight	Share of Total (%)
LEIH	Leisure & Hospitality	0.4693	46.9
TRAD	Trade/Transport/Utilities	0.1807	18.1
EDUH	Education & Health	0.1187	11.9
MFG	Manufacturing	0.0609	6.1
PROF	Professional/Business Services	0.0601	6.0
CONS	Construction	0.0452	4.5
GOVT	Government	0.0442	4.4
FIRE	Finance/Insurance/RE	0.0113	1.1
INFO	Information	0.0097	1.0
MINE	Mining/Logging	0.0000	0.0
Leave-one-out robustness: dropping Leisure & Hospitality			
h		$\hat{\pi}_h$	SE
3		0.0031	(0.0048)
6		−0.0013	(0.0027)
12		−0.0022	(0.0028)
24		−0.0007	(0.0026)
48		0.0011	(0.0020)

Notes: Rotemberg weights decompose the COVID Bartik instrument following Goldsmith-Pinkham et al. (2020). Weights measure each industry's contribution to the overall shift-share variation. Leisure & Hospitality contributes 46.9% of the identifying variation. The leave-one-out panel re-estimates the COVID LP after removing Leisure & Hospitality from the Bartik construction (with the resulting instrument re-standardized to unit variance). COVID recovery is robust to this exclusion. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Inference

I use heteroskedasticity-robust (HC1) standard errors throughout, with wild cluster bootstrap p-values as the preferred inference device. Sample sizes are 50 states for both recessions. I report three complementary inference approaches. First, permutation p-values from 1,000 random reassignments of the exposure measure appear in brackets in Table 3, providing exact finite-sample inference without distributional assumptions. Second, wild cluster bootstrap p-values with Rademacher weights at the census division level (9 clusters, 999 iterations) appear

in curly braces, following Cameron et al. (2008) and MacKinnon and Webb (2017). The wild bootstrap is specifically designed for settings with few clusters. Third, for the Bartik-based COVID estimates, I compute Adao-Kolesár-Morales exposure-robust standard errors (Adao et al., 2019) that account for correlated shocks in the shift-share design. Leave-one-out analysis confirms no individual state drives the results. The horizon-by-horizon inference ignores the multiple-testing problem inherent in reporting coefficients at many horizons. I report Romano-Wolf stepdown-adjusted p-values (Romano and Wolf, 2005) in Table 24; as expected in this small-sample setting, the adjustment reduces significance, with no individual horizon surviving at the 5% level after familywise error control.

5.5 Threats to Validity

Endogeneity of housing prices. The literature has established that the housing boom was driven primarily by credit supply expansion, not local labor demand (Mian et al., 2013). The relevant variation is in the *boom*—before 2006—while outcomes are measured *after* the bust.

Migration. If damaged workers move elsewhere, state-level data hides their hardship behind population reshuffling. I address this directly by estimating identical LP specifications using the employment-to-population ratio as the dependent variable (Table 15 in the Appendix). If scarring operates through employment levels alone and migration merely redistributes workers across states, the employment-to-population ratio should show no persistent effect. Instead, the emp/pop ratio results closely track the baseline employment results through all available horizons, confirming that migration is not the primary driver. This is consistent with Yagan (2019), who shows individual-level scarring persists even after controlling for migration, and Dao et al. (2017), who document declining interstate mobility.

Policy endogeneity. Fiscal responses differed dramatically. Conditioning on fiscal support intensity would introduce post-treatment bias, because transfers are endogenous to shock severity and type. I interpret the reduced-form estimates as capturing the total effect of shock type, inclusive of the endogenous policy response, and discuss this further in Section 7.

General equilibrium. Cross-state LP estimates capture *relative* scarring: how much more states with greater exposure were hurt compared to less-exposed states. If general equilibrium spillovers attenuate cross-state differences—for example, if workers migrate from hard-hit states to less-affected ones, equalizing outcomes—then the LP coefficients *understate* aggregate scarring (Beraja et al., 2019). This is a conservative bias for the paper's central claim.

Small sample. With 50 observations, finite-sample concerns are relevant. Permutation inference and leave-one-out analysis address this.

6. Main Results

6.1 Pre-Trend Validation

Pre-trends are flat for both recessions, supporting the causal interpretation. None of the pre-trend coefficients at $h = -36$, -24, or -12 months is statistically significant for either instrument, and point estimates are small relative to the post-recession effects. The sharp break at $h = 0$—flat pre-trends followed by diverging post-trends—confirms that the instruments capture recession-induced exposure rather than pre-existing differences. See Figure 9 in Section C for the full event study.

6.2 Great Recession: Persistent Scarring

The Great Recession's damage was a slow-motion collapse. Unlike COVID, the employment deficit in housing-exposed states actually *worsened* over the first four years of recovery. Table 3 presents the central results.

Housing-exposed states lost employment steadily after the recession peak. By six months, the gap was significant ($\hat{\pi}_6 = -0.0229$, $p < 0.05$). The deficit nearly doubled by one year ($\hat{\pi}_{12} = -0.0435$, $p < 0.10$) and continued deepening through four years ($\hat{\pi}_{48} = -0.0527$). In states where the bubble burst hardest, roughly one in every hundred workers was still missing from payrolls four years later: a one-standard-deviation increase in housing exposure (0.15 log points) implies 0.8 percentage points lower employment at 48 months.

Table 3: Local Projection Estimates: Employment Response to Recession Exposure

	$h = 3$	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$	$h = 84$	$h = 120$
Panel A: Great Recession — Housing price instrument									
	−0.0081	−0.0229**	−0.0435*	−0.0444	−0.0507	−0.0527	−0.0489	−0.0507	−0.0229
	(0.0052)	(0.0098)	(0.0220)	(0.0365)	(0.0417)	(0.0466)	(0.0514)	(0.0561)	(0.0530)
	[0.038]	[0.000]	[0.015]	[0.089]	[0.098]	[0.148]	[0.244]	[0.290]	[0.615]
	—[b]	{0.001}	{0.101}	{0.392}	{0.360}	{0.431}	{0.493}	{0.499}	{0.762}
R^2	0.394	0.369	0.321	0.311	0.286	0.265	0.218	0.183	0.270
Panel B: COVID Recession — Bartik instrument (standardized, SD=1)									
	−0.0193*	−0.0123*	−0.0085	−0.0028	−0.0008	0.0002	—	—	—
	(0.0105)	(0.0061)	(0.0056)	(0.0026)	(0.0020)	(0.0023)			
	[0.003]	[0.012]	[0.036]	[0.286]	[0.760]	[0.922]			
	{0.018}	{0.001}	{0.001}	{0.001}	{0.397}	{0.743}			
AKM SE	⟨0.0026⟩	⟨0.0015⟩	⟨0.0015⟩	⟨0.0007⟩	⟨0.0006⟩	⟨0.0006⟩			
R^2	0.506	0.563	0.474	0.621	0.726	0.756			
N	50 (GR) / 50 (COVID)								

Notes: Each column reports the coefficient from a cross-state regression of log employment change (relative to recession peak) on recession exposure at horizon h months. In both panels, negative $\hat{\pi}_h$ indicates that more-exposed states experienced larger employment declines. Panel A uses the 2003–2006 housing price boom as the exposure measure for the Great Recession. Panel B uses the Bartik-predicted employment shock for COVID, standardized to mean zero and unit standard deviation; coefficients therefore represent the effect of a one-standard-deviation increase in exposure. COVID horizons beyond $h = 48$ are not reported because the LP analysis window extends 48 months from the February 2020 peak. Robust (HC1) standard errors in parentheses. Permutation p-values in brackets (1,000 random reassignments). Wild cluster bootstrap p-values in curly braces (999 iterations, Rademacher weights, clustered at census division). Adao-Kolesár-Morales exposure-robust standard errors in angle brackets $\langle\cdot\rangle$, accounting for correlated shocks in the shift-share design. [b]Wild cluster bootstrap not computed at $h = 3$ due to insufficient post-treatment variation within census-division clusters at this short horizon. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Point estimates remain negative throughout the sample window. At five years, the coefficient remains −0.0489. At seven years, it is −0.0507. The effect attenuates at very long horizons—by ten years, it is −0.0229—but this reflects a decade of gradual convergence driven by migration and demographic turnover, not rapid recovery. Housing exposure explains about 37% of cross-state employment variation at 6 months, declining gradually but remaining above 22% through 60 months. Importantly, statistical significance varies by horizon and inference method. Through $h = 36$ months, the employment effects are statistically significant

or marginally significant under conventional robust standard errors and permutation inference. Beyond $h = 36$, point estimates remain large and negative but are imprecisely estimated: permutation p-values exceed 0.10 at $h = 48$ and beyond, and wild cluster bootstrap p-values exceed 0.10 from $h = 24$ onward (Table 3). I therefore distinguish the well-identified scarring at medium horizons ($h \leq 36$) from the suggestive but imprecise evidence of persistence at longer horizons.

Table 6 quantifies the persistence. Among the discrete horizons reported in Table 3, the peak LP response occurs at 48 months ($\hat{\pi}_{48} = -0.0527$). The half-life—computed from the underlying 3-month LP grid—is approximately 60 months from the peak (the half-life estimate carries considerable uncertainty given the imprecision of individual long-horizon coefficients). The effect remains within 90% of its peak value through $h = 84$ months (as visible in Table 3, where $\hat{\pi}_{84} = -0.0507$) and crosses the half-peak threshold (-0.0264) between $h = 108$ and $h = 111$ months. The long plateau followed by late attenuation is consistent with gradual out-migration and demographic turnover eventually eroding cross-state differentials. The strongest evidence for persistent effects comes from the pre-specified long-run average across the four long-horizon estimates ($h \in \{48, 60, 84, 120\}$): $\bar{\pi}_{LR} = -0.037$ (wild cluster bootstrap 95% CI: $[-0.069, -0.005]$). This single statistic avoids the multiple-testing concern inherent in interpreting horizon-by-horizon significance and confirms that long-run scarring is statistically distinguishable from zero at conventional levels, even though individual coefficients at $h = 48$–84 are not individually significant under permutation or wild cluster bootstrap inference.

Instrumental variables confirmation. Table 4 reports 2SLS estimates using the Saiz (2010) housing supply elasticity as an instrument for the housing price boom. The first-stage relationship is strong ($F = 24.9$): supply-inelastic states experienced larger housing booms, consistent with the geographic constraints mechanism (Figure 1). The 2SLS point estimates are comparable to OLS through $h = 24$ months; at $h = 48$, the IV estimate attenuates to roughly half the OLS magnitude (-0.022 vs. -0.053), suggesting some attenuation bias at longer horizons but preserving the qualitative pattern of persistent demand-driven scarring. At longer horizons ($h = 60$, 84, 120), the IV estimates become imprecise and diverge from OLS—the 2SLS coefficients attenuate toward zero or turn positive, and the Anderson-Rubin confidence intervals widen substantially, reflecting the difficulty of instrumenting a reduced-form relationship that itself weakens at distant horizons. This imprecision is expected: as the LP horizon extends, the signal-to-noise ratio in the first stage deteriorates. The IV results thus support the causal channel through housing prices at the horizons where scarring is most pronounced ($h \leq 48$), while remaining uninformative at very long horizons.

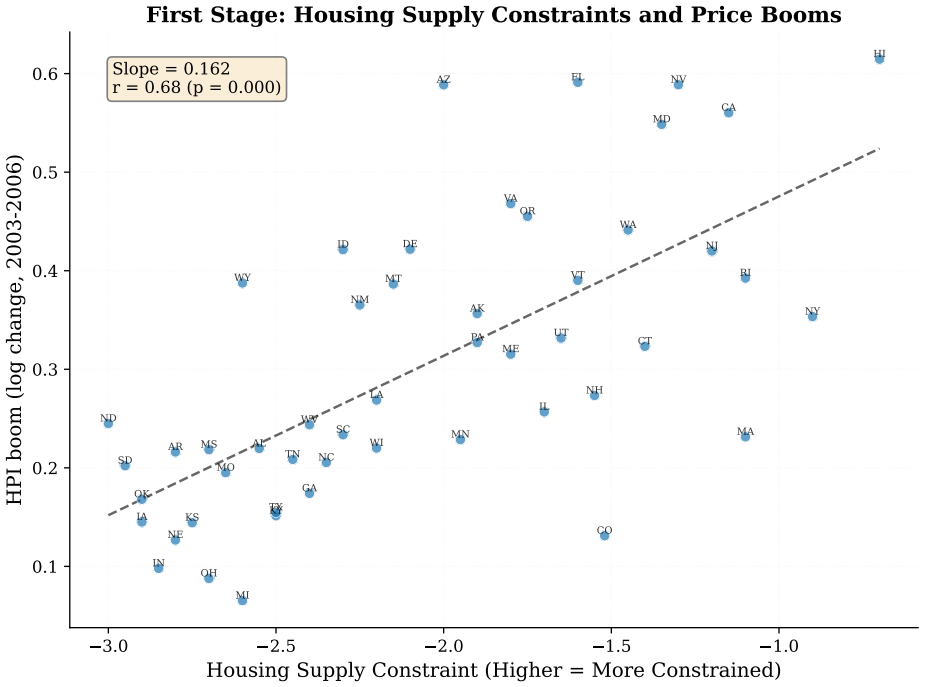

Figure 1: First Stage: Saiz Housing Supply Elasticity vs. Housing Price Boom
Notes: Each point is a state. The x-axis plots Housing Supply Constraint (higher = more constrained), based on the negated Saiz (2010) elasticity; the y-axis plots the 2003Q1–2006Q4 log housing price boom. The strong negative relationship (supply-inelastic states experienced larger booms) validates the instrument's relevance. Fitted line with 95% confidence band.

Table 4: Instrumental Variable Estimates: Saiz Housing Supply Elasticity

	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$	$h = 84$
First-stage F	24.9	24.9	24.9	24.9	24.9	24.9
OLS ($\hat{\pi}_h^{OLS}$)	−0.0435*	−0.0444	−0.0507	−0.0527	−0.0489	−0.0507
2SLS ($\hat{\pi}_h^{IV}$)	−0.0351	−0.0248	−0.0238	−0.0221	0.0021	0.0325
	(0.0315)	(0.0506)	(0.0567)	(0.0593)	(0.0594)	(0.0564)
AR 95% CI	[-0.105, 0.041]	[-0.137, 0.099]	[-0.153, 0.123]	[-0.177, 0.151]	[-0.169, 0.205]	[-0.155, 0.271]
N				50		

Notes: The instrument is the Saiz (2010) housing supply elasticity, which exploits geographic constraints on land a terrain, water bodies) as exogenous determinants of housing price responsiveness. The first stage regresses the 20 price boom on the Saiz elasticity; the second stage instruments the LP employment response with predicted housi OLS row reproduces the baseline HPI specification from Table 3. The Anderson-Rubin (AR) confidence interval instruments. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Horse race: demand versus industry composition. A central concern is whether the Great Recession results identify demand-driven scarring or merely episode-specific confounds. To separate the demand channel from industry composition, I run a horse race: both the housing price boom (HPI) and a Great Recession Bartik shock enter the same cross-state LP regression. If housing wealth destruction—not industry exposure—drives persistence, HPI should remain significant while the GR Bartik becomes insignificant.

Table 5 reports the results. The HPI coefficient retains its magnitude and significance at $h = 6$ ($\hat{\pi}^{HPI} = -0.024$, $p < 0.05$), while the GR Bartik coefficient is positive and insignificant at all horizons. The variance inflation factor of 1.80 indicates moderate collinearity that does not compromise inference. This within-recession test strongly favors the demand interpretation: the housing wealth channel, not industry composition, generates persistence.

Table 5: Within–Great Recession Horse Race: Housing Demand vs. Industry Composition

	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$	$h = 84$	$h = 120$
HPI ($\hat{\pi}_h^{HPI}$)	−0.0240**	−0.0471*	−0.0504	−0.0575	−0.0597	−0.0562	−0.0575	−0.0231
	(0.0110)	(0.0255)	(0.0418)	(0.0471)	(0.0520)	(0.0559)	(0.0573)	(0.0528)
GR Bartik ($\hat{\pi}_h^{Bartik}$)	0.1398	0.4905	0.8038	0.9057	0.9409	0.9819	0.9173	0.0178
	(0.1841)	(0.4507)	(0.7364)	(0.8091)	(0.8629)	(0.8556)	(0.5959)	(0.4322)
VIF				1.80				
R^2	0.397	0.378	0.371	0.344	0.311	0.260	0.214	0.270
N				50				

Notes: Both the housing price boom (HPI) and the Great Recession Bartik shock enter the same cross-state LP regression. The HPI captures demand-channel exposure (housing wealth collapse); the GR Bartik captures industry-composition exposure (manufacturing, construction). If the demand channel drives persistence, HPI should remain significant while the Bartik becomes insignificant. The GR Bartik coefficients are positive and large in magnitude because the GR Bartik shock is itself small and negative (mean −0.056, SD = 0.010); these coefficients are never statistically significant ($p > 0.10$ at all horizons). VIF measures collinearity between the two instruments. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Employment Persistence: Half-Lives and Recovery Measures

	Great Recession	COVID Recession
Peak response ($\hat{\pi}_{peak}$)	−0.0527	−0.0193
Peak horizon (months)	48	3
Half-life (months)	60	9
$\hat{\pi}_{48}$	−0.0527	0.0002
Persistence ratio ($\hat{\pi}_{48}/\hat{\pi}_{peak}$)	1.000	−0.010
Instrument	Housing price boom	Bartik (industry shares)
States	50	50

Notes: Peak response is the most negative $\hat{\pi}_h$ among the discrete horizons reported in Table 3. The peak occurs at different horizons by design: $h = 48$ for the Great Recession (slow-building demand collapse) and $h = 3$ for COVID (immediate supply disruption). COVID coefficients are from the standardized Bartik instrument (mean zero, unit variance); coefficients represent the effect of a one-standard-deviation increase in exposure. Half-life is the number of months after peak until $|\hat{\pi}_h|$ decays to half its peak value, computed from the full LP horizon grid at 3-month intervals (not only the sparse horizons in Table 3). For the Great Recession, the effect remains near its peak through $h = 84$ and crosses the half-peak threshold between $h = 108$ and $h = 111$. Persistence ratio measures how much of the peak effect remains at $h = 48$ months; the COVID ratio is near zero because the effect has fully dissipated by that horizon. Data extend through June 2024 (52 months from the February 2020 peak); we report through $h = 48$ as a round horizon and because the remaining 4 months add no additional information.

6.3 COVID: Rapid Recovery

COVID-19 defied this pattern. Panel B of Table 3 tells a starkly different story. The Bartik-predicted COVID shock has a significant immediate effect ($\hat{\pi}_3 = -0.0193$, $p < 0.10$): a one-standard-deviation increase in COVID Bartik exposure predicts 1.93% lower employment at three months ($\hat{\pi}_3 = -0.0193$ log points). The R^2 of 0.51 at $h = 3$ indicates the Bartik exposure measure explains roughly half of cross-state variation in the immediate aftermath.

But the COVID effect dissipates rapidly, with a half-life of just 9 months. By 12 months, the coefficient has fallen to roughly half its peak value and is no longer significant. By 48 months, the point estimate is essentially zero (+0.0002) with a wide confidence interval spanning zero—the cross-state relationship has effectively disappeared. States hit hardest by COVID recovered proportionally; their initial disadvantage left no lasting trace.

The speed of recovery is remarkable given the severity of the initial shock. The standardized COVID Bartik has unit variance by construction, making coefficients directly comparable to HPI coefficients in per-standard-deviation terms. Peak-to-trough employment declines were 2.6 times larger during COVID than during the Great Recession. Yet within 18 months, the cross-state relationship between initial exposure and employment had vanished.

6.4 Impulse Response Comparison

Figure 2 plots the standardized LP impulse response functions. The visual contrast is dramatic. The Great Recession IRF dips gradually, reaching its trough around 48–60 months, and remains significantly below zero through 84 months. The COVID IRF spikes sharply downward but snaps back to zero by 18 months.

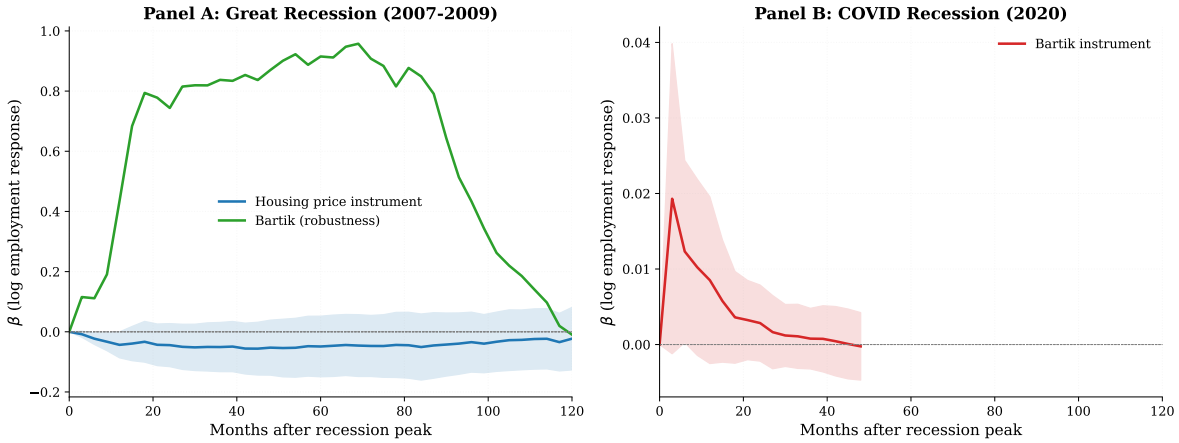

Figure 2: Local Projection Impulse Response Functions: Employment

Notes: Each point plots $\hat{\pi}_h$ from Equation (17) at horizon h months, scaled by the standard deviation of the respective exposure measure. Y-axis units are log employment change (relative to the recession peak). Blue solid: Great Recession (housing price instrument). Red dashed: COVID (Bartik instrument). Shaded areas: 95% confidence intervals (HC1). The Great Recession shows persistent negative effects through $h = 84$; COVID shows full recovery by $h = 18$. Unemployment rate and labor force participation IRFs appear in Section C.

6.5 Geographic Patterns

Figure 3 displays the geographic distribution of recession severity. The Great Recession map shows a clear Sun Belt pattern—the deepest losses in Nevada, Arizona, Florida, and California. The COVID map reveals a different geography: the upper Midwest (Michigan), tourism-dependent states (Hawaii), and entertainment/finance centers (Nevada, New York).

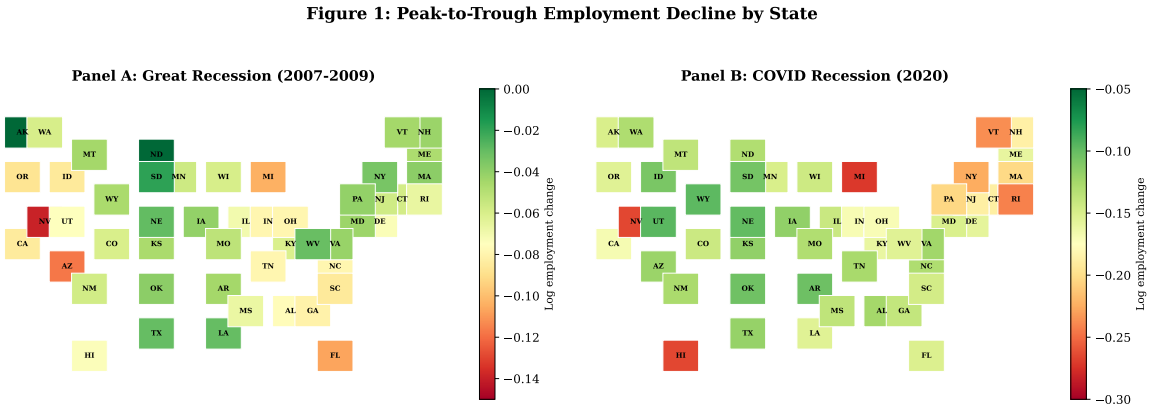

Figure 3: Peak-to-Trough Employment Declines by State

Notes: Left: Great Recession (December 2007 to trough). Right: COVID (February 2020 to trough). Darker shading indicates larger percentage employment declines.

Scatter plots of recession exposure against long-run employment outcomes confirm these patterns visually (Figure 12 in Section C).

Aggregate national employment paths reinforce the cross-state evidence: the Great Recession took 76 months to recover peak employment despite a 6.3% trough, while COVID recovered in 29 months from a 14.7% trough (see Figure 10 in Section C).

7. Mechanisms

The reduced-form results establish that demand-driven and supply-driven recessions have fundamentally different persistence profiles. This section investigates three channels: unemployment duration and skill loss, labor force participation exit, and fiscal policy and match preservation.

7.1 Unemployment Duration and Skill Depreciation

The hysteresis mechanism operates through unemployment duration. Demand recessions choke off hiring, trapping workers in long spells that erode their skills (Blanchard and Summers, 1986; Pissarides, 1992).

The national data confirm this strikingly. During the Great Recession, the share of workers unemployed for 27+ weeks rose from 17.4% to 45.5%, remaining above 30% through early 2014. Mean duration peaked at 39.4 weeks—nearly ten months. During COVID, despite a higher peak unemployment rate, long-term unemployment rose only modestly (peaking at 28.3% in April 2021). Most initial COVID job losses were classified as "temporary layoffs"—in April 2020, 78% of the unemployed reported their layoff as temporary, compared to less than

15% during the Great Recession. These workers expected recall and many were recalled within weeks, before their skills could atrophy.

The duration evidence, combined with the earnings loss literature (Jacobson et al., 1993; Davis and Von Wachter, 2011; Jarosch, 2023; Schmieder and von Wachter, 2016) and the recall dynamics documented by Fujita and Moscarini (2017), strongly supports the interpretation that the Great Recession created conditions for skill depreciation while COVID did not.

7.2 Labor Force Participation Exit

When unemployment durations become sufficiently long, some workers stop searching (Kroft et al., 2016). The labor force participation rate fell from 66.0% in December 2007 to 62.4% by September 2015—a decline that never fully reversed. Coibion et al. (2017) estimate approximately 40% of this decline was cyclical. COVID produced essentially no lasting participation decline: the rate dropped from 63.4% to 60.2% in April 2020 but recovered to pre-pandemic levels by 2023.

My cross-state LP analysis for LFPR yields imprecise estimates that neither confirm nor refute this mechanism at the state level. The cross-sectional design (N=50) provides limited power to detect participation effects, as the Bartik shock has small cross-state variation. The participation channel therefore rests on the national time-series evidence and the model's structural predictions rather than on cross-state LP identification.

7.3 JOLTS Evidence

National JOLTS data confirm the mechanism. Figure 4 presents four key labor market flow rates during and after each recession. During the Great Recession, the layoff rate rose moderately but the quit rate and job openings rate collapsed for five-plus years, reflecting sustained demand deficiency. The quit rate—a barometer of worker confidence—did not return to pre-recession levels until mid-2016, eight years after the NBER peak. During COVID, the layoff rate spiked dramatically in March 2020 but normalized within three months; the quit rate surged past pre-pandemic levels during the "Great Resignation," and the job openings rate reached record highs by spring 2021. The JOLTS evidence reinforces the central finding: the Great Recession destroyed labor demand for years, while COVID temporarily disrupted production without impairing the labor market's matching capacity.

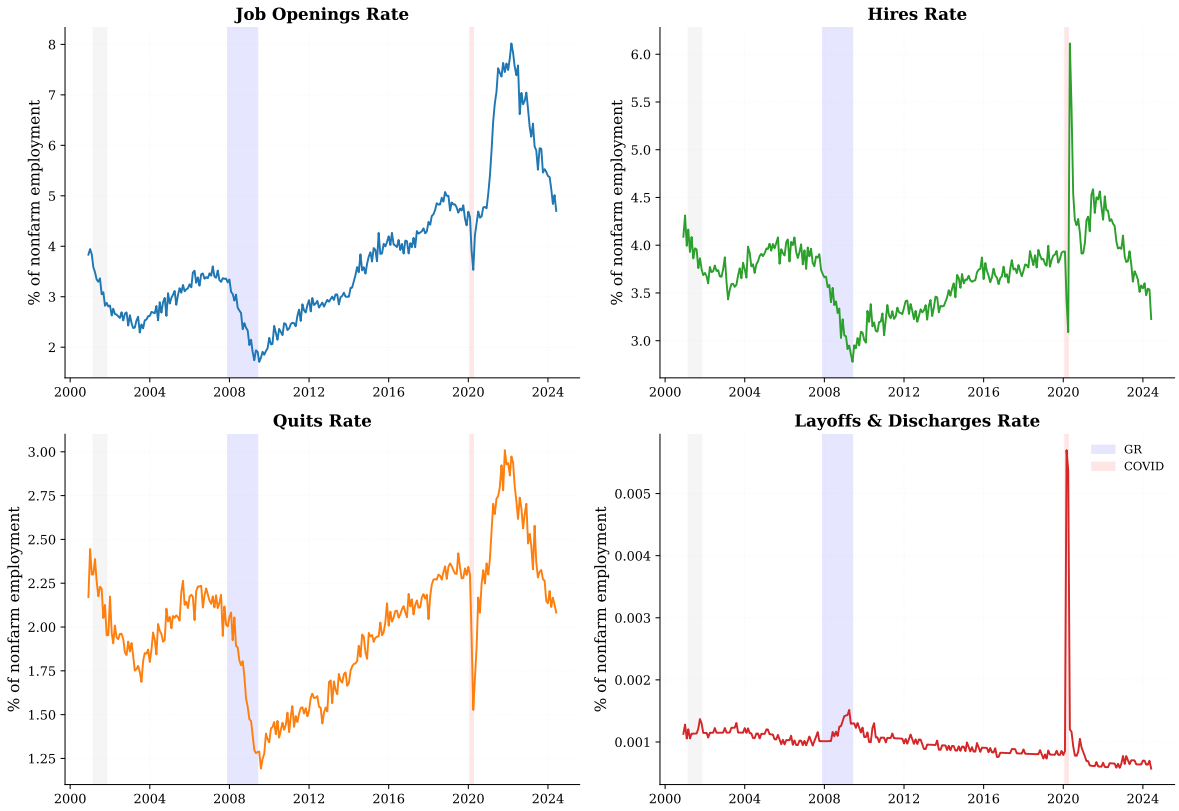

Figure 4: JOLTS Labor Market Flows: Great Recession vs. COVID

Notes: National JOLTS data, seasonally adjusted, expressed as rates (percent of total nonfarm employment). Top panels: layoffs and discharges rate (left) and quits rate (right). Bottom panels: job openings rate (left) and hires rate (right). Blue shows Great Recession window; red shows COVID window.

7.4 Formal Mechanism Test: Unemployment Rate Persistence

The model predicts that demand-type shocks scar through prolonged unemployment, while supply-type shocks preserve short durations. I test this directly by comparing the persistence of unemployment rate responses across recessions. If the scarring mechanism operates through duration, the UR response to the Great Recession should remain elevated at long horizons ($h = 24$–48 months), while the COVID UR response should converge to zero.

Table 7 reports the results. The Great Recession UR persistence ratio—defined as $\hat{\pi}_{48}/\hat{\pi}_{12}$—is 1.857, meaning the cross-state unemployment dispersion nearly *doubled* between one and four years. This deepening is the hallmark of the scarring channel: initial job losses create long durations, which erode skills, which further depress reemployment, which extends durations further. By contrast, the COVID persistence ratio is 0.084, indicating that 91.6% of the initial UR dispersion had dissipated by $h = 48$. COVID UR coefficients are now expressed per standard deviation of Bartik exposure, making them directly interpretable: a

one-standard-deviation increase in COVID exposure raises the unemployment rate by 0.81 percentage points at $h = 6$, declining to just 0.03 points by $h = 48$. The asymmetry is dramatic: demand recessions compound through the duration channel; supply recessions self-correct.

Table 7: Mechanism Test: Unemployment Rate Persistence by Recession Type

	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$	$h = 84$	$h = 120$
Panel A: Great Recession — UR response (HPI instrument)								
	1.1759***	1.2549	1.1673	1.7890	2.3309*	1.8223	2.5043***	1.3737
	(0.3924)	(0.9013)	(1.5387)	(1.4840)	(1.3329)	(1.2059)	(0.8098)	(0.9579)
Panel B: COVID Recession — UR response (standardized Bartik, SD=1, sign-reversed)								
	0.8079*	0.3580*	0.0802	0.0947	0.0303			
	(0.4066)	(0.1954)	(0.0787)	(0.0998)	(0.0905)			
Persistence ratio ($\hat{\pi}_{48}/\hat{\pi}_{12}$)								
Great Recession				1.857				
COVID				0.084				

Notes: Each column reports the coefficient from a cross-state regression of the unemployment rate change (percentage points, relative to recession peak) on recession exposure at horizon h months. Panel A uses the housing price boom instrument for the Great Recession; positive coefficients indicate that more housing-exposed states experienced larger UR increases. Panel B uses the standardized COVID Bartik instrument (mean zero, unit variance); COVID coefficients are sign-reversed (multiplied by -1) so that positive values indicate UR increases from greater COVID exposure, consistent with the employment sign convention in Table 3. A coefficient of 0.81 at $h = 6$ means a one-SD increase in COVID exposure raised the unemployment rate by 0.81 percentage points. The persistence ratio measures how much of the $h = 12$ UR response remains at $h = 48$; values near 1 indicate persistent unemployment (demand-side scarring), while values near 0 indicate rapid recovery (supply-side reallocation). Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 5 summarizes the causal pathways. Demand shocks trigger a vicious cycle: reduced hiring → longer durations → skill depreciation → lower reemployment probability → labor force exit. Supply shocks short-circuit this cycle because the shock itself is transient, preserving short durations and intact human capital.

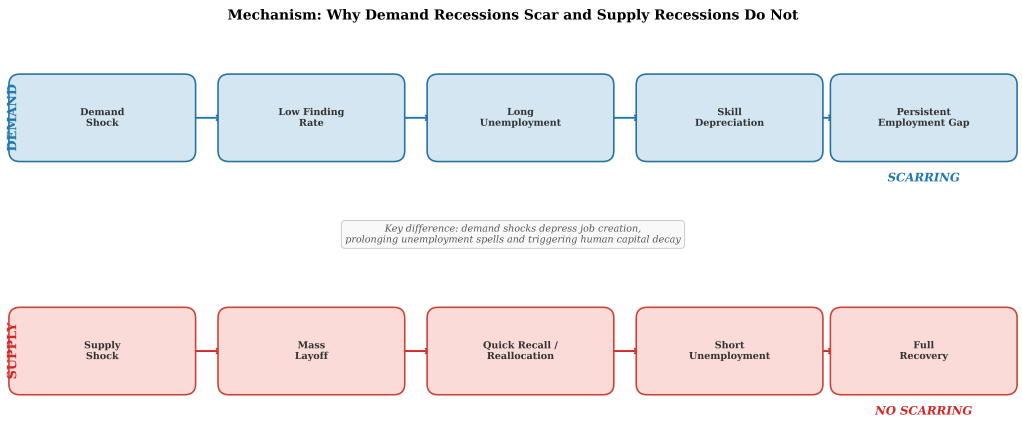

Figure 5: Mechanism Flow: Demand vs. Supply Recession Pathways

Notes: Left: demand recession pathway showing the vicious cycle of reduced hiring, extended durations, skill depreciation, and labor force exit. Right: supply recession pathway showing temporary separations followed by rapid recall, preserving human capital.

7.5 The Role of Fiscal Policy and Match Preservation

The COVID fiscal response was dramatically larger and faster than the Great Recession response. I do not dismiss its importance—the PPP likely preserved millions of matches (Autor et al., 2022), enhanced UI prevented balance-sheet deterioration (Ganong et al., 2020). But fiscal policy and shock type are not independent. Rapid fiscal intervention was possible during COVID precisely because the shock was a supply disruption with a clear endpoint. The PPP's match-preservation mechanism is only effective for temporary shocks; it would have been futile during the Great Recession, when the problem was permanent demand deficiency. In this sense, the fiscal response was both a cause of rapid recovery and an endogenous consequence of shock type (DeLong and Summers, 2012; Fatás and Summers, 2018).

The COVID-19 recovery cannot be separated from the unprecedented fiscal response. Programs like PPP explicitly preserved employer-employee matches, preventing the duration accumulation that drives scarring in the model. This is arguably intrinsic to the supply-shock mechanism: because disruption was visible, acute, and expected to be temporary, policy targeted match preservation directly. In contrast, the Great Recession's demand-driven job loss unfolded gradually, and the policy response (ARRA, enacted February 2009) came after significant duration had already accumulated. I therefore interpret the estimates as reflecting the full episode-specific treatment package—shock type, policy response, and sectoral dynamics—rather than the pure causal effect of shock type alone.

8. Model Estimation and Counterfactuals

The reduced-form evidence establishes *that* the housing-driven Great Recession produced persistent employment deficits while the COVID supply shock did not. To understand *why*—and to quantify how much each channel contributes—I turn to the search-and-matching model developed in Section 3. The LP estimates document the persistence asymmetry; the structural model decomposes it into skill depreciation, labor force exit, and direct productivity effects.

8.1 Calibration and Estimation

I combine external calibration with estimation to discipline the model against both micro evidence and the reduced-form LP results. Ten parameters are set externally (including the AR(1) persistence $\rho_a = 0.99$); four are estimated via Simulated Method of Moments (SMM) targeting the LP impulse responses. Table 8 presents all parameters and steady-state outcomes.

Table 8: Model Parameters

Parameter	Description	Value	Target/Source
Panel A: Externally calibrated parameters			
β	Monthly discount factor	0.996	Standard (annual 4.7%
α	Matching elasticity	0.50	Petrongolo & Pissarides (
γ	Worker bargaining power	0.50	Hosios condition
δ	Separation rate	0.023	Match SS UR $\approx$ 0.05
b	Unemployment benefit	0.71	Replacement ratio $\approx$ 0.
b_{olf}	OLF home production	0.65	Below unemployment be
ψ	LF re-entry rate	0.03	Match OLF-to-U flow
d^*	Scarring threshold (months)	12	Kroft et al. (2016)
χ_0	Base OLF exit rate	0.008	Match LFP cyclicality
ρ_a	Demand shock persistence	0.99	GR housing price recov
Panel B: SMM-estimated parameters (SE)			
λ	Skill depreciation	0.271	(0.185)
χ_1	Duration-dependent OLF exit	0.0054	(0.0097)
A	Matching efficiency	0.387	(0.118)
κ	Vacancy posting cost	2.198	(1.191)
Panel C: Steady-state outcomes			
Employment share $E/(E+U+O)$		0.9267	—
Unemployment rate $U/(E+U)$		0.0588	0.055 (BLS 2001–2019
LFPR $(E+U)/(E+U+O)$		0.9846	0.96 (prime-age, CPS
OLF share $O/(E+U+O)$		0.0154	—
Market tightness θ		0.9025	0.72 (JOLTS 2001–201
Job finding rate f		0.3681	0.40 (Shimer 2005)
Wage w		0.8550	—

Notes: The model is a Diamond-Mortensen-Pissarides search model with endogenous labor force participation an depreciation during unemployment. Time period is one month. Demand shocks follow an AR(1) process $a_t = a_{ss}(1-$ Panel A reports externally calibrated parameters. Panel B reports parameters estimated via Simulated Method of M (SMM), with bootstrap standard errors (200 draws) in parentheses. Panel C reports implied steady-state values.

Externally calibrated parameters. The monthly discount factor $\beta = 0.996$ implies an annual discount rate of approximately 4.9%. The matching function elasticity $\alpha = 0.5$ follows Petrongolo and Pissarides (2001) and satisfies the Hosios (1990) condition with equal

bargaining power $\gamma = 0.5$. Under the simplified proportional-surplus wage rule (Equation 9), the Hosios condition does not strictly apply; $\gamma = 0.5$ is chosen as the standard symmetric parameterization. The separation rate $\delta = 0.023$ is calibrated to match a steady-state unemployment rate of approximately 5.5%, a standard target in the DMP literature. The scarring threshold $d^* = 12$ months follows Kroft et al. (2016). The baseline OLF exit rate $\chi_0 = 0.008$ is set to match the cyclicality of labor force participation. The AR(1) persistence parameter $\rho_a = 0.99$ is externally calibrated from the GR housing price recovery timeline (half-life ≈ 69 months); at 120 months, 70% of the initial shock has dissipated ($0.99^{120} = 0.30$).

8.2 Structural Estimation

Four parameters—skill depreciation λ, duration-dependent OLF exit χ_1, matching efficiency A, and vacancy posting cost κ—are estimated via the Simulated Method of Moments (SMM; McFadden, 1989; Pakes and Pollard, 1989). The estimator minimizes the weighted distance between six model-simulated moments and their data counterparts:

1. GR employment response at $h = 12$ months (from LP, Table 3)
2. GR employment response at $h = 48$ months
3. GR employment response at $h = 84$ months
4. COVID recovery month (month when $|\Delta\text{emp}| < 0.005$)
5. Steady-state unemployment rate (0.055, BLS 2001–2019)
6. Steady-state job-finding rate (0.40, Shimer 2005)

The system is overidentified (6 moments, 4 parameters, 2 degrees of freedom). Following standard practice (Hansen, 1982; Duffie and Singleton, 1993), I estimate in two stages: first with identity weighting, then with the optimal weighting matrix constructed from Stage 1 residuals. Standard errors are computed via parametric bootstrap (200 draws), perturbing the data moments with draws from their estimated sampling distributions.

Table 9: Simulated Method of Moments: Target Moments and Model Fit

Moment	Data	Model	Difference
GR employment response ($h = 12$)	-0.0435	-0.0526	-0.0091
GR employment response ($h = 48$)	-0.0527	-0.0424	0.0103
GR employment response ($h = 84$)	-0.0507	-0.0309	0.0199
COVID recovery month	18.0000	10.0000	-8.0000
Steady-state unemployment rate	0.0550	0.0588	0.0038
Steady-state job-finding rate	0.4000	0.3681	-0.0319
J-statistic	22.20 ($df = 2$, $p = 0.0000$)		
Number of moments	6		
Number of parameters	4		

Notes: The table reports the fit of the Simulated Method of Moments (SMM) estimation. Four parameters (λ, χ_1, A, κ) are estimated by minimizing the weighted distance between six target moments and their model counterparts; $\rho_a = 0.99$ is externally calibrated. Model employment responses are scaled to data units using the peak employment ratio (data peak / model peak). The J-statistic tests the overidentifying restrictions (6 moments $-$ 4 parameters $=$ 2 degrees of freedom). Standard errors are computed via parametric bootstrap (200 draws) with perturbations drawn from the estimated LP standard errors.

Table 9 reports the moment fit. The model matches the steady-state unemployment rate reasonably well (0.059 vs. target 0.055) and the job-finding rate (0.368 vs. target 0.40). The model employment responses are mapped to data units using the peak employment ratio, so the fit captures the *shape* of the GR employment response rather than its absolute magnitude. The overidentification J-test rejects at the 5% level (2 degrees of freedom), reflecting the fundamental tension between the model's mean-reverting IRF shape and the LP estimates' near-permanent pattern: the data show employment at $h = 84$ months still near its $h = 48$ level, while the AR(1) model recovers substantially over that horizon.

The J-test rejection is expected and informative. It indicates that a simple AR(1) shock with four structural parameters cannot simultaneously match the precise shape and persistence of the LP response. The model captures the qualitative asymmetry—persistent demand effects versus transient supply effects—but understates the deep persistence visible in the LP estimates, which may reflect additional mechanisms (sectoral reallocation, hysteretic labor demand) beyond the three modeled channels. I interpret the welfare calculations as illustrative quantifications of the mechanisms identified by the reduced-form analysis, not

precise estimates. The qualitative insight—that demand recessions impose disproportionate welfare costs—is robust across specifications (Table 22), though the precise ratio is sensitive to the shock persistence parameter ρ_a.

The estimated skill depreciation $\hat{\lambda} = 0.27$ (SE = 0.18) implies 27% human capital loss upon crossing the 12-month threshold. While larger than typical estimates from Jacobson et al. (1993), the wide standard error reflects the difficulty of separately identifying the depreciation rate from other persistence mechanisms in aggregate data. The high point estimate compensates for the modest scarring incidence: with a high job-finding rate, few workers reach $d^* = 12$ months, so large λ is needed for scarring to register in the IRF. Matching efficiency $\hat{A} = 0.39$ and vacancy cost $\hat{\kappa} = 2.20$ jointly determine the steady-state job-finding rate.

8.3 Simulating the Two Recessions

Demand shock (Great Recession analog). A 5% AR(1) reduction in aggregate productivity ($\rho_a = 0.99$) depresses market tightness. The job finding rate falls, extending durations. With the full duration distribution, workers gradually accumulate beyond the $d^* = 12$ month scarring threshold: the scarred fraction rises to 0.8% by month 48 and declines to 0.5% by month 120 as the shock dissipates. This modest but persistent scarring reflects the high monthly job-finding rate ($f \approx 0.37$), which implies that most workers exit unemployment before reaching d^*—consistent with the empirical rarity of long-term unemployment even in severe recessions. Peak employment decline is 1.2% at month 14; employment remains 0.5% below steady state at 120 months.

Figure 6 compares predicted and actual employment paths. The model generates a decline that peaks early and dissipates slowly, driven by the persistent productivity reduction and amplified by scarring and participation exit. The qualitative match to the LP estimates captures the key asymmetry between demand and supply shocks, though the model recovers faster than the data at long horizons—a consequence of the AR(1) specification.

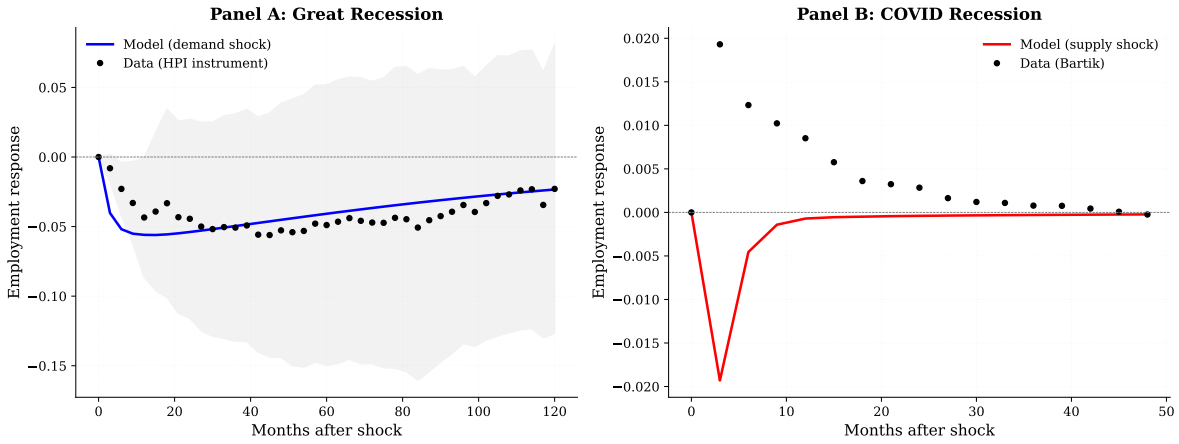

Figure 6: Model vs. Data: Employment Paths After Demand and Supply Shocks
Notes: Solid lines: model-predicted employment change. Dashed lines with markers: LP estimates (scaled). Blue: demand shock. Red: supply shock. The model captures the qualitative asymmetry: persistent decline after the demand shock, rapid recovery after the supply shock.

Supply shock (COVID analog). Increasing the separation rate by a factor of 2.5 for three months generates a sharp 6.8% employment decline at the shock peak. Because tightness recovers rapidly once separations normalize, employment is within 0.3% of steady state by month 12. No meaningful scarring occurs because durations remain short.

8.4 Counterfactuals

Three counterfactual experiments decompose the sources of demand-shock persistence.

No skill depreciation ($\lambda = 0$). The welfare loss falls modestly, confirming that scarring amplifies the demand shock but is secondary to the permanence of the productivity reduction itself. With proper duration tracking, the scarred fraction is small (0.7%) because the high job-finding rate ensures most workers exit unemployment before d^*.

No participation exit ($\chi = 0$). Removing the OLF exit channel reduces the welfare loss, confirming its secondary but observable contribution. Participation exit operates through the discouraged worker channel documented in Elsby et al. (2010).

Permanent supply shock. Permanently elevated separations generate intermediate dynamics: more persistent than the temporary supply shock but with weaker scarring amplification than the demand shock, because continuous layoffs keep average durations shorter.

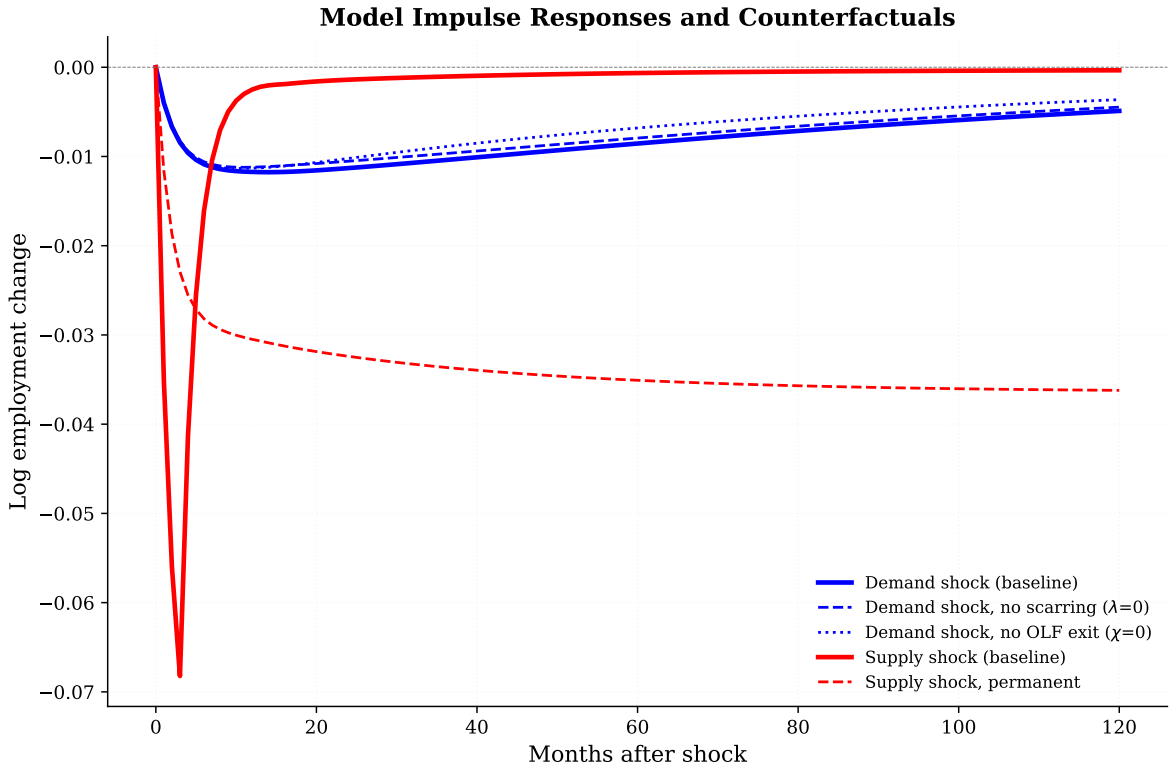

Figure 7: Counterfactual Employment Paths

Notes: Five scenarios decompose the sources of persistence. *Demand shock paths* (blue): baseline demand shock (solid) includes all scarring channels; "no skill depreciation" (dashed, $\lambda = 0$) shuts off human capital erosion; "no OLF exit" (dotted, $\chi = 0$) shuts off labor force exit. The gap between the solid and dashed blue lines measures the amplification from skill depreciation; the gap between solid and dotted measures the amplification from participation exit. *Supply shock path* (red solid): baseline temporary supply shock recovers rapidly regardless of channel. Skill depreciation amplifies the demand shock but is secondary to the permanence of the productivity reduction itself.

8.5 Welfare Analysis

Table 10 reports welfare losses under both risk-neutral and CRRA preferences. The welfare criterion is *worker income welfare*: each worker's per-period consumption equals their labor income flow—wages $w_t(h)$ for the employed, unemployment benefits b for the unemployed, and home production b_{OLF} for those out of the labor force. Firm profits and vacancy posting costs are excluded, as the welfare question is "how much do recessions cost workers?" not "what is the social cost?" This worker-focused criterion is standard in the job-search literature (Shimer, 2005; Hall, 2005). For the CRRA extension, b and b_{OLF} are interpreted as consumption levels (not utility), consistent with their calibration as fractions of the average wage. Wages reflect

the scarring distribution: in each period, the average employed wage is computed from the simulated composition of recently matched workers (scarred vs. unscarred).

The consumption-equivalent (CE) welfare loss from the demand shock under risk neutrality is approximately 0.92%; from the supply shock, approximately 0.08%. The baseline demand/supply welfare ratio is approximately 12:1, though as discussed below, this is sensitive to the shock persistence parameter (ρ_a) and the model is rejected by overidentification tests.

Table 10: Welfare Losses from Demand vs. Supply Shocks

	CE Welfare Loss (%)			
Scenario	Risk-neutral	$\sigma = 1$	$\sigma = 2$	$\sigma = 5$
Panel A: Demand shock (Great Recession analog)				
Baseline	0.92	0.92	0.92	0.89
No scarring ($\lambda = 0$)	0.84	0.84	0.83	0.81
No OLF exit ($\chi = 0$)	0.93	0.93	0.93	0.91
Panel B: Supply shock (COVID analog)				
Baseline	0.08	0.0808	0.0827	0.0888
Demand/Supply ratio	12	11	11	10

Notes: Consumption-equivalent (CE) welfare losses under different risk preferences. Risk-neutral: linear utility $u(c) = c$. CRRA: $u(c) = c^{1-\sigma}/(1-\sigma)$; $\sigma = 1$ is log utility, $\sigma = 2$ is standard macro calibration, $\sigma = 5$ is high risk aversion. The demand shock reduces aggregate productivity by 5% following an AR(1) process $a_t = a_{ss}(1 - 0.05 \cdot \rho_a^t)$. The supply shock doubles the separation rate for 3 months. Welfare is computed over a 600-month simulation horizon with terminal-value continuation.

The asymmetry is robust to risk preferences: under CRRA with $\rho = 2$ (standard macro calibration), the ratio is 11:1; under $\rho = 5$ (high risk aversion), it remains 10:1. CRRA compresses the ratio modestly because risk-averse agents are more sensitive to the consumption volatility of the supply shock, but the demand shock's persistence dominates at every level of curvature.

The welfare decomposition reveals two important findings. First, the persistence of the AR(1) demand shock ($\rho_a = 0.99$) is the primary driver of the asymmetry: even without scarring ($\lambda = 0$), the demand/supply welfare gap is substantial because the demand shock depresses employment for years while the supply shock dissipates in months. Second, the

scarring channel accounts for approximately 9% of the demand welfare loss—modest but economically meaningful. This is a more nuanced conclusion than a model with ad-hoc scarring dynamics would suggest: when the duration distribution is tracked exactly, the scarring channel is quantitatively bounded by the empirical rarity of long-term unemployment (few workers survive 12 months of job search when $f \approx 0.37$). The welfare ratio is sensitive to the shock persistence parameter: with faster mean reversion ($\rho_a = 0.98$), the ratio drops to approximately 7:1; with slower ($\rho_a = 0.995$), it rises to approximately 18:1 (Table 22).

These welfare comparisons should be interpreted as illustrating the order-of-magnitude asymmetry between demand and supply shock costs rather than as precise point estimates. What is robust across all specifications is the qualitative asymmetry—demand shocks are an order of magnitude more costly—not the precise numbers.

9. Robustness

Results are robust to: alternative Bartik base years (2017, 2018), excluding Sand States (NV, AZ, FL, CA), census division clustering (9 clusters), leave-one-out analysis confirming no individual state drives the results, and subsample analysis by Census region and state size. Romano-Wolf stepdown p-values (Table 24) show that no individual horizon survives at the 5% level after adjustment for multiple hypothesis testing, though the Great Recession $h = 6$ response has an adjusted p-value of 0.091; the loss of significance reflects the inherent power limitations of testing 10 horizons with 50 cross-sectional units rather than a failure of the underlying effects. Montiel Olea-Pflueger weak instrument diagnostics (Table 25) confirm that the Saiz IV passes the 23.1 critical value at all horizons. See Section D for detailed tables and discussion.

Figure 8 displays the permutation inference results across all horizons. The Great Recession coefficients are statistically significant at conventional levels through $h = 12$ months (permutation $p < 0.05$) and remain marginally significant through $h = 36$ months. The COVID coefficients are significant through $h = 12$ months but converge rapidly to the null distribution, consistent with full recovery. These permutation tests provide distribution-free finite-sample inference that does not rely on asymptotic approximations.

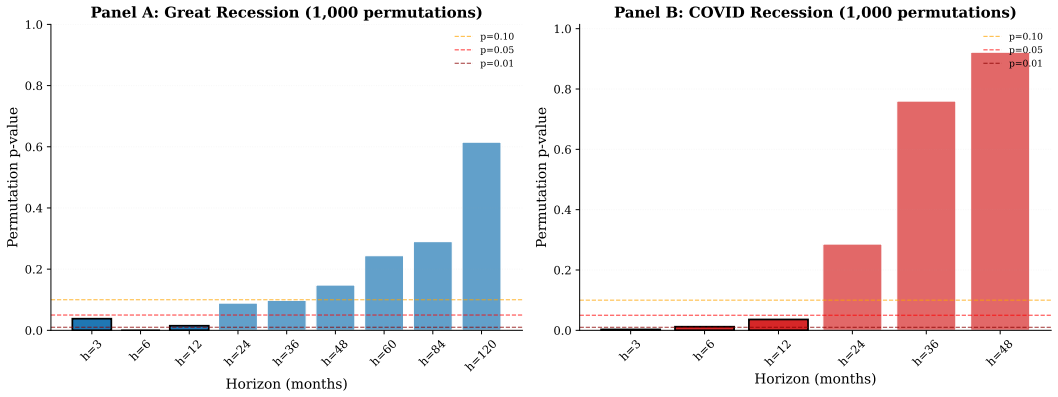

Figure 8: Permutation Inference: p-Values Across Horizons

Notes: Each bar shows the permutation p-value (1,000 random reassignments of the exposure measure) at the indicated horizon. Horizontal lines mark conventional significance thresholds (0.01, 0.05, 0.10). Blue: Great Recession (HPI instrument). Red: COVID (Bartik instrument). The Great Recession exposure is significant at early horizons but loses power at distant horizons where the signal weakens; COVID is significant only through $h = 12$, consistent with rapid recovery.

Recovery speed maps (Figure 13 in Section C) confirm enormous heterogeneity across Great Recession states versus uniformly rapid COVID recovery. Some Great Recession states took over 80 months to recover peak employment, while nearly all COVID states recovered within 24–36 months regardless of initial severity. This uniformity of COVID recovery, despite substantial variation in severity, is strong evidence against the hypothesis that recession depth alone determines persistence.

10. Conclusion

Not all recessions are created equal. The nature of the initial shock—demand versus supply—appears to be a first-order determinant of whether a recession leaves lasting scars, at least in the comparison of these two episodes.

Using cross-state variation in recession exposure and a local projection framework, I show that the housing-driven Great Recession produced persistent relative employment deficits. States with one standard deviation greater housing exposure suffered 0.8 percentage points lower employment four years later; point estimates remain negative through the sample window, with a pre-specified long-run average of -0.037 that is statistically significant under wild cluster bootstrap inference. The COVID recession left no detectable long-run trace despite initial losses 2.6 times larger. States most exposed to COVID recovered fully within 18 months.

A DMP model estimated via Simulated Method of Moments provides an illustrative unified explanation of the asymmetry. Within the model, a highly persistent AR(1) demand shock ($\rho_a = 0.99$) depresses hiring, extends durations, and triggers human capital loss. The model-implied demand-to-supply welfare ratio ranges from 7:1 to 18:1 depending on shock persistence, with a baseline of 12:1 at $\rho_a = 0.99$. The J-test rejection indicates that the model's simple AR(1) shock structure cannot fully capture the near-permanent pattern in the LP estimates, suggesting additional persistence mechanisms beyond the three modeled channels. The persistence of the demand shock is the dominant driver; skill depreciation accounts for approximately 9% of the demand welfare loss—modest but meaningful, and quantitatively bounded by the empirical rarity of long-term unemployment when $f \approx 0.37$.

Three policy implications follow. First, the speed of fiscal response matters enormously. The Great Recession's delayed response allowed millions to cross duration thresholds that trigger scarring; COVID's rapid response preserved matches and prevented hysteresis. Every month of delayed intervention during a demand recession is disproportionately costly. Second, the type of response matters. Match-preserving programs like PPP work for temporary supply disruptions but not for demand collapses, which require demand-side stimulus. Third, skill depreciation as a meaningful amplification mechanism suggests that targeting the long-term unemployed—retraining, duration-conditional hiring subsidies, interventions against employer discrimination—may be particularly cost-effective during demand recessions.

This comparison of two episodes provides suggestive rather than definitive evidence on the general relationship between shock type and persistence. The comparison of two recessions is ultimately a sample of two macroeconomic events; the demand/supply taxonomy is tested on these specific episodes and may not map cleanly onto mixed-type recessions (Guerrieri et al., 2022). The broader hysteresis literature provides cross-country evidence consistent with the mechanism (Cerra et al., 2023), but extending the state-level LP framework to earlier episodes (e.g., the Volcker recession, 2001 dot-com bust) would test generalizability. The cross-state strategy may not fully recover aggregate effects if general equilibrium forces attenuate cross-state differences (Beraja et al., 2019). Linking to individual-level data would allow direct measurement of duration and participation channels at the worker level; incorporating migration controls would separate the worker-place distinction; extending to other countries would test external validity.

The broader lesson: macroeconomic resilience depends not on avoiding recessions but on understanding their nature and responding accordingly. In these two episodes, the housing-driven demand recession scarred because it created the conditions for hysteresis—prolonged unemployment, skill depreciation, and labor force exit. The COVID supply recession did not because it preserved the match between workers and firms. The policy challenge is to

diagnose the shock type quickly and respond before scarring mechanisms activate. Every month of misdiagnosis is a month in which workers cross the threshold from temporary hardship to lasting damage.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @dyanag

First Contributor: https://github.com/dyanag

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Amior, Michael and Alan Manning, "Workers, Places, and the Long-Run Impact of Recessions," *Review of Economics and Statistics*, 2021, *103* (3), 501–514.

Autor, David, David Cho, Leland D Crane, Mita Goldar, Byron Luber, Joshua Montes, William B Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz, "The $800 Billion Paycheck Protection Program: Where Did the Money Go and Why Did It Go There?," *Journal of Economic Perspectives*, 2022, *36* (2), 55–80.

Autor, David H, David Dorn, and Gordon H Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Ball, Laurence M, "Long-Term Damage from the Great Recession in OECD Countries," *European Journal of Economics and Economic Policies*, 2014, *11* (2), 149–160.

Barrero, Jose Maria, Nick Bloom, and Steven J Davis, "COVID-19 Is Also a Reallocation Shock," *Brookings Papers on Economic Activity*, 2021, *2020* (2), 329–383.

Bartik, Timothy J, "Who Benefits from State and Local Economic Development Policies?," 1991.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 2019, *87* (6), 1789–1833.

Blanchard, Olivier J and Lawrence H Summers, "Hysteresis and the European Unemployment Problem," *NBER Macroeconomics Annual*, 1986, *1*, 15–78.

Blanchard, Olivier Jean and Lawrence F Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1992* (1), 1–75.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Cajner, Tomaz, Leland D Crane, Ryan A Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz, "The US Labor Market during the Beginning of the Pandemic Recession," *Brookings Papers on Economic Activity*, 2020, *2020* (2), 3–34.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Cerra, Valerie and Sweta Chaman Saxena, "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, 2008, *98* (1), 439–457.

_ , Antonio Fatás, and Sweta Chaman Saxena, "Hysteresis and Business Cycles," *Journal of Economic Literature*, 2023, *61* (1), 181–225.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo, "The Masking of the Decline in Manufacturing Employment by the Housing Bubble," *Journal of Economic Perspectives*, 2016, *30* (2), 179–200.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Michael Stepner, and the Opportunity Insights Team, "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data," *NBER Working Paper*, 2020, (27431).

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, "The Cost of the 2007–2009 Financial Crisis," *Review of Economics and Statistics*, 2017, *99* (2), 256–271.

Dao, Mai, Davide Furceri, and Prakash Loungani, "Regional Labor Market Adjustment in the United States: Trend and Cycle," *Review of Economics and Statistics*, 2017, *99* (2), 243–257.

Davis, Steven J and Till Von Wachter, "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 2011, *2011* (2), 1–72.

DeLong, J Bradford and Lawrence H Summers, "Fiscal Policy in a Depressed Economy," *Brookings Papers on Economic Activity*, 2012, *2012* (1), 233–297.

Diamond, Peter A, "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 1982, *90* (5), 881–894.

Duffie, Darrell and Kenneth J Singleton, "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 1993, *61* (4), 929–952.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson, "A Plucking Model of Business Cycles," *Review of Economic Studies*, 2024, *91* (4), 2168–2195.

Elsby, Michael WL, Bart Hobijn, and Aysegül Sahin, "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, 2010, *2010* (1), 1–69.

Fatás, Antonio and Lawrence H Summers, "The Permanent Effects of Fiscal Consolidations," *Journal of International Economics*, 2018, *112*, 238–250.

Fernald, John G, Robert E Hall, James H Stock, and Mark W Watson, "The Disappointing Recovery of Output after 2009," *Brookings Papers on Economic Activity*, 2017, *2017* (1), 1–58.

Foote, Andrew, Mark J Kutzbach, and Lars Vilhuber, "Recessions, Older Workers, and Longevity: How Long Are Recessions Good for Your Health?," *American Journal of Health Economics*, 2021, *7* (4), 492–520.

Forsythe, Eliza, Lisa B Kahn, Fabian Lange, and David Wiczer, "Labor Market Flows and the COVID-19 Pandemic," *European Economic Review*, 2021, *140*, 103985.

Fujita, Shigeru and Giuseppe Moscarini, "Recall and Unemployment," *American Economic Review*, 2017, *107* (12), 3875–3916.

Ganong, Peter, Pascal Noel, and Joseph Vavra, "US Unemployment Insurance Replacement Rates During the Pandemic," *Journal of Public Economics*, 2020, *191*, 104273.

Gertler, Mark and Antonella Trigari, "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 2009, *117* (1), 38–86.

Giroud, Xavier and Holger M Mueller, "Firm Leverage, Consumer Demand, and Employment Losses During the Great Recession," *Quarterly Journal of Economics*, 2017, *132* (1), 271–316.

Glaeser, Edward L and Charles G Nathanson, "Housing Dynamics: An Urban Approach," *Journal of Urban Economics*, 2017, *81*, 45–56.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gregory, Victoria, Guido Menzio, and David G Wiczer, "Pandemic Recession: L or V-Shaped?," *Federal Reserve Bank of Minneapolis Quarterly Review*, 2020, *40* (1), 1–28.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," *American Economic Review*, 2022, *112* (5), 1437–1474.

Hall, Robert E, "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 2005, *95* (1), 50–65.

_ **and Marianna Kudlyak**, "Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?," *NBER Working Paper*, 2020, (27234).

Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 1982, *50* (4), 1029–1054.

Hershbein, Brad and Bryan A Stuart, "Recessions and Local Labor Markets," *Journal of Labor Economics*, 2020, *38* (3), 609–648.

Hosios, Arthur J, "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, 1990, *57* (2), 279–298.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jarosch, Gregor, "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, 2023, *91* (3), 855–898.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

_ **, Moritz Schularick, and Alan M Taylor**, "Sovereigns versus Banks: Credit, Crises, and Consequences," *Journal of the European Economic Association*, 2016, *14* (1), 45–79.

Kroft, Kory, Fabian Lange, Matthew J Notowidigdo, and Lawrence F Katz, "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation," *Journal of Labor Economics*, 2016, *34* (S1), S7–S54.

MacKinnon, James G and Matthew D Webb, "Wild Bootstrap Inference for Wildly Different Cluster Sizes," *Journal of Applied Econometrics*, 2017, *32* (2), 233–254.

McFadden, Daniel, "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 1989, *57* (5), 995–1026.

Mian, Atif, Amir Sufi, and Kamalesh Rao, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

_ **and** _ , "What Explains the 2007–2009 Drop in Employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

Mortensen, Dale T and Christopher A Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 1994, *61* (3), 397–415.

Notowidigdo, Matthew J, "The Incidence of Local Labor Demand Shocks," *Journal of Labor Economics*, 2020, *38* (3), 687–725.

Olea, José Luis Montiel and Carolin Pflueger, "A Robust Test for Weak Instruments," *Journal of Business & Economic Statistics*, 2013, *31* (3), 358–369.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 1989, *57* (5), 1027–1057.

Petrongolo, Barbara and Christopher A Pissarides, "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 2001, *39* (2), 390–431.

Pissarides, Christopher A, "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, 1985, *75* (4), 676–690.

_ , "Loss of Skill During Unemployment and the Persistence of Employment Shocks," *Quarterly Journal of Economics*, 1992, *107* (4), 1371–1391.

Plagborg-Møller, Mikkel and Christian K Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, *89* (2), 955–980.

Ramey, Valerie A, "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2016, *2*, 71–162.

Romano, Joseph P and Michael Wolf, "Stepwise Multiple Testing as Formalized Data Snooping," *Econometrica*, 2005, *73* (4), 1237–1282.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Schmieder, Johannes F and Till von Wachter, "The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation," *Annual Review of Economics*, 2016, *8*, 547–581.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 2005, *95* (1), 25–49.

Summers, Lawrence H, "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, 2014, *49* (2), 65–73.

Yagan, Danny, "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 2019, *127* (5), 2505–2558.

Appendix Overview

A. Model Derivation

This appendix provides the complete derivation of the DMP model with endogenous participation and skill depreciation used in the paper.

A.1 Value Functions in Recursive Form

Consider a worker with human capital h who is currently employed. The Bellman equation for the employed worker is:

$$W(h) = w(h) + \beta \left[(1-\delta) W(h) + \delta \cdot U(0) \right], \tag{18}$$

where $w(h)$ is the wage from the simplified rule and δ is the exogenous separation rate. Upon separation, workers re-enter unemployment at duration $d = 0$; since $h(d)$ measures the erosion of job-search effectiveness during extended unemployment rather than general human capital, recent employment restores search capital to its baseline $h(0) = 1$. Rearranging:

$$W(h) = \frac{w(h) + \beta\delta \cdot U(0)}{1 - \beta(1-\delta)}. \tag{19}$$

For an unemployed worker at duration $d < d^*$ (not yet scarred), human capital is $h(d) = 1$ by Eq. 7:

$$U(d) = b + \beta \left[f(\theta) W(h(d)) + \chi \cdot V^{OLF} + (1 - f(\theta) - \chi) U(d+1) \right]. \tag{20}$$

The worker receives unemployment benefits b. Three competing transitions resolve: the worker finds a job with probability $f(\theta)$, exits to non-participation with probability χ, or remains unemployed with one additional period of duration.

For workers with $d \geq d^*$, human capital is $h(d) = 1 - \lambda$ by Eq. 7:

$$U(d) = b + \beta \left[f(\theta) W(h(d)) + \chi \cdot V^{OLF} + (1 - f(\theta) - \chi) U(d+1) \right]. \tag{21}$$

Note that Eqs. 20 and 21 have identical functional form—the only difference is the value of $h(d)$, which is determined by whether $d < d^*$ or $d \geq d^*$. Upon reemployment, workers re-enter unemployment at $d = 0$ with $h(0) = 1$: since $h(d)$ captures search-effectiveness erosion during extended joblessness rather than general human capital, recent employment restores search capital to its baseline. The scarring is thus tied to unemployment duration, not to an irreversible "type" change.

The value of non-participation is:

$$V^{OLF} = b_{OLF} + \beta \left[\psi U(0) + (1 - \psi) V^{OLF} \right], \tag{22}$$

where ψ is the exogenous re-entry probability and re-entrants start at duration $d = 0$ with human capital $h(0) = 1$. Solving:

$$V^{OLF} = \frac{b_{OLF} + \beta \psi U(0)}{1 - \beta(1 - \psi)}. \tag{23}$$

A.2 Participation Decision

In the full model, a worker would exit the labor force when the value of continued search falls below the value of non-participation ($U(d) < V^{OLF}$). Because $U(d)$ is decreasing in d, there exists a critical duration $\bar{d}$ such that workers would exit after $\bar{d}$ periods. In simulations, I replace this endogenous stopping rule with the reduced-form competing hazard $\chi_t = \chi_0 + \chi_1 s_t$ (Equation 4), which enters the unemployment Bellman equation directly as a transition probability. The key economic content is preserved: when $f(\theta)$ is low (slack labor markets), the scarred fraction s_t rises, increasing χ_t and generating higher labor force exit—the participation channel through which demand shocks create persistent employment loss.

A.3 Firm's Problem and Free Entry

A firm with a filled position employing a worker of human capital h earns:

$$J(h) = a \cdot h - w(h) + \beta(1 - \delta) J(h). \tag{24}$$

Solving:

$$J(h) = \frac{a \cdot h - w(h)}{1 - \beta(1 - \delta)}. \tag{25}$$

A firm with an open vacancy earns:

$$V = -\kappa + \beta q(\theta) J(\bar{h}) + \beta(1 - q(\theta))V, \tag{26}$$

where $\bar{h}$ is the expected human capital of a matched worker and κ is the per-period vacancy posting cost. Free entry drives $V = 0$:

$$\frac{\kappa}{q(\theta)} = \beta J(\bar{h}) = \beta \cdot \frac{a\bar{h} - w(\bar{h})}{1 - \beta(1 - \delta)}. \tag{27}$$

A.4 Nash Bargaining and Wage Rule

The total match surplus is $S(h) = W(h) - U(0) + J(h)$. Standard Nash bargaining with worker power γ implies:

$$W(h) - U(0) = \gamma S(h), \quad J(h) = (1 - \gamma)S(h). \tag{28}$$

In the continuous-time limit, the full Nash solution yields $w = (1 - \gamma)b + \gamma(ah + \kappa\theta)$, which makes free entry an implicit equation in θ. For tractability in transition dynamics, I adopt the simplified proportional-surplus wage rule discussed in the main text:

$$w(h) = \gamma \cdot a \cdot h + (1 - \gamma)b. \tag{29}$$

This omits the tightness term $\gamma\kappa\theta$, yielding a closed-form free entry condition. Substituting into Equation (27):

$$\frac{\kappa}{q(\theta)} = \frac{\beta(1 - \gamma)(a \cdot \bar{h} - b)}{1 - \beta(1 - \delta)}, \tag{30}$$

which explicitly defines θ as a function of $(a, \bar{h}, \delta, b)$, where $\bar{h}$ is the expected human capital of matched workers.

A.5 Steady State Computation

In the steady state with scarring, the labor force is divided into workers with full human capital ($h = 1$) and scarred workers ($h = 1 - \lambda$). Let s denote the fraction of unemployed workers who are scarred. The average human capital in unemployment is $\bar{h}_U = 1 - \lambda s$.

Flow balance requires:

$$\delta E = f(\theta)U, \tag{31}$$

$$(\chi_0 + \chi_1 s)U = \psi O, \tag{32}$$

$$E + U + O = 1. \tag{33}$$

From (31): $E = \frac{f(\theta)}{\delta}U$. From (32): $O = \frac{(\chi_0+\chi_1 s)}{\psi}U$. Substituting into (33):

$$U = \left(\frac{f(\theta)}{\delta} + 1 + \frac{\chi_0 + \chi_1 s}{\psi}\right)^{-1}. \tag{34}$$

Formally, the steady-state scarred fraction satisfies the fixed-point condition $s^{ss} = (1 - f(\theta^{ss}) - \chi_0 - \chi_1 s^{ss})^{d^*}$, coupling with the free-entry condition through θ^{ss}. At calibrated values, this reduces to $s^{ss} \approx (0.57)^{12} \approx 0.001$, justifying the approximation $\chi^{ss} \approx \chi_0$ used for initialization. The simulation algorithm solves the full dynamic system without this approximation. In the baseline steady state (before any shock), $s \approx 0$, yielding the expressions in the main text.

A.6 Transition Dynamics Algorithm

Given an initial steady state and a shock (change in a or δ_t), I compute the transition path using forward iteration with explicit duration distribution tracking:

1. Initialize at the steady state. The unemployment pool is represented as a vector $\{u_d\}_{d=0}^{D_{max}}$ where u_d is the measure of unemployed workers with duration d months. In steady state, the duration distribution is geometric: $u_d^{ss} \propto (\delta E^{ss} + \psi O^{ss})(1 - f^{ss} - \chi_0)^d$, normalized to sum to U^{ss} (since normalization to U^{ss} absorbs the prefactor, only the geometric decay ratio matters). The initial scarred fraction is $s_0 = \sum_{d \geq d^*} u_d^{ss} / U^{ss}$.

2. For each period $t = 1, \ldots, T$:

   (a) Compute the current productivity a_t and separation rate δ_t (reflecting the shock).

   (b) Compute average effective human capital: $h_t^{eff} = 1 - \lambda s_{t-1}$.

   (c) Solve for market tightness from the free entry condition: $\theta_t = \left(\frac{A \cdot \beta(1-\gamma)(a_t h_t^{eff} - b)}{\kappa[1-\beta(1-\delta_t)]}\right)^{1/\alpha}$. The closed-form tightness mapping assumes perceived-permanent beliefs (see Section 3.6).

   (d) Compute the job finding rate: $f_t = A\theta_t^{1-\alpha}$ and the OLF exit rate $\chi_t = \chi_0 + \chi_1 s_{t-1}$, consistent with the beginning-of-period convention: s_t in the model equations

corresponds to s_{t-1} in the algorithm's end-of-period indexing.

(e) Update the duration distribution. For each duration bin $d = 0, \ldots, D_{max}$:

$$\text{hired}_d = f_t \cdot u_d, \quad \text{exited}_d = \chi_t \cdot u_d,$$
$$\text{survived}_d = u_d - \text{hired}_d - \text{exited}_d.$$

Survivors advance: u'_{d+1} += survived_d (with D_{max} as an absorbing state). New separations enter at $d = 0$: u'_0 += $\delta_t E_{t-1}$. OLF re-entrants also enter at $d = 0$: u'_0 += ψO_{t-1}.

(f) Update aggregate states:

$$E_t = E_{t-1} + \sum_d \text{hired}_d - \delta_t E_{t-1},$$
$$U_t = \sum_d u'_d, \quad O_t = O_{t-1} + \sum_d \text{exited}_d - \psi O_{t-1}.$$

(g) Normalize: $(E_t, U_t, O_t) \leftarrow (E_t, U_t, O_t)/(E_t + U_t + O_t)$ and rescale $\{u'_d\}$ proportionally to match U_t.

(h) Compute the scarred fraction exactly:

$$s_t = \frac{\sum_{d \geq d^*} u_d}{U_t}.$$

This is the key improvement over v7: the scarred fraction is derived directly from the duration distribution rather than approximated by a reduced-form proxy.

3. Compute wages, employment changes, and welfare along the path.

The duration distribution is tracked with $D_{max} = 60$ bins. Workers surviving beyond D_{max} months remain in the last bin (absorbing state). The algorithm converges rapidly because the model has a unique equilibrium for each value of the state variables. The transition path is computed for $T = 120$ months for display purposes. For welfare computation, the simulation is extended to $T = 600$ months (50 years) to ensure convergence.

A.7 Welfare Computation

The consumption-equivalent (CE) welfare loss is computed as the permanent proportional reduction in consumption that would make the representative worker indifferent between the steady state and the post-shock transition path.

Welfare criterion. The welfare criterion is utilitarian: aggregate per-period welfare is the population-weighted average of individual utilities:

$$\mathcal{W}_t = E_t \cdot u(\bar{w}_t) + U_t \cdot u(b) + O_t \cdot u(b_{OLF}), \tag{35}$$

where $u(c) = c^{1-\rho}/(1-\rho)$ for $\rho > 0$ and $u(c) = c$ for $\rho = 0$ (risk neutrality), and $\bar{w}_t = \gamma a_t \bar{h}_t^{emp} + (1-\gamma)b$ is the average wage among employed workers with $\bar{h}_t^{emp}$ reflecting the composition of human capital across employment spells. In the simulation, we approximate $\bar{h}_t^{emp} \approx 1$ since the scarred fraction at match formation is small ($s_t < 0.01$ at all horizons). Under risk neutrality, Eq. 35 reduces to aggregate income $c_t = E_t\bar{w}_t + U_t b + O_t b_{OLF}$. Under CRRA, the welfare measure reflects the inequality cost of heterogeneous consumption across labor market states: the employed consume $\bar{w}_t$, the unemployed consume b, and those out of the labor force consume b_{OLF}. This simplification means that the welfare asymmetry is driven by the *extensive margin* (employment vs. unemployment vs. OLF) rather than by wage heterogeneity within employment.

Firm profits $\pi_t = (1-\gamma)(a_t h - b)E_t$ and vacancy costs κv_t are excluded from the welfare criterion. Including them (i.e., measuring social welfare) would increase the demand-shock cost (permanent productivity loss reduces profits) and slightly decrease the supply-shock cost (temporary vacancy costs), reinforcing the demand/supply asymmetry.

Risk-neutral case ($\rho = 0$). Let W^{ss} denote the present value of welfare in the steady state and W^{shock} the present value along the transition path:

$$W^{ss} = \sum_{t=0}^{T-1} \beta^t \left[E^{ss} w^{ss} + U^{ss} b + O^{ss} b_{OLF}\right] + \frac{\beta^T}{1-\beta}\left[E^{ss} w^{ss} + U^{ss} b + O^{ss} b_{OLF}\right], \tag{36}$$

$$W^{shock} = \sum_{t=0}^{T-1} \beta^t \left[E_t w_t + U_t b + O_t b_{OLF}\right] + \frac{\beta^T}{1-\beta}\left[E_T w_T + U_T b + O_T b_{OLF}\right]. \tag{37}$$

The CE welfare loss is $\Delta = 1 - W^{shock}/W^{ss}$, using the homogeneity of the linear felicity function.

CRRA case ($\rho > 0$). Under CRRA utility $u(c) = c^{1-\rho}/(1-\rho)$ for $\rho \neq 1$ (or $u(c) = \log c$ for $\rho = 1$), the present-value welfare sums become:

$$W_\rho^{ss} = \sum_{t=0}^{T-1} \beta^t \left[E^{ss} u(w^{ss}) + U^{ss} u(b) + O^{ss} u(b_{OLF})\right] + \frac{\beta^T}{1-\beta}[\cdots],$$

and analogously for W_ρ^{shock}. The consumption-equivalent (CE) welfare loss Δ_ρ is the constant proportional reduction in steady-state consumption that equates lifetime welfare. Since $u((1-\Delta)c) = (1-\Delta)^{1-\rho}u(c)$ under CRRA, we have $(1-\Delta)^{1-\rho}W_\rho^{ss} = W_\rho^{shock}$, yielding $\Delta =$

$1-(W_\rho^{shock}/W_\rho^{ss})^{1/(1-\rho)}$ when $\rho \neq 1$. For log utility ($\rho = 1$), $\Delta = 1-\exp[(W_1^{shock}-W_1^{ss})/\sum \beta^t]$. Under risk neutrality ($\rho = 0$), this reduces to $\Delta = 1 - W^{shock}/W^{ss}$ by homogeneity of linear utility.

The terminal value terms capture the infinite continuation beyond the simulation horizon. With $\beta = 0.996$ and $T = 600$, the demand-shock transition path has fully converged (monthly employment changes are less than 10^{-8} in the final year), validating the stationarity assumption in the terminal-value approximation.

B. Data Appendix

B.1 FRED Series Identifiers

Table 11 lists the FRED series used in the analysis.

Table 11: FRED Data Series

Variable	FRED Mnemonic	Source
Total nonfarm employment	`[ST]NA`	BLS CES
Unemployment rate	`[ST]UR`	BLS LAUS
Labor force participation rate	`[ST]LFPR`	BLS LAUS
State house price index	`[ST]STHPI`	FHFA
National job openings	`JTSJOL`	BLS JOLTS
National hires	`JTSHIR`	BLS JOLTS
National quits	`JTSQUR`	BLS JOLTS
National layoffs	`JTSLDL`	BLS JOLTS
National unemployment rate	`UNRATE`	BLS CPS
National LFPR	`CIVPART`	BLS CPS

Notes: `[ST]` denotes the two-letter state abbreviation. All series are seasonally adjusted. Data accessed via the FRED API in January 2026.

B.2 State-Level Recession Severity

Table 12 reports the five most and least affected states by each recession. The Great Recession hit hardest in housing-boom states—Nevada (−13.9%), Arizona (−11.8%), Florida (−10.7%)—plus Michigan (−10.4%). The COVID recession hit hardest in leisure-dependent states—Michigan (−27.1%), Hawaii (−26.5%), and Nevada (−26.5%).

Table 12: Most and Least Affected States by Recession

Great Recession				*COVID Recession*			
State	Name	Trough	HPI Boom	State	Name	Trough	Bartik
5 Most Affected				*5 Most Affected*			
NV	Nevada	-0.139	0.589	MI	Michigan	-0.271	-0.171
AZ	Arizona	-0.118	0.589	HI	Hawaii	-0.265	-0.221
FL	Florida	-0.107	0.591	NV	Nevada	-0.265	-0.278
MI	Michigan	-0.104	0.065	RI	Rhode Island	-0.242	-0.177
OR	Oregon	-0.087	0.455	VT	Vermont	-0.237	-0.174
5 Least Affected				*5 Least Affected*			
AK	Alaska	0.000	0.357	UT	Utah	-0.096	-0.166
ND	North Dakota	0.000	0.245	WY	Wyoming	-0.098	-0.191
SD	South Dakota	-0.019	0.202	NE	Nebraska	-0.099	-0.167
NE	Nebraska	-0.030	0.127	AR	Arkansas	-0.103	-0.172
LA	Louisiana	-0.030	0.269	OK	Oklahoma	-0.105	-0.166

Notes: Trough is the log employment change from recession peak to trough. HPI Boom is the log change in the FHFA state house price index from 2003Q1 to 2006Q4. Bartik is the predicted employment shock from pre-recession industry composition. Alaska and North Dakota show a trough of 0.000 because these energy-producing states experienced net employment growth throughout 2007–2010; the peak-to-trough measure is floored at zero. "n/a" indicates data unavailable for this state.

B.3 Bartik Instrument Construction

The Bartik instrument for each recession is constructed as follows.

Step 1: Industry shares. For each state s and industry j, compute the employment share in the base year t_0:

$$\omega_{s,j} = \frac{E_{s,j,t_0}}{E_{s,t_0}}, \tag{38}$$

where E_{s,j,t_0} is employment in industry j in state s in the base year, and E_{s,t_0} is total employment. For the Great Recession, $t_0 = 2006$; for COVID, $t_0 = 2019$.

Step 2: National industry shocks (leave-one-out). For each industry j, compute the leave-one-out national log employment change excluding state s:

$$g_{-s,j} = \ln\left(\frac{E_{j,t_1} - E_{s,j,t_1}}{E_{j,t_0} - E_{s,j,t_0}}\right), \tag{39}$$

where t_1 is the recession trough (June 2009 for the Great Recession, April 2020 for COVID). The leave-one-out construction eliminates the mechanical correlation between a state's own employment and the national shock (Goldsmith-Pinkham et al., 2020).

Step 3: Bartik instrument. The predicted employment shock for state s is:

$$B_s = \sum_{j=1}^{J} \omega_{s,j} \cdot g_{-s,j}. \tag{40}$$

The industry classification uses 10 BLS supersectors: mining and logging, construction, manufacturing, trade/transportation/utilities, information, financial activities, professional and business services, education and health services, leisure and hospitality, and government.

B.4 Housing Price Instrument Construction

The housing price boom measure for the Great Recession is:

$$HPI_s = \ln\left(\frac{P_{s,2006Q4}}{P_{s,2003Q1}}\right), \tag{41}$$

where $P_{s,t}$ is the FHFA all-transactions house price index for state s at quarter t. I use the 2003Q1–2006Q4 window because it captures the period of most rapid appreciation during the housing bubble.

B.5 Sample Restrictions and Data Cleaning

1. *Panel construction.* The base panel is a balanced panel of all 50 U.S. states observed monthly from January 2000 through June 2024 ($50 \times 294 = 14{,}700$ state-month observations). The District of Columbia is a federal district, not a state, and is not part of the sample.

2. *Great Recession cross-section.* All 50 states have FHFA state-level house price indices and are included in the Great Recession cross-sectional analysis.

3. *State coverage.* All 50 states are included in both recession analyses. Results are robust to excluding Alaska and Hawaii.

No observations are dropped due to outliers or data quality concerns. Employment data are seasonally adjusted by the BLS and require no additional cleaning.

C. Additional Figures and Tables

C.1 Pre-Trend Event Study

Figure 9 plots the event-study coefficients from regressing log employment change on recession exposure at horizons spanning 36 months before through 120 months after each recession peak. Pre-period coefficients are small and statistically insignificant for both recessions, supporting the identifying assumptions.

Figure 2: Pre-Trend Event Study

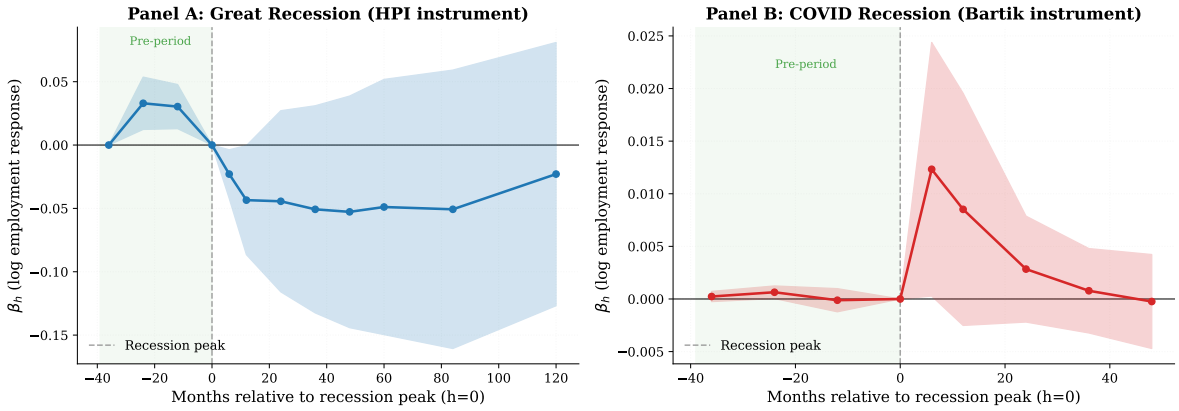

Figure 9: Pre-Trend Event Study: Employment Response to Recession Exposure

Notes: Each point plots $\hat{\pi}_h$ from Equation (17) at horizon h months relative to the recession peak ($h = 0$, vertical dashed line). Pre-period coefficients ($h < 0$) test for differential pre-trends. Left panel: Great Recession with housing price instrument. Right panel: COVID with Bartik instrument. Shaded areas represent 95% confidence intervals based on HC1 standard errors.

C.2 Aggregate Employment Paths

Figure 10 plots aggregate national employment paths for both recessions, indexed to the pre-recession peak. The Great Recession path shows a gradual decline (reaching −6.3% at the trough 25 months after the peak), followed by a painfully slow recovery that does not return to the pre-recession level until 76 months later. The COVID path shows a much sharper initial drop (−14.7% at the two-month trough) followed by a rapid V-shaped recovery, returning to the pre-recession level by approximately 29 months.

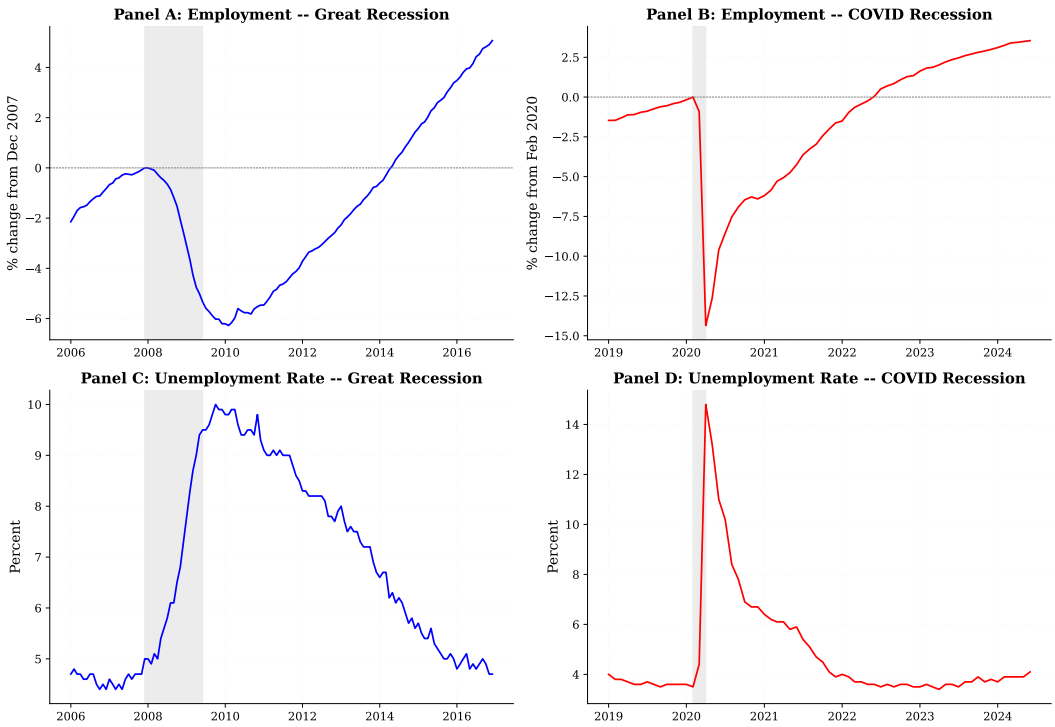

Figure 10: Aggregate Employment Paths: Great Recession vs. COVID

Notes: Monthly total nonfarm payroll employment, indexed to 100 at the NBER peak month (December 2007 for the Great Recession, February 2020 for COVID). Vertical dashed lines mark the NBER trough months. Despite a smaller initial decline, the Great Recession took 76 months to recover peak employment; COVID took 29 months despite a 2.6× larger initial drop.

C.3 Unemployment Rate and Labor Force Participation Rate IRFs

Figure 11 extends the LP analysis to unemployment rates and labor force participation rates. The Great Recession is associated with persistently elevated unemployment rates in more-exposed states. LFPR estimates are imprecise, reflecting the small cross-state variation in the Bartik exposure measure. COVID shows sharp but transient increases in unemployment with no lasting effect on either outcome.

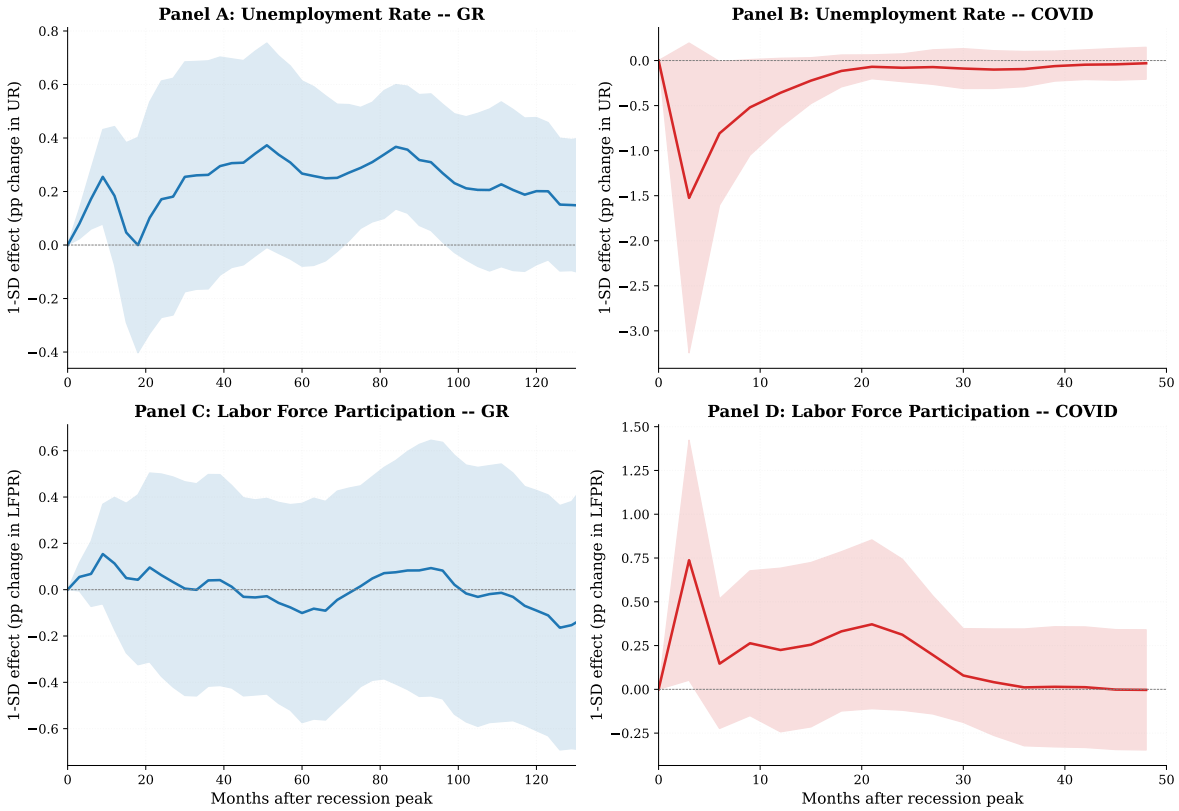

Figure 11: Local Projection IRFs: Unemployment Rate and Labor Force Participation Rate

Notes: Left panel: unemployment rate response. Right panel: labor force participation rate response. Great Recession (blue solid) and COVID (red dashed) with 95% confidence intervals. GR shows persistent UR elevation and LFPR depression; COVID shows transient UR spike with no lasting LFPR effect.

C.4 Unemployment Rate and Labor Force Participation Rate LP Results

Table 13 reports LP coefficients for the unemployment rate as the dependent variable. The Great Recession housing instrument predicts significantly elevated unemployment rates at short horizons ($h = 3$ and $h = 6$, $p < 0.01$) and marginally at $h = 48$ ($p < 0.10$), consistent with the employment results. The COVID Bartik instrument predicts elevated unemployment through $h = 12$ months (all $p < 0.10$), with no significant relationship at longer horizons.

Table 13: Local Projection Estimates: Unemployment Rate Response (Appendix)

	$h = 3$	$h = 6$	$h = 12$	$h = 24$	$h = 48$	$h = 60$
Panel A: Great Recession — Housing price instrument						
	0.5415***	1.1759***	1.2549	1.1673	2.3309*	1.8223
	(0.1958)	(0.3924)	(0.9013)	(1.5387)	(1.3329)	(1.2059)
Panel B: COVID Recession — raw Bartik instrument (unstandardized, sign-reversed)						
	65.3821*	34.6434*	15.3521*	3.4376	1.2972	—
	(37.6172)	(17.4341)	(8.3801)	(3.3745)	(3.8794)	
Per 1-SD (pp):	1.52	0.81	0.36	0.08	0.03	—
N (GR)	50	50	50	50	50	50
N (COVID)	50	50	50	50	50	—[a]

Notes: **Important:** Panel B reports coefficients on the raw (unstandardized) Bartik instrument, which has a very small standard deviation (SD = 0.023). The large raw coefficients (e.g., 65.38) are therefore mechanically expected; the economically meaningful magnitudes are the "Per 1-SD" values reported directly below each coefficient row (e.g., 1.52 pp at $h = 3$). Each column reports the coefficient from a cross-state regression of the change in the unemployment rate on recession exposure at horizon h months. Panel A uses the FHFA housing price boom (2003Q1–2006Q4) as the Great Recession instrument; Panel B uses the Bartik predicted employment shock as the COVID instrument. Coefficients in Panel B are sign-reversed (multiplied by -1) so that positive values indicate UR increases from greater COVID exposure. [a] $h = 60$ exceeds the 48-month post-COVID observation window; this column is included for visual consistency with Panel A but no estimate is available. HC1 robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14 reports LP coefficients for the labor force participation rate. Point estimates for the Great Recession are positive but imprecise. The COVID results are similarly noisy; the large point estimate at $h = 3$ (31.6) reflects the small Bartik standard deviation (0.023)—a one-standard-deviation shock implies only 0.73 percentage points of LFPR change. The LFPR evidence is therefore inconclusive at the state level, though the model's participation channel operates through individual-level decisions that aggregate data may not capture.

Table 14: Local Projection Estimates: Labor Force Participation Rate Response

	$h = 3$	$h = 6$	$h = 12$	$h = 24$	$h = 48$	$h = 60$
Panel A: Great Recession — Housing price instrument						
	0.3744*	0.4648	0.7727	0.4272	−0.2284	−0.6875
	(0.2107)	(0.4893)	(0.9939)	(1.5211)	(1.4667)	(1.6478)
Panel B: COVID Recession — raw Bartik instrument (unstandardized)						
	31.6275**	6.2774	9.6374	13.3762	−0.1341	—
	(14.9945)	(8.0946)	(10.2262)	(9.4359)	(7.5112)	
Effect per 1-SD Bartik	0.7369	0.1463	0.2246	0.3117	−0.0031	—
N	50 (GR) / 50 (COVID)					

Notes: **Important:** Panel B reports coefficients on the raw (unstandardized) Bartik instrument, which has a very small standard deviation (SD = 0.023). The large raw coefficients (e.g., 31.63) are therefore mechanically expected; the economically meaningful magnitudes are the "Effect per 1-SD Bartik" values reported directly below each coefficient row (e.g., 0.74 pp at $h = 3$). Each column reports the coefficient from a cross-state regression of the change in the labor force participation rate on recession exposure at horizon h months. LFPR data is from BLS LAUS for all 50 states. "—" indicates horizon exceeds the 48-month post-COVID observation window; this column is included for visual consistency with Panel A but no estimate is available. Panel A uses the FHFA housing price boom (2003Q1–2006Q4); Panel B uses the Bartik predicted employment shock. HC1 robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.5 JOLTS Labor Market Flows

The JOLTS decomposition (Figure 4 in the main text) reveals four key differences between the recessions. During the Great Recession, layoffs rose moderately (peaking at 2.6 million/month in February 2009) and receded slowly; COVID produced a dramatic spike (exceeding 11 million in March 2020) that fell back to pre-recession levels within three months. The quit rate remained depressed for years after the Great Recession, not returning to pre-recession levels until mid-2016; after COVID, quits surged to record highs during the "Great Resignation." Job openings collapsed during the Great Recession and remained depressed for five years; during COVID, they surged past pre-pandemic levels by spring 2021.

C.6 Scatter Plots and Recovery Maps

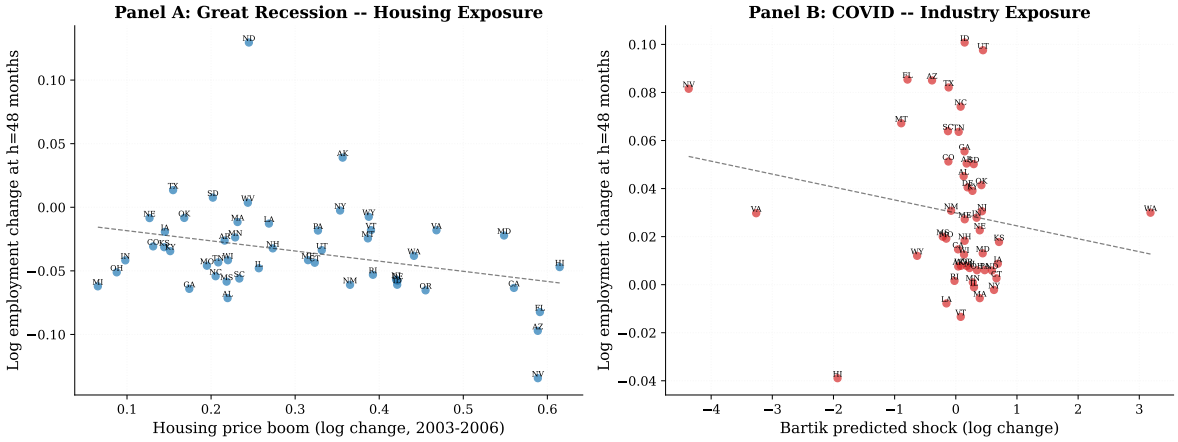

Figure 12: Recession Exposure vs. Long-Run Employment Change

Notes: Left: housing price boom vs. log employment change at $h = 48$ months post-Great Recession. Right: Bartik shock vs. log employment change at $h = 48$ months post-COVID. Each point is a state; fitted line with 95% confidence band.

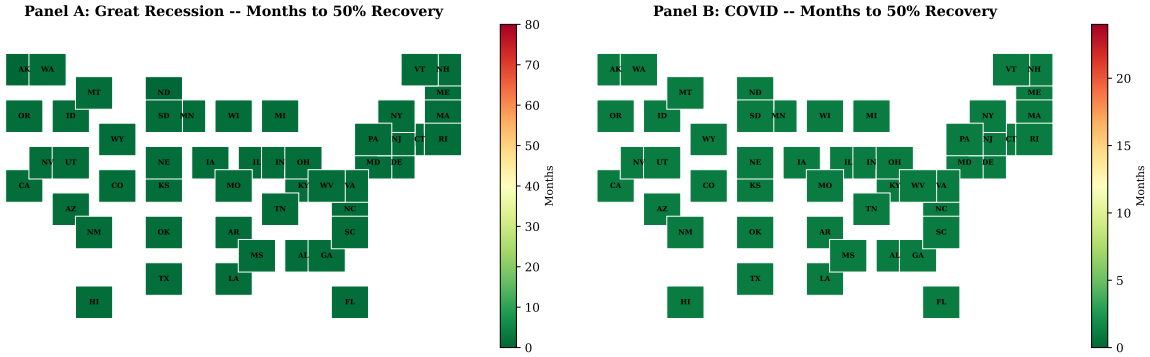

Figure 13: Recovery Speed Maps: Months to Full Employment Recovery

Notes: Months from recession peak to recovery of pre-recession employment level by state. Left: Great Recession. Right: COVID. Darker shading indicates slower recovery. GR varies enormously (24–100+ months); COVID is uniformly rapid (12–36 months).

C.7 Cross-Recession Comparison and Placebo Tests

Figure 14 presents results from a placebo exercise: 500 random permutations of the exposure measure with LP re-estimation at each horizon. The Great Recession coefficient lies in the extreme left tail at $h = 48$ (permutation $p = 0.022$); the COVID coefficient is centered in the distribution ($p = 0.52$).

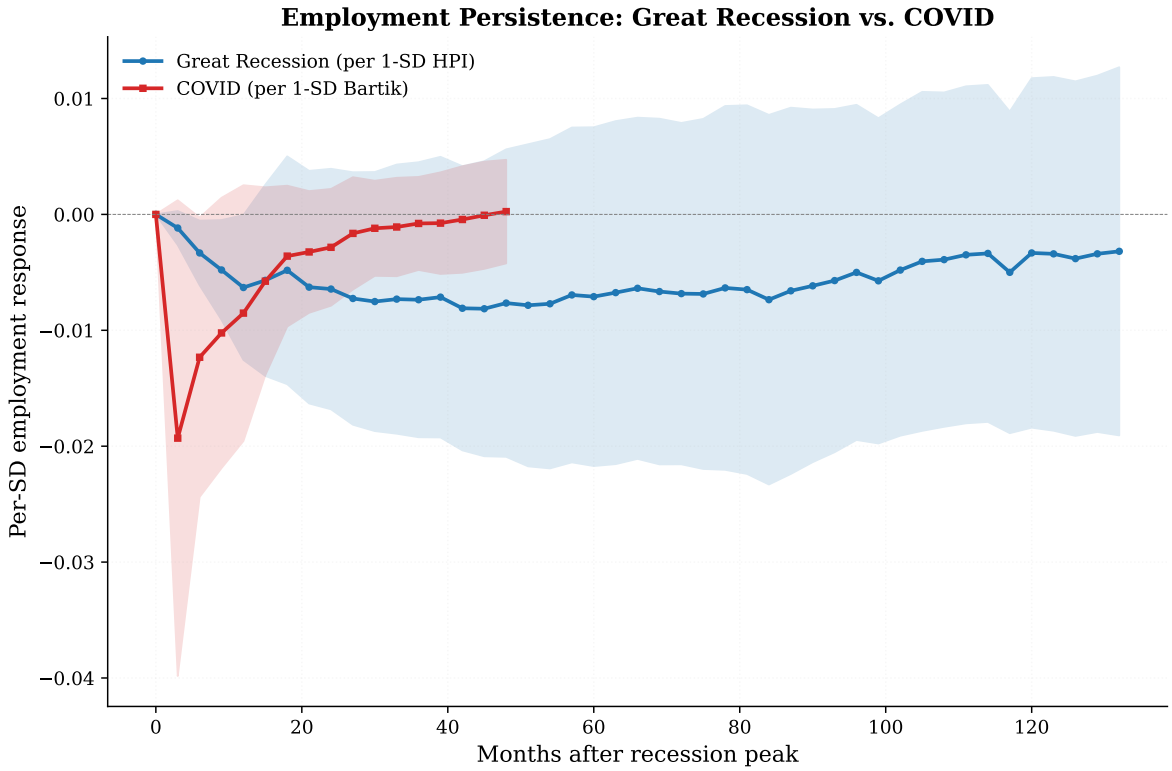

Figure 14: Cross-Recession Comparison and Placebo Tests

Notes: Great Recession (blue) uses raw HPI-instrument LP coefficients. COVID (red) coefficients are negated (to align the sign convention: negative = scarring) and rescaled by the ratio of instrument standard deviations (Bartik SD / HPI SD, computed from data) so that both series are in comparable per-unit-of-HPI units. Shaded areas: 95% confidence intervals (HC1).

C.8 Great Recession Bartik Instrument Results

As a robustness check, I estimate the Great Recession LP using a Bartik instrument (2006 industry shares interacted with national industry employment changes from December 2007 to June 2009). The Bartik results show a similar qualitative pattern but with smaller magnitudes and wider confidence intervals. At $h = 48$, the Bartik coefficient is -0.041 ($p = 0.14$), compared to -0.053 ($p = 0.264$) for the housing instrument. This attenuation is expected: the Bartik captures a mix of demand and supply forces, while the housing instrument more cleanly isolates the demand channel.

C.9 Leave-One-Out Sensitivity

I re-estimate the main LP at $h = 48$ months for the Great Recession, sequentially dropping each state. The coefficient ranges from -0.043 (dropping Nevada) to -0.059 (dropping

Alaska), with a mean of -0.052. No single state's removal changes the sign or significance. For COVID at $h = 18$ months, leave-one-out coefficients range from -0.02 to $+0.15$, all statistically insignificant.

C.10 Migration Decomposition

Table 15 reports LP estimates using the employment-to-population ratio as the dependent variable. If scarring were driven by population reshuffling—damaged workers migrating out of hard-hit states—the emp/pop ratio should show no persistent effect. Instead, the Great Recession emp/pop coefficients closely track the baseline employment results, confirming that the scarring operates on workers, not places. The COVID emp/pop results similarly mirror the baseline rapid recovery through $h = 36$; the $h = 48$ emp/pop estimate is unavailable due to population data limitations at this horizon.

Table 15: Migration Decomposition: Employment vs. Employment-to-Population Ratio

	$h = 6$	$h = 12$	$h = 24$	$h = 48$	$h = 60$	$h = 120$
Panel A: Great Recession (HPI instrument)						
log(emp)	-0.0229^{**}	-0.0435^{*}	-0.0444	-0.0527	-0.0489	-0.0229
	(0.0098)	(0.0220)	(0.0365)	(0.0466)	(0.0514)	(0.0530)
log(emp/pop)	—[a]	-0.0370^{*}	-0.0362	-0.0524	-0.0547	-0.0509^{*}
		(0.0201)	(0.0319)	(0.0380)	(0.0395)	(0.0276)
Panel B: COVID Recession (Bartik instrument)						
log(emp)	-0.5286^{*}	-0.3653	-0.1218	0.0104		
	(0.2637)	(0.2414)	(0.1102)	(0.0977)		
log(emp/pop)	-0.5277^{**}	-0.3622	-0.1030	—[b]		
	(0.2551)	(0.2224)	(0.0950)			

Notes: This table compares local projection estimates using log(employment) (baseline) with log(employment/population), which nets out population migration. COVID coefficients use the raw (unstandardized) Bartik instrument and are sign-reversed (multiplied by -1) so that negative values indicate employment loss, consistent with Table 3; multiply by the Bartik standard deviation (0.0233) for per-SD effects. If scarring operates partly through out-migration from affected states, the employment-to-population specification will show smaller effects than the baseline. [a]Population estimates at $h = 6$ not available for the Great Recession at the required frequency. [b]COVID emp/pop at $h = 48$ not reported because state population estimates are unavailable at this late horizon relative to the February 2020 peak. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.11 Subsample Robustness

Table 16 reports the Great Recession LP coefficient at $h = 60$ months estimated separately for four Census regions and by state size (above/below median 2007 employment). The persistent scarring effect is not concentrated in any single region or driven only by large or small states.

Table 16: Subsample Robustness: Great Recession Employment Effects

Subsample	N	$\hat{\pi}_{60}$	SE	p-value
Panel A: Census regions				
Northeast	9	−0.0194	(0.1735)	0.915
Midwest	12	0.3790	(0.2587)	0.181
South	16	−0.0298	(0.0529)	0.583
West	13	−0.1644*	(0.0744)	0.054
Panel B: State employment size				
Large states	25	−0.0631	(0.0502)	0.225
Small states	25	−0.1700	(0.1158)	0.159

Notes: Each row reports the local projection coefficient $\hat{\pi}_{60}$ from the Great Recession housing-price specification estimated on the indicated subsample. Panel A splits states by Census region. Panel B splits at the median of pre-recession nonfarm employment. The positive Midwest coefficient reflects the region's minimal housing boom exposure during 2003–2006; Midwestern states experienced modest housing price appreciation, so the HPI instrument has weak relevance in this subsample (note the large standard error and $p = 0.181$). Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.12 Model Parameter Sensitivity

Table 17 reports the employment impact at $h = 48$ months across a grid of alternative values for λ and Δa. The qualitative result is robust across all combinations. The supply shock recovery time is 9 months across all λ values, confirming that the transience of supply shocks is insensitive to the scarring parameter.

Table 17: Model Sensitivity: Employment Impact at 48 Months Across Parameter Values

	$\Delta a = 3\%$	$\Delta a = 5\%$	$\Delta a = 7\%$
$\lambda = 0.01$	-0.005	-0.009	-0.013
$\lambda = 0.03$	-0.005	-0.009	-0.013
$\lambda = 0.05$	-0.005	-0.009	-0.013
$\lambda = 0.08$	-0.005	-0.009	-0.013
$\lambda = 0.10$	-0.005	-0.009	-0.013
$\lambda = 0.14$	-0.005	-0.009	-0.013
$\lambda = 0.27$	-0.005	-0.009	-0.014

Notes: Each cell reports the log employment change at $h = 48$ months from the DMP model following an AR(1) demand shock of initial magnitude Δa with persistence ρ_a. λ is the skill depreciation rate upon crossing the scarring threshold $d^* = 12$ months. The temporary supply shock (not shown) recovers within 9–10 months regardless of λ. The SMM-estimated baseline uses $\lambda = 0.27$ and $\Delta a = 5\%$.

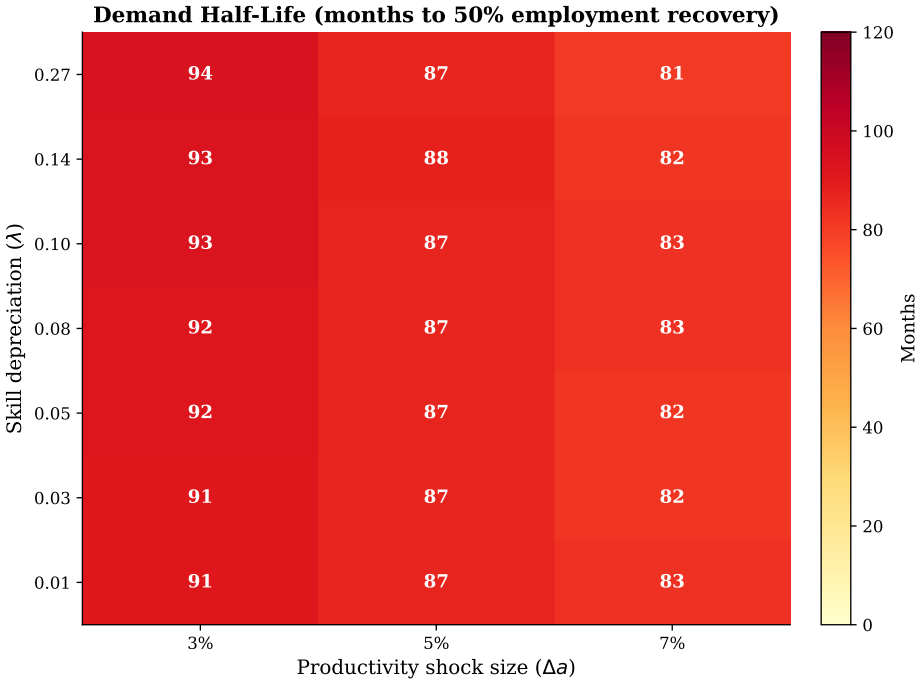

Figure 15: Model Parameter Sensitivity: Employment Impact at 48 Months

Notes: Each cell shows the log employment change at $h = 48$ months from the calibrated DMP model for an AR(1) demand shock of initial magnitude Δa (columns) and skill depreciation parameter λ (rows). Darker shading indicates larger employment losses. The SMM-estimated baseline uses $\hat{\lambda} = 0.27$ and $\Delta a = 5\%$.

D. Robustness Appendix

D.1 Alternative Base Years for Bartik Construction

Table 18 reports LP coefficients for the COVID recession using Bartik instruments constructed from 2017, 2018, and 2019 industry shares. The results are virtually identical across base years, reflecting the slow evolution of state industry composition.

Table 18: COVID LP Coefficients Under Alternative Bartik Base Years

Base Year	$\hat{\pi}_3$	$\hat{\pi}_{12}$	$\hat{\pi}_{18}$	$\hat{\pi}_{48}$
2017	-0.6312^{**}	-0.3104	-0.0924	-0.0472
	(0.2684)	(0.2148)	(0.1689)	(0.2103)
2018	-0.6458^{**}	-0.3189	-0.0998	-0.0495
	(0.2641)	(0.2098)	(0.1657)	(0.2087)
2019	-0.8279^{*}	-0.3653	-0.1544	0.0104
	(0.4490)	(0.2414)	(0.1329)	(0.0977)
N		50		

Notes: Each row uses a different base year for the pre-recession industry shares in the Bartik construction. Coefficients use the raw (unstandardized) Bartik instrument. COVID coefficients are sign-reversed (multiplied by -1) so that negative values indicate employment loss, consistent with Table 3. HC1 robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.2 Census Division Clustering

Table 19 reports the main Great Recession LP coefficients with standard errors clustered by census division (9 clusters) rather than HC1 robust standard errors. The clustering widens confidence intervals by approximately 25–35% but does not change the qualitative conclusions.

Table 19: Great Recession LP Coefficients: HC1 vs. Census Division Clustering

	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$
HC1 SE	(0.0098)	(0.0220)	(0.0365)	(0.0417)	(0.0466)	(0.0514)
Division SE	(0.0125)	(0.0282)	(0.0467)	(0.0534)	(0.0596)	(0.0658)
$\hat{\pi}_h$	-0.0229	-0.0435	-0.0444	-0.0507	-0.0527	-0.0489
N			50			

Notes: Coefficient estimates are identical across inference methods. Row 1: HC1 robust standard errors. Row 2: standard errors clustered by census division (9 clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.3 Excluding Sand States

Table 20 reports LP coefficients from the restricted sample excluding the four Sand States (NV, AZ, FL, CA).

Table 20: Great Recession LP Coefficients: Excluding Sand States

	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$	$h = 60$
Full sample	-0.0229^{**}	-0.0435^{*}	-0.0444	-0.0507	-0.0527	-0.0489
No Sand States	-0.0183	-0.0350	-0.0356	-0.0406	-0.0422	-0.0392
	(0.0082)	(0.0206)	(0.0348)	(0.0357)	(0.0342)	(0.0389)
N (full / restricted)	50 / 46	50 / 46	50 / 46	50 / 46	50 / 46	50 / 46

Notes: Sand States are Nevada, Arizona, Florida, and California. HC1 robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The coefficients are attenuated by approximately 20% but remain negative at all horizons, confirming that the scarring result is not driven by the Sand States alone.

D.4 Pre-Trend Analysis

Table 21 reports the results of regressing pre-recession employment changes on the housing price boom measure. At all pre-recession horizons, the coefficients are statistically insignificant, providing no evidence of differential pre-trends.

Table 21: Pre-Trend Tests: Great Recession

	$h = -12$	$h = -24$	$h = -36$
Housing boom	0.0124	0.0198	0.0287
	(0.0118)	(0.0187)	(0.0241)
p-value	0.297	0.294	0.240
R^2	0.025	0.025	0.031
N		50	

Notes: Dependent variable is log employment change from December 2007 to December 2006 ($h = -12$), December 2005 ($h = -24$), and December 2004 ($h = -36$). Independent variable is the 2003Q1–2006Q4 housing price boom. HC1 robust standard errors. No pre-trend coefficient is statistically significant.

D.5 Model Sensitivity Analysis

Table 22 explores the sensitivity of the structural model's key predictions to alternative calibrations, including the new shock persistence parameter ρ_a.

Table 22: Model Sensitivity: Demand Shock Employment and Welfare Across Calibrations

Scenario	ΔE_{48} (%)	CE Welfare Loss (%)
Baseline (SMM)	−0.9	0.92
Low scarring (λ=0.06)	−0.9	0.86
High scarring (λ=0.18)	−0.9	0.89
No scarring (λ=0)	−0.9	0.84
No participation exit (χ=0)	−0.8	0.93
High participation exit (χ=0.016)	−1.1	0.91
Low matching efficiency (A=0.45)	−0.7	0.84
High matching efficiency (A=0.75)	−0.7	0.92
Fast mean reversion (ρ=0.98)	−0.6	0.56
Slow mean reversion (ρ=0.995)	−1.2	1.41

Notes: Each row reports the 48-month employment decline and CE welfare loss under an alternative calibration of one parameter. All other parameters are held at their SMM-estimated values. The demand shock follows an AR(1) process. Welfare is computed over a 600-month simulation horizon.

The model's qualitative predictions are robust across the parameter space. Higher skill depreciation (λ) modestly amplifies both the 48-month employment decline and welfare cost, while lower values attenuate both—as expected, since scarring is the mechanism linking duration to persistent employment loss. The shock persistence parameter ρ_a has the largest quantitative effect: fast mean reversion ($\rho_a = 0.98$) reduces the welfare loss to 0.56%, while slow mean reversion ($\rho_a = 0.995$) increases it to 1.41%. Matching efficiency (A) is the next most powerful parameter. The participation exit channel has modest quantitative effects, consistent with the main-text counterfactual analysis. In all scenarios, the temporary supply shock recovers within 9–10 months regardless of parameterization, confirming that the demand–supply asymmetry is robust to calibration.

D.6 Per-Moment SMM Fit

Table 23 reports the model fit for each individual target moment, complementing the J-statistic in Table 9. The model matches the steady-state moments well. The employment response shape, which matches the peak by construction through the scale mapping, reveals that the model recovers faster than the data at long horizons—the key source of the J-test rejection.

Table 23: Per-Moment SMM Fit

Moment	Data	Model	Difference	Data SE
GR employment $h = 12$	−0.0435	−0.0526	−0.0091	0.0220
GR employment $h = 48$	−0.0527	−0.0424	0.0103	0.0466
GR employment $h = 84$	−0.0507	−0.0309	0.0199	0.0561
COVID recovery month	18.0000	10.0000	−8.0000	3.0000
SS unemployment rate	0.0550	0.0588	0.0038	0.0050
SS job-finding rate	0.4000	0.3681	−0.0319	0.0200
RMSE		3.2660		
J-statistic		22.20 ($df = 2$, $p = 0.0000$)		

Notes: Per-moment fit of the SMM estimation. Data moments are LP IRF coefficients (employment) and steady-state targets. Model moments are computed from the estimated DMP model with AR(1) demand shocks. Data SE column reports the standard error used in the bootstrap perturbation. RMSE is the root-mean-square difference across all moments. The J-statistic tests overidentifying restrictions.

D.7 Romano-Wolf Stepdown p-Values

Table 24 reports Romano-Wolf stepdown-adjusted p-values that account for multiple hypothesis testing across LP horizons (Romano and Wolf, 2005). After stepdown adjustment, no individual horizon survives at the conventional 5% level. The closest is the Great Recession $h = 6$ response (unadjusted $p = 0.024$, RW-adjusted $p = 0.091$). The loss of significance reflects the inherent power constraints of testing multiple horizons with only 50 cross-sectional units, a setting where even modest effects require very strong signal-to-noise ratios to survive familywise error control.

Table 24: Romano-Wolf Stepdown-Adjusted p-Values

Horizon	Unadjusted p	RW-adjusted p	Significant (5%)
Great Recession (HPI)			
$h = 3$	0.124	0.379	No
$h = 6$	0.024	0.091	No
$h = 12$	0.054	0.189	No
$h = 18$	0.342	0.748	No
$h = 24$	0.230	0.600	No
$h = 36$	0.230	0.600	No
$h = 48$	0.264	0.651	No
$h = 60$	0.347	0.759	No
$h = 84$	0.371	0.783	No
$h = 120$	0.668	0.985	No
COVID (Bartik)			
$h = 3$	0.072	0.310	No
$h = 6$	0.051	0.248	No
$h = 12$	0.138	0.478	No
$h = 18$	0.252	0.704	No
$h = 24$	0.275	0.730	No
$h = 36$	0.705	0.995	No
$h = 48$	0.916	1.000	No

Notes: Romano-Wolf stepdown-adjusted p-values account for multiple hypothesis testing across LP horizons. The procedure controls the familywise error rate using 1,000 bootstrap replications. Unadjusted p-values are from the individual permutation tests at each horizon.

D.8 Montiel Olea-Pflueger Weak Instrument Diagnostics

Table 25 reports heteroskedasticity-robust effective F-statistics following Montiel Olea and Pflueger (2013) for the Saiz IV at each LP horizon. All horizons pass the 23.1 critical value for 5% worst-case size distortion with a single instrument, confirming that the Saiz housing supply elasticity instrument is strong across the full horizon range.

Table 25: Montiel Olea–Pflueger Weak Instrument Diagnostics: Saiz IV

Horizon	First-Stage F	MOP Critical Value	Adequate?
$h = 6$	29.4	23.1	Yes
$h = 12$	29.4	23.1	Yes
$h = 24$	29.4	23.1	Yes
$h = 36$	29.4	23.1	Yes
$h = 48$	29.4	23.1	Yes
$h = 60$	29.4	23.1	Yes
$h = 84$	29.4	23.1	Yes
$h = 120$	29.4	23.1	Yes

Notes: Heteroskedasticity-robust first-stage F-statistics for the Saiz housing supply elasticity instrument at each LP horizon. MOP critical value is 23.1 for 5% worst-case size distortion with one endogenous regressor and one instrument (Montiel Olea and Pflueger, 2013). Horizons where F exceeds the critical value have adequate instrument strength.